



SAPIENZA
UNIVERSITÀ DI ROMA

Sapienza University of Rome

Department of Computer, Control and Management Engineering
PhD in Data Science

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Coupling Processes and Travelling Ionospheric Disturbances

Detection, Characterization and Forecasting with GNSS
Observations

Thesis Advisor
Dr. Claudio Cesaroni

Candidate
Marco Guerra
2067098

Academic Year 2025-2026

*"You miss 100% of the shots you
don't take – Wayne Gretzky"*
– Michael Scott

Abstract

The ionosphere is a dynamic region where energy and momentum from the Sun, magnetosphere, and lower atmosphere converge, generating complex plasma variability that affects global communication and navigation systems. Among the most significant manifestations of this variability are Traveling Ionospheric Disturbances (TIDs), large-scale wave-like perturbations in electron density capable of propagating thousands of kilometers. Despite decades of research, the mechanisms governing their generation and evolution remain only partially understood due to the multiplicity of coupling processes and observational constraints.

This thesis develops a reproducible, interpretable, and data-driven framework for detecting and forecasting Large-Scale TIDs (LSTIDs) using Global Navigation Satellite System (GNSS) observations. The work is structured as a progressive methodological pipeline: the quantification of measurement and processing uncertainties, the development of an autonomous LSTID detection algorithm based on GNSS keograms and image-processing techniques, and the creation of probabilistic forecasting models linking ionospheric disturbances to solar and geomagnetic drivers. These models are optimized via ensemble learning, calibrated for reliability, and interpreted through explainable AI techniques to trace predictions back to their physical causes.

Through the integration of empirical analysis, physical interpretation, and machine learning, this thesis advances the capability to monitor and predict ionospheric disturbances in near real time. As a result, it bridges observational ionospheric research with operational space-weather forecasting, contributing to both fundamental understanding and applied geospace monitoring.

Contents

Introduction	6
1. Scientific Background	8
1.1 Ionosphere.....	8
1.1.1 Ionospheric Structure and Regional Classification	8
1.2 Drivers of Ionospheric Variability	10
1.2.1 Quiet-Time Variability	10
1.2.2 Disturbed Ionospheric Conditions.....	12
1.3 Vertical coupling in the atmosphere-ionosphere system	14
1.4 Radio propagation in the ionosphere	17
1.4.1 HF Sounding of the ionosphere.....	18
1.4.2 GNSS and Total Electron Content	19
1.5 Travelling ionospheric disturbances	23
1.5.1 Forcing from above	23
1.5.2 Forcing from below	24
1.5.3 TID Detection.....	26
2. Examples of coupling processes	27
2.1 Hunga Tonga Volcanic Eruption	27
2.1.1 Data and methods	28
2.1.2 Results and discussion.....	29
2.2 Türkiye Earthquake.....	31
2.2.1 Data and methods	33
2.2.2 Results and discussion.....	33
2.3 Severe Convection Linked to Jet-Front over Europe.....	39
2.3.1 Data and Methods.....	39
2.3.2 Results and discussion.....	40
3. Errors in TID detection	43
3.1 The synthetic arc database	43
3.2 Detrending techniques comparison.....	45
3.3 Elevation cutoff impact.....	48
3.4 Near-Real-Time Observables.....	50
4. LSTID Nowcasting	54
4.1 The detection process layout.....	55
4.2 Data gathering and station selection with K-Means	56
4.3 Keogram generation.....	57
4.4 LSTID autonomous identification algorithm.....	58
4.5 Validation on case events.....	61
4.6 Statistical analysis	65

5. LSTID Forecasting	71
5.1 From LSTID catalogues to target variables	71
5.1.1 T-FORS Catalogue	71
5.1.2 ARISTIDE Catalogue	72
5.2 From physical drivers to input features	73
5.2.1 T-FORS input features	73
5.2.2 ARISTIDE input features	74
5.3 The Machine Learning Models	77
5.3.1 T-FORS – CatBoost	78
5.3.2 ARISTIDE – QUANT	80
5.4 Results and comparisons	82
5.4.1a T-FORS – Hyperparameter Tuning	83
5.4.1a T-FORS – Metrics, Explainability, and Calibration	83
5.4.2a ARISTIDE – Hyperparameter Tuning	87
5.4.2b ARISTIDE – Metrics, Explainability, and Calibration	87
5.4.3 Model comparison on case event	94
Conclusions and future work	100
Bibliography	104
Appendix A	115
A.1 Ionosonde Measurements for the Türkiye Earthquake	115
A.2 Additional results for the August 2023 tropospheric storm	118
A.2.1 GNSS receivers	118
A.2.2 Doppler sounder	119
A.2.3 Ionosondes	120
A.2.3 Atmospheric weather system	121
Appendix B	122
B.1 Savitzky-Golay and polinomial detrending parameters	122
B.1.1 Polynomial Detrending	122
B.1.2 Savitzky–Golay Filtering	123
Appendix C	125
C.1 Impact of GNSS network size on keogram IPP per pixel	125
C.2 Investigating the amplitude threshold impact	127
C.3 Investigating the minimum solidity impact	131
C.4 Investigating the maximum time separation impact	134
C.5 Investigating the max area of the largest concavity impact	136

Introduction

The ionosphere is a dynamic plasma environment where solar, magnetospheric, and atmospheric energy inputs converge and interact. Its complex electrodynamics governs the propagation of electromagnetic waves including GNSS signals and thus plays a critical role in global communication and navigation systems. During enhanced solar or geomagnetic activity periods, a wide range of phenomena affects the spatial and temporal distribution of the ionospheric plasma. Among these, travelling ionospheric disturbances (TIDs) stand out as coherent wave-like perturbations capable of propagating across thousands of kilometres, altering the ionospheric total electron content (TEC) and degrading the performance of trans-ionospheric systems.

Despite their operational importance and decades of observational research, the mechanisms governing TID generation and propagation remain only partially understood. Multiple coupling mechanisms modulate their occurrence, from magnetospheric forcing and auroral heating to atmospheric wave dissipation and background ionospheric conditions, each characterized by different spatial and temporal scales. This complexity has historically limited our capacity to detect and forecast TIDs systematically. Traditional analyses, often based on isolated events or limited regional coverage, have provided valuable qualitative insights but lack the reproducibility and automation required for operational monitoring.

The central motivation of this thesis lies in addressing this gap: to develop a reproducible, interpretable, and data-driven framework for detecting and forecasting large-scale TIDs (LSTIDs). The research combines physical modelling and statistical learning to describe the genesis and propagation of LSTIDs as an ionospheric response to solar and geomagnetic forcing, using GNSS-based TEC observations and ionosonde measurements as core inputs. It was conducted within the framework of two research projects: T-FORS (funded by the European Union - GA-101081835) and ARISTIDE (Programma regionale FESR Lazio 2021-2027). Both of them contributed to a specific phase of the methodological pipeline, from event detection to short-term forecasting. The work combines several peer-reviewed publications I authored and co-authored, and original analyses, integrating them into a scientific pipeline that links empirical observations to physical interpretation.

The thesis opens with the theoretical and instrumental background. Chapter 1 reviews the fundamental processes governing the ionosphere, emphasizing its role within the Sun-Earth system. It discusses the main mechanisms of solar wind-magnetosphere coupling, geomagnetic storms, and atmospheric wave coupling, which generate or modulate TIDs. This theoretical overview integrates the study within the broader context of space-weather research, where understanding ionospheric variability is essential for scientific and operational purposes.

Chapter 2 presents a series of representative case studies that illustrate the diversity of ionospheric disturbances and provide the foundation for the methodological developments that follow. Each event combines GNSS-derived TEC observations with complementary instruments such as ionosondes to capture how different coupling mechanisms manifest as TIDs. These case analyses, to which I contributed, all published in peer-reviewed journals (Barta et al., 2025; Haralambous et al., 2023; Verhulst et al., 2022), demonstrate the typical morphology and variability of TIDs, highlight instrumental limitations, and present the physical parameters that most reliably describe the ionospheric response. In doing so, Chapter 2 transforms isolated observational events into a structured knowledge base that constrains the design of the detrending experiments in Chapter 3 and guides the feature extraction logic of the autonomous detector in Chapter 4.

Building on this foundation, Chapter 3 explores how methodological uncertainties impact the detection of ionospheric disturbances. Using synthetic TEC data, it investigates the performance of different detrending strategies and quantifies their impact on the detectability and amplitude of TID signatures. Synthetic arcs allow the construction of an observable that includes a known modeled disturbance, allowing a rigorous comparison of filtering choices, near real-time observables, and observational geometry and their influence on the retrieved amplitude and period. This analysis provided clear guidelines for designing subsequent detection algorithms. The results, published in Guerra et al. (2024), represent the first step of the main thesis outcome: before detecting or forecasting, one must first understand how the data processing distorts or reveals the physical process of interest.

Chapter 4, based on Guerra et al. (2025), translates this methodological groundwork into an autonomous detection system for LSTID events. Using GNSS keograms (two-dimensional latitude–time representations of TEC deviations from the background variability), the developed algorithm combines spectral analysis and image-processing techniques to identify the spatio-temporal patterns associated with LTID propagation. This framework transforms dense GNSS networks into an operational monitoring system capable of detecting disturbances automatically and consistently. The resulting multi-year catalog of LSTID occurrences is the reference dataset for later forecasting efforts. Validation was performed by studying the climatological behavior of LSTID with respect to known relationships with the geomagnetic and solar conditions. The methods and results presented in this chapter have been published in Guerra et al. (2025), covering the full detection methodology, an example of detection for different areas and case events, and a climatological study to assess the physical soundness of the results. All of the previous are presented in this thesis.

The final part of the thesis, Chapter 5, focuses on the transition from detection to forecasting. Here, two complementary models are introduced: T-FORS, which is based on an independent catalog of LSTID derived from HF measurements, and ARISTIDE, which is based on the catalog obtained with the algorithm presented in Chapter 4. T-FORS is designed to forecast LSTID by leveraging geomagnetic and near-real-time LSTID activity indices, such as the regional HF-INTEU, along with disturbance velocity and propagation direction derived from single ionosondes. In contrast, ARISTIDE addresses the larger-scale geomagnetic conditions, where the orientation of the interplanetary magnetic field, the strength of the auroral electrojets, and the level of solar activity modulate the probability of LSTID formation. Both models rely on ensemble learning architectures optimized through Bayesian parameter search and calibrated to produce probabilistic forecasts. Explainable artificial intelligence methods are applied to trace each prediction back to its physical drivers, ensuring transparency and interpretability. The development, optimization, and calibration of T-FORS is based on Ventriglia et al. (2025), while the ARISTIDE model is presented here for the first time.

Together, these models embody the dual nature of ionospheric predictability: one anchored to immediate observational response, the other reflecting the slower magnetospheric coupling that precedes disturbance onset. Their integration marks the completion of the thesis work, which started from quantifying measurement errors, moved to autonomous detection, and finally leveraged the previous development to produce short-term probabilistic forecasting.

This thesis consolidates three years of research accomplished in different research collaborations into a manuscript that connects fundamental understanding, methodological innovation, and operational application. It represents a continuous progression: from quantifying the errors that limit the measurements, through designing algorithms that autonomously detect disturbances, to developing interpretable models capable of forecasting them. In doing so, I contributed to the scientific study of ionospheric coupling and the technological evolution of space-weather forecasting systems, bridging the gap between physical insight and operational capability.

Chapter 1

Scientific Background

1.1 Ionosphere

The ionosphere is a region of the Earth's upper atmosphere, extending roughly from 60 km to over 1,000 km in altitude, where solar radiation ionizes atmospheric particles, creating a weakly ionized cold plasma composed of free electrons and ions. This region plays a vital role in space weather, radio communication, and satellite-based navigation systems, as it significantly affects the propagation of electromagnetic waves. Structurally, the ionosphere is shaped by complex interactions between solar radiation, geomagnetic forces, and atmospheric dynamics, leading to distinct spatial and temporal variations in its electron density. The layered nature of the ionosphere reflects the interplay between ionization sources, recombination, and transport processes at different altitudes. In addition to this altitude-based stratification, the ionosphere exhibits significant horizontal variability across geomagnetic latitudes, with distinct features present at equatorial, mid-latitude, and polar regions and in daytime and nighttime regions. Understanding this layered and dynamic structure is essential for characterizing ionospheric variability under quiet and disturbed geophysical conditions.

1.1.1 Ionospheric Structure and Regional Classification

The ionosphere is vertically structured into distinct regions, namely the D, E, and F layers, characterized by varying electron density profiles as a function of altitude. Figure 1.1 shows the different ionospheric layers and the height-dependent temperature and density for the different ion populations constituting the ionosphere. These layers were initially identified through radio sounding techniques and were later confirmed by satellite measurements. Each layer is defined by a different altitude range and ion composition due to different solar irradiance and atmospheric conditions (King, 1956; Nicholls, 1993; Mathews et al., 1997).

The D-region occupies altitudes approximately between 60 and 90 km and represents the lowest and least ionized ionospheric layer. Ionization within this region predominantly arises from solar Lyman-alpha emissions and soft X-ray radiation during daytime hours. Due to the high neutral particle density at these altitudes, the ionization is highly collisional and short-lived. Consequently, the D-region significantly absorbs low-frequency (LF and VLF) radio waves and largely dissipates during nighttime when the ionizing radiation ceases (Barr, 1982).

Above the D-region lies the E-region, extending from approximately 90 to 150 km altitude. Ionization within the E-region is primarily driven by solar soft X-ray and far ultraviolet (FUV) radiation, with NO^+ and O_2^+ ions dominating its composition. This region exhibits pronounced diurnal variation, typically achieving peak densities around noon. Notably, the E-region also hosts sporadic-E layers, transient enhancements in ionization that reflect high-frequency radio signals. These layers result from wind shears, metallic ion deposition from meteoric ablation, and electrojet currents prevalent at low and high latitudes.

The F-region is the highest and most ionized portion of the ionosphere, extending roughly from 150 km to beyond 1,000 km in altitude. During daytime conditions, it can be subdivided

into two sub-layers: the F1-layer, located between 150 and 200 km, and the F2-layer, typically observed between 250 and 400 km altitude. The F2-layer exhibits the highest electron densities and is mainly composed of atomic oxygen ions (O^+). Unlike the D- and E-regions, whose electron density profiles can be approximated by a simple Chapman layer determined primarily by local ionization and recombination processes, the F-region is strongly influenced by transport mechanisms, such as diffusion along magnetic field lines and large-scale thermospheric winds. As a result, its electron density distribution cannot be accurately described by a purely photochemical equilibrium. Instead, it reflects a complex balance between production, loss, and plasma transport, making the F-region highly dynamic and sensitive to variations in solar EUV flux, geomagnetic activity, and neutral atmosphere circulation.




Figure 1.1: Altitudinal distribution of ionospheric plasma, ion populations and ion/electron temperature retrieved from the IRI-2012 model for local noon (03 UT) during the spring equinox over Japan during solar minimum. Daytime layers, along with their names (D, E, F1, and F2), are also shown. Taken from Shiokawa (2023)

In addition to altitude-based layering, the ionosphere can be classified according to magnetic latitude into distinct regions in both the Southern and Northern hemispheres. These include the low latitude region, the mid-latitude region, and the high latitude region. The low latitude region extends approximately 20° to 30° from the magnetic equator, while the mid-latitude region spans magnetic latitudes from 30° to 60° . The high latitude or polar region is further subdivided into the auroral region, approximately from 60° to 70° , and the polar cap region from 70° to 90° magnetic latitude.

1.2 Drivers of Ionospheric Variability

The ionosphere, being at the frontier between Earth and space, is extremely dynamic and exhibits variability across a wide range of temporal and spatial scales due to forcing from below and above. These variations are broadly categorized into quiet-time and disturbed conditions, each driven by distinct physical mechanisms (Paul, 1995; Kouris & Fotiadis, 2002; Zolesi & Cander, 2014). Photoionization by solar radiation and the morphology of the geomagnetic field are the dominant drivers under quiet conditions, governing the large-scale structure and variability of the ionosphere. During disturbed periods, however, the ionospheric state is additionally influenced by both enhanced coupling between the Earth's subsystems (lithosphere, atmosphere, ionosphere, magnetosphere, and plasmasphere) and direct solar-terrestrial interactions, such as variations in solar-wind pressure, interplanetary magnetic field orientation, coronal mass ejections, and solar flares. These processes inject energy and momentum into the coupled system, driving large-scale ionospheric disturbances and altering its electrodynamic balance.

1.2.1 Quiet-Time Variability

Quiet-time ionospheric variability arises from processes that act over distinct temporal scales, ranging from minutes to years. These processes are largely governed by solar radiation and atmospheric dynamics under geomagnetically undisturbed conditions, and they manifest as structured variations in electron density across latitude, longitude, altitude, and local time. While the ionosphere is often described in equilibrium photoionization, its behavior is modulated by a combination of solar input, thermospheric winds, composition changes, and transport mechanisms.

On diurnal timescales, ionospheric variability is dominated by changes in solar zenith angle (χ), which affects the photoionization rate. This response generally conforms to Chapman's theory at low and mid-latitudes, where the ion production rate increases as the Sun rises, peaking near local noon, and then decreases post-sunset due to recombination. (Munro, 1962; Aravindan & Iyer, 1990; Ezquer et al., 2004). Consequently, the vertically integrated number of electrons in the ionosphere, named Total Electron Content (TEC), follows a daily cycle, with a rapid morning buildup, a mid-day maximum, and a slower decay through the evening. Figure 1.2 shows the global TEC distribution for different local times and solar activity conditions. This pattern varies with latitude: equatorial and low-latitude regions receive more direct sunlight and exhibit stronger daytime TEC peaks than higher latitudes. (Soicher & Gorman, 1985; Gail et al., 1993; Wilkinson et al., 2001)

In addition to daily changes in electron density, diurnal variability also drives the formation of persistent spatial structures. One of the most prominent is the Equatorial Ionization Anomaly (EIA), a band of enhanced electron density that forms symmetrically around the geomagnetic equator ($\sim \pm 15^\circ$). The EIA, clearly identifiable in Figure 1.2, is caused by the so-called “fountain effect”, whereby ionized plasma is lifted vertically by the combined effect of the eastward electric field and the meridional geomagnetic field ($E \times B$ drift) and then redistributed poleward and downward along magnetic field lines. The EIA is most pronounced during daytime and contributes to sharp latitudinal gradients in electron density across equatorial regions. (Stolle et al., 2008; Paul & DasGupta, 2010). In contrast, the mid-latitude is characterized by the nighttime trough, a recurring depletion of electron density between roughly 55° and 70° geomagnetic latitude, which arises from the interplay of recombination and poleward transport of plasma. The trough forms in the absence of strong auroral activity and serves as an indicator of background thermospheric circulation (Scali, 1992; He et al., 2011).

On seasonal timescales, ionospheric electron density varies with Earth's tilt and orbit, modulating solar insolation in each hemisphere. At midlatitudes, increased sunlight during summer typically leads to higher ionization levels, while winter is characterized by reduced solar input and thus lower densities (Mansilla et al., 2005; Liu et al., 2009). Nevertheless, the F region does not always follow this expected pattern. Several anomalies characterize its behavior:

- The winter anomaly refers to the counterintuitive observation that F2-layer peak densities (NmF2) are often higher in the winter hemisphere than in the summer during solstice periods. This anomaly is mainly caused by lower recombination rates in the colder and denser winter thermospheres (Torr et al., 1980; Lee et al., 2011).
- The annual anomaly describes the global tendency for higher NmF2 values during the December solstice than during the June solstice (Deminov et al., 2021).
- The semiannual anomaly is characterized by F2-layer density peaks occurring near the equinoxes (March and September) rather than during the summer solstice. This anomaly is caused by the enhanced ion–neutral coupling efficiency during equinoxes, as well as globally symmetric wind and wave conditions that uplift plasma to higher altitudes, where recombination rates are lower due to a less dense neutral atmosphere (Rishbeth et al., 2000; Pavlov & Pavlova, 2017)

These anomalies highlight the importance of non-photochemical processes in modulating seasonal ionospheric variability, particularly in the F-region.




Figure 1.2: Distribution of total electron density for different levels of solar activity at 03:00 and 12:00 UT. The bottom panel shows the solar sunspot number in grey and the F10.7 index in red. The vertical lines correspond to the days for which the ionospheric maps are plotted and color-coded according to the map titles. Figure taken from Maletcki (2023).

On solar-cycle timescales, the ionosphere reflects the approximately 11-year modulation of solar EUV output associated with the sunspot cycle. The F10.7 solar flux index, which tracks solar radio emissions at 10.7 cm, is a reliable proxy for EUV variability and correlates strongly with global and regional TEC. The three rows of Figure 1.2 represent the global TEC distribution for three different levels of solar activity, represented by the F10.7 index shown in the bottom panel of Figure 1.2. During solar maxima, TEC almost doubles with respect to solar minimum, resulting in higher plasma densities, increased ionospheric delay for GNSS signals, and stronger coupling with the magnetosphere (Balan et al., 1996; Balan et al., 1993; Vaishnav et al., 2019). Observational records demonstrate that global TEC and sunspot numbers evolve accordingly and that this relationship exists across geographic regions. The correlation is strong enough that F10.7 is routinely used as an input to empirical ionospheric models such as IRI and NeQuick. These long-term modulations influence baseline ionospheric behavior and are essential for forecasting background electron density conditions. (Smith & King, 1981; Hall et al., 2020)

1.2.2 Disturbed Ionospheric Conditions

Disturbed ionospheric conditions arise when external energy inputs from space weather or atmospheric sources alter the ionospheric system equilibrium. These events disrupt the balance of ionization, plasma transport, and thermospheric dynamics that characterize quiet-time variability. They are primarily driven by enhanced solar and geomagnetic activity, including solar flares, coronal mass ejections (CMEs), high-speed solar wind streams (HSSs), and energetic particle precipitation. Additional disturbances originate from below, such as atmospheric waves launched by surface or lower-atmospheric processes.

Solar flares, associated with magnetic reconnection events in the solar corona, emit bursts of ultraviolet and X-ray radiation that reach the Earth in ~ 8 minutes, directly impacting the D- and E-regions of the ionosphere (Donnelly, 1976). These rapid enhancements in shortwave irradiance can cause sudden ionospheric disturbances (SIDs), such as Shortwave Fadeouts or dramatic increases in Total Electron Content (TEC).

Larger-scale and longer-lasting disturbances arise from CMEs and HSSs, which interact with Earth's magnetosphere through solar wind–magnetosphere–ionosphere coupling. When CMEs reach Earth (typically 1–3 days after ejection), they compress the magnetosphere and inject energy through magnetic reconnection at the dayside magnetopause. The interplanetary magnetic field (IMF), particularly its southward B_z component, determines how effectively solar wind energy couples into the system (Liu & Li, 2002). This interaction initiates geomagnetic storms, complex events that unfold in distinct phases: initial, main, and recovery.

The initial phase is marked by the arrival of a solar wind pressure front, often from a CME or interplanetary shock. This phase produces a sudden compression of Earth's magnetosphere and can be observed as a storm sudden commencement (SSC) in ground-based magnetometer data (Willis, 1964; Motoba et al., 2003). While ionospheric responses at this stage are generally modest, electric fields can begin penetrating to lower latitudes, and slight TEC perturbations may occur due to dynamic pressure enhancements (Wang et al., 2010).

The main phase begins when sustained southward IMF transfers energy and particles into the magnetosphere. This injection triggers enhanced magnetospheric convection, strong electric field penetration, and intensified magnetospheric-ionospheric current systems. One of the primary features of this phase is the intensification of the auroral electrojet, a pair of high-latitude, eastward and westward ionospheric currents located in the auroral oval. These electrojets are driven by magnetospheric convection and field-aligned currents and are responsible for dramatic perturbations in the geomagnetic field's horizontal (H) component (Feldstein et al., 2006). The enhanced electrojets lead to significant Joule heating in the high-

latitude thermosphere, elevating neutral temperatures, increasing their altitudinal extent, and modifying global wind circulation. Due to coupling processes, this thermospheric disturbance alters the ionospheric composition, notably reducing the O/N_2 ratio and increasing recombination rates, contributing to widespread electron density depletions at midlatitudes (Bates, 1974; Rodger et al., 2001; Crowley & Meier, 2013; Spogli et al., 2024). At lower latitudes, the main phase is characterized by Prompt Penetration Electric Fields (PPEFs), which disrupt the F-region dynamo and uplift plasma to higher altitudes, increasing the EIA and sometimes producing super-fountain effects. These changes are often highly asymmetric, depending on local time, season, and storm strength (Tsurutani et al., 2008; Balan et al., 2011; Lu et al., 2012).

The recovery phase begins as solar wind conditions stabilize and energy input into the magnetosphere diminishes. However, the ionosphere exhibits delayed responses (Yermolaev, 2014; Chen et al., 2015). The global wind circulation, modified by thermospheric expansion during the main phase, generates Disturbance Dynamo Electric Fields (DDEFs). These fields evolve over several hours and often oppose the previously dominant PPEFs, inducing further redistribution of ionospheric plasma. Meanwhile, the auroral electrojets gradually weaken but may remain elevated due to substorm activity or slow magnetospheric relaxation. TEC gradients, troughs, and enhancements formed during the main phase often persist deep into the recovery stage. Figure 1.3 shows the ionospheric response at middle latitude to the Mother's Day storm of 2024 (Spogli et al., 2024). The vertical TEC (vTEC) time evolution, along with its 27-day median, shows a strong density depletion of TEC starting on the 11th of May and continuing until the 14th. The black squares, showing GUVI-measured O/N_2 ratio, demonstrate the relationship between the decreased neutral ratio and electron density depletion.




Figure 1.3: Panel (a) illustrates, for the interval from 12:00 UT on 9 May 2024 to 23:50 UT on 14 May 2024, the temporal evolution of the regional mean and standard deviation of vTEC derived from IONORING maps (blue line). For comparison, the 27-day median values are shown as a quiet-time reference (red dashed line). The panel also includes the daily averages of the $[O]/[N_2]$ ratio retrieved from GUVI observations over the same area (black symbols). Panel (b) displays the vTEC departures from the corresponding 27-day medians, presented as a function of both time and latitude. Figure taken from Spogli et al. (2024).

A clear consequence of this complex redistribution is the generation of ionospheric irregularities and fluctuations in electron density. In particular, LSTIDs frequently emerge during geomagnetic storms, originating from high-latitude heating and propagating

equatorward with coherent wavefronts and characteristic phase speeds of several hundred meters per second (Tsugawa et al., 2004; Ding et al., 2008). While small-scale irregularities and GNSS scintillation can occur, particularly at high latitudes (Beniguel & Adam, 2007; Tiwari et al., 2013), the dominant features during storms are large-scale deformations of ionospheric morphology, shaped by electric fields, thermospheric winds, and composition changes (Fuller-Rowell, 2011; Fuller-Rowell et al., 2013).

In addition to solar wind-driven disturbances, the ionosphere can be perturbed by atmospheric waves launched by impulsive geophysical events such as earthquakes, tsunamis, volcanic eruptions, or anthropogenic activities such as big explosions or rocket launches. These events excite upward propagating acoustic and gravity waves, which can reach ionospheric heights within minutes. The resulting disturbances appear as coherent, wave-like perturbations in TEC, often forming concentric patterns detectable by GNSS networks. These non-geomagnetic TIDs demonstrate the sensitivity of the ionosphere to Lithosphere–Atmosphere–Ionosphere Coupling (LAIC), and can occur independently of geomagnetic storm conditions (Jakowski et al., 2006; Jin et al., 2019; Shalimov et al., 2019).

1.3 Vertical coupling in the atmosphere-ionosphere system

Vertical coupling in the geospace environment refers to the transfer of energy and momentum from the lower atmosphere into the upper atmosphere and ionosphere through wave propagation. This process links atmospheric layers across a wide range of spatial and temporal scales and is primarily mediated by atmospheric waves that grow in amplitude as they travel upward due to the exponential decrease in neutral density with altitude (Volland, 1974; Yiğit & Medvedev, 2015).

The key wave types involved in vertical coupling are mostly acoustic and gravity waves, each generated by distinct sources and characterized by different propagation properties. Figure 1.4 highlights the frequency ranges of gravity and acoustic waves, along with the Brunt–Väisälä and acoustic cutoff frequencies.




Figure 1.4: Frequency and period domains of acoustic and gravity waves. The acoustic cut-off frequency, denoted ω_a , marks the lower limit above which acoustic waves ($\omega > \omega_a$) can propagate vertically in the atmosphere. The Brunt–Väisälä (buoyancy) frequency, ω_b , represents the upper bound for gravity waves, which occur at $\omega < \omega_b$. The intermediate band, where $\omega_b < \omega < \omega_a$ (shaded area), corresponds to evanescent waves, which can travel only in the horizontal direction. Image taken from Astafyeva (2019).

Gravity waves are buoyancy-driven oscillations with frequencies below the Brunt–Väisälä frequency (~ 2.8 mHz in the lower atmosphere). They are mainly generated by meteorological processes such as frontal systems, convection, jet streams, and orographic

forcing (Jewtougoff et al., 2013). Gravity waves can also be excited by tsunamis, volcanic eruptions, and plumes (Hickey, 2011; Wei et al., 2015). These waves typically propagate obliquely and can reach thermospheric altitudes in 30–60 minutes, where their interaction with the ionized environment alters neutral winds and background densities. Depending on their frequency and propagation direction, gravity waves may undergo filtering or ducting, and their transmission into the ionosphere depends critically on the background wind and temperature structure (Yeh et al., 1979; Boška & Šauli, 2001; Miyoshi & Fujiwara, 2008).




Figure 1.5: Variation of propagation time and sound speed (c_s) with altitude.

Acoustic waves have frequencies above the acoustic cutoff (~ 3.3 mHz) and are typically generated by impulsive, high-energy sources such as large earthquakes, seismic waves, explosions, or the vertical displacement of the ocean surface during tsunamigenic events. Acoustic waves propagate nearly vertically at the speed of sound, reaching ~ 800 – 1000 m/s at thermospheric heights. Figure 1.5 shows the sound speed and travel time of acoustic waves as a function of atmospheric altitude. Due to their high vertical speed and short travel time to ionospheric altitudes (often just 7 to 10 minutes), acoustic waves are among the earliest indicators of rapid energy transfer from the ground to the upper atmosphere.

In addition to acoustic and gravity waves, Lamb waves are another mode of vertical coupling, especially in the case of large-scale impulsive events. Due to the rarity of the generating event, Lamb waves have been detected only in a handful of cases. Lamb waves are non-dispersive surface waves that travel horizontally along the Earth's surface at speeds around 300–320 m/s, maintaining coherence over global distances. Unlike gravity or acoustic waves, Lamb waves remain confined in the troposphere but can indirectly influence the upper atmosphere by exciting secondary waves with their passage. Their signature is often delayed but persistent, and they play a key role in modulating the long-range ionospheric response to sources such as explosive volcanic eruptions (Kulichkov et al., 2022).

The efficiency of vertical coupling depends not only on the wave source and type but also on background atmospheric and ionospheric conditions. In regions dominated by neutral density, wave propagation is shaped by filtering effects and can be either suppressed or amplified by background atmospheric winds. As neutral waves propagate upward into regions where the ionized component becomes significant, their ability to couple with plasma evolves with altitude, governed by the height-dependent balance between neutral-plasma collision frequencies and electron and ion gyrofrequencies. Neutral-plasma collisions dominate in the lower thermosphere, typically below ~ 90 km. At these altitudes, both electrons and ions experience collision frequencies higher than their respective gyrofrequencies, preventing charged particles from completing a full gyration around the magnetic field. In this unmagnetized regime, the plasma is effectively decoupled from the magnetic field and moves

coherently with the neutral atmosphere. Here, momentum coupling efficiently transfers wave-induced perturbations in density and winds to the plasma. At intermediate altitudes, roughly between 90 and 150 km, due to the difference in ion and electron gyrofrequency, electrons become magnetized while ions remain collision-dominated and unmagnetized. In this transitional regime, the plasma response to neutral perturbations becomes species-specific, as only electrons begin to feel the influence of the magnetic field, while ions remain tied to the neutral background. At higher altitudes, above ~ 150 km, where the bulk of the ionospheric plasma resides, both electrons and ions are fully magnetized. In this region, plasma motion becomes strongly constrained along magnetic field lines, and wave-plasma coupling efficiency becomes highly anisotropic. The ability of neutral waves to redistribute plasma now depends not only on their amplitude and scale but also on the orientation of the wave vector relative to the magnetic field. Here, wave-induced disturbances extend beyond the neutral component, driving spatially structured changes in electron density through a combination of ion-neutral collisions, plasma transport, and composition-dependent recombination. Figure 1.6 shows the modeled anisotropy of ionospheric disturbances caused by an acoustic neutral wave excited by an earthquake. Specifically, Figure 1.6a shows how the relative orientation of the acoustic wave vector and magnetic field is responsible for the anisotropy in the ionospheric response as plasma motion is restricted along magnetic field lines at such altitudes. Figure 1.6b and Figure 1.6c further confirm this as they represent horizontal slices at 280 km of electron density fluctuations for the scenario without and with the magnetic field, respectively.




Figure 1.6: Influence of the geomagnetic field on the evolution of coseismic ionospheric disturbances for Northern Hemisphere conditions. (a) Vertical cross-section of electron density perturbations (ΔN_e) 13 minutes after the earthquake. The solid black arrow marks the magnetic field vector B , while the dotted arrow shows the neutral perturbation velocity. (b–c) Horizontal slices of ΔN_e at 280 km altitude: (b) without and (c) with geomagnetic field effects. (d) Ionospheric radiation pattern at 280 km altitude and 200 km epicentral distance, with blue denoting negative polarity and red positive polarity. Adapted and modified from Rolland et al. (2013).

1.4 Radio propagation in the ionosphere

Radio waves travelling through the ionosphere interact with its plasma and, as a result, they experience bending (refraction), slowing down (group delay), distortion, and, in some cases, reflection or scattering. The theoretical basis for understanding these effects is given by Equation 1.1, which describes the refractive index of a cold, magnetized plasma. Equation 1.1 was first developed by Edward Appleton and later expanded by Douglas Hartree to include the influence of magnetic fields and wave polarization. The full form of the Appleton-Hartree equation is:

$$n^2 = 1 - \frac{X}{1 - iZ - \frac{Y_T^2}{2(1 - X - iZ)} \pm \sqrt{\frac{Y_T^2}{4(1 - X - iZ)^2 + Y_L}}} \quad (1.1)$$

$$X = \frac{\omega_0^2}{\omega^2}, \quad Y = \frac{\omega_B}{\omega}, \quad Z = \frac{\nu}{\omega}, \quad \omega = 2\pi f, \quad \omega_0 = \sqrt{\frac{N_e q_e^2}{\epsilon_0 m_e}}, \quad \omega_B = \frac{B q_e}{m_e}$$

where n is the refractive index, N_e is the electron density, ϵ_0 is the vacuum permittivity, m_e is the electron mass, ν is the collision frequency between neutral and charged particles, and q_e is the electron charge. X , Y , and Z are dimensionless parameters, with X being the ratio between the plasma frequency (ω_0) and the wave frequency (ω), Y represents the effect of the magnetic field through the electron gyrofrequency (ω_B) along the transverse and longitudinal components of the wave vector (Y_T and Y_L , respectively) and Z accounts for collisions between electrons and neutrals. Although the full expression is complicated, its physical meaning is straightforward: it shows how wave propagation depends on the local electron density, the magnetic field, and the wave frequency. Equation 1.1 explains the presence of two propagation modes (ordinary and extraordinary), the bending of waves at lower frequencies, and the cutoff or reflection of waves under certain conditions. At high frequencies, such as those used in GNSS, typically above 1 GHz, the value of X is very small, and the ionosphere acts like a weakly refracting medium. Thus, the GNSS signal travels through the ionosphere with only a small delay depending on the total electron content along the path.

In contrast, High-frequency (HF) radio waves, typically in the range of 3 to 30 MHz, are reflected by the ionosphere under certain conditions. This behavior makes long-distance communication possible without satellites, as radio signals can bounce between the ionosphere and the Earth's surface. The key to understanding this reflection lies in how the ionosphere affects the refractive index of the medium, which is given in Equation 1.1. In a simplified, collisionless form ($\nu = 0$) and neglecting the Earth's magnetic field ($B = 0$), Equation 1.1 reduces to:

$$n^2 = 1 - X = 1 - \frac{\omega_0^2}{\omega^2} \quad (1.2)$$

Equation 1.2 shows how the plasma frequency defines the cutoff below which radio waves cannot propagate through the ionosphere. If the wave frequency ω is lower than the local plasma frequency ω_0 , then:

$$n^2 < 0 \Rightarrow n \text{ is imaginary} \quad (1.3)$$

this means that the wave becomes evanescent, meaning that it cannot propagate and instead is reflected toward the Earth. The altitude at which $\omega = \omega_0$ is called reflection height. This height changes with electron density, which varies with altitude, time of day, season, solar activity, and geographic location. In practice, the critical frequency f_c (or ω_c) of the ionosphere is the highest frequency that can be reflected vertically. The critical frequency can be obtained by rewriting Equation 1.2 and setting the condition $n^2 = 0$:

$$0 = 1 - \frac{\omega_0^2}{\omega^2} \Rightarrow \omega_c^2 = \omega_0^2 \Rightarrow 4\pi^2 f_c^2 = \frac{N_e q_e^2}{\epsilon_0 m_e} \Rightarrow f_c = \sqrt{\frac{N_e q_e^2}{4\pi^2 \epsilon_0 m_e}} \approx 9\sqrt{N_e} \quad (1.4)$$

For example, if the maximum electron density in the ionosphere is 10^{12} el/m^3 , by applying Equation 1.4, one obtains a critical frequency around 9 MHz . All waves below this frequency emitted vertically will be reflected, while waves above this frequency will penetrate the ionosphere and continue into space.

1.4.1 HF Sounding of the ionosphere

High-frequency (HF) sounding techniques, particularly ionosondes, utilize the reflection properties of radio waves to study ionospheric electron density vertical profiles. An ionosonde transmits short HF pulses across a range of frequencies, typically from 1 to 20 MHz, measuring the time it takes for each pulse to be reflected in the ionosphere and return. By sweeping across frequencies, ionosondes identify critical frequencies and virtual reflection heights, and display them as ionograms. Figure 1.7 shows an ionogram obtained from the EISCAT ionosonde on 20 May 1994 at 11:22 UT.




Figure 1.7: Ionogram recorded by the EISCAT Dynasonde on 20 May 1994 at 11:22 UT. The virtual reflection heights of the main ionospheric layers are indicated: $h'E$, $h'F1$, and $h'F2$. The green trace represents the ordinary (O) mode, while the red trace corresponds to the extraordinary (X) mode. The critical frequencies foE , $foF1$, $foF2$, and $fxF2$ are also marked. Figure taken from Davis (1996).

Close to the maximum plasma frequency, echoes often split into two traces, corresponding to the ordinary and extraordinary waves, due to the presence of the Earth's magnetic field. The two traces are visible in Figure 1.7 in green and red. The two propagation modes are separated by approximately half the electron gyrofrequency ($\omega_{ce} = \frac{eB}{m_e c}$), with ordinary waves corresponding to lower frequencies (e.g., f_oF2) and extraordinary to higher ones (e.g., f_xF2).

The height scale of an ionogram assumes that radio waves propagate vertically at the speed of light in a vacuum, thus resulting in measured reflection altitudes that exceed the actual reflection points due to the dispersive nature of the ionosphere. These apparent altitudes are termed virtual heights (e.g., $h'F2$ of Figure 1.7). The virtual height h' is related to the measured group travel time τ of the radio pulse as follows:

$$h' = \frac{c\tau}{2} \quad (1.5)$$

Inversion techniques are required to correct virtual heights for refraction to determine the real electron density profiles. This technique involves solving an inverse problem using the frequency-dependent bending of radio paths, often under the assumption of monotonic or parabolic electron density distributions.

Oblique sounding methods complement vertical ionosonde techniques by transmitting HF waves at oblique angles. Such waves can reflect even at frequencies higher than the vertical critical frequency because the effective path through the ionosphere is longer, enabling reflection at lower electron densities. This behavior is described by the secant law:

$$f_{oblique} = \frac{f_c}{\cos\theta} \quad (1.6)$$

where θ is the angle of incidence relative to the vertical. In addition to virtual heights and critical frequencies, a crucial parameter in HF sounding is the Maximum Usable Frequency (MUF), and its related factor M . MUF represents the highest frequency usable for a radio link over a given distance, considering a single reflection path. The factor M is usually considered for a standard distance of 3000 km, $M(3000)$, often specified per ionospheric layer, such as $M(3000)F2$ or $M(3000)D$. To derive $M(3000)$, the $MUF(3000)$ obtained from vertical ionograms is used by applying the following relation with the critical frequency (f_c):

$$M(3000) = \frac{MUF(3000)}{f_c} \quad (1.7)$$

Traditionally, $MUF(3000)$ is obtained graphically by overlaying standard transmission curves onto ionograms and selecting the transmission curve tangent to the ionogram trace of the layer of interest.

1.4.2 GNSS and Total Electron Content

Another technology influenced by radio wave propagation in the ionosphere is the Global Navigation Satellite Systems (GNSS), an infrastructure that provides continuous positioning, navigation, and timing (PNT) services worldwide. These systems operate through constellations of satellites that transmit precise timing signals, which are received and processed by ground-based receivers. By measuring the time delays from multiple satellites, a GNSS receiver can determine its position along with its clock offset. GNSS constellations typically operate in medium Earth orbit (MEO), at altitudes of $\sim 20,000$ km, providing global coverage with favorable visibility and orbit stability. In addition to global GNSS constellations, GPS (United States), GLONASS (Russia), BeiDou (China), and Galileo (Europe), there are regional navigational systems such as NavIC (India) and QZSS (Japan). Regardless of their scale or configuration, all these systems rely on the propagation of radio signals through the Earth's atmosphere to deliver positioning and timing information to users on the ground. Each GNSS satellite broadcasts signals on two or more microwave frequencies, typically in the L-

band (1–2 GHz), including a code (or pseudorandom noise, PRN) on top of a carrier phase component.

As GNSS signals travel from a satellite to a ground-based receiver through the ionosphere, free electrons alter the speed and path of the radio wave. This interaction introduces positioning errors, but it also offers a way to observe and study the ionosphere. In particular, the delay introduced by the ionosphere is proportional to the TEC along the signal path, a measure of the integrated electron density along the signal trajectory. Accordingly, TEC is defined as the line integral of the electron density N_e along the signal path s and is usually expressed in TEC units (TECu), where 1 TECu corresponds to 10^{16} em^{-2} :

$$sTEC = \int_{sat}^{rec} N_e(s) ds \quad (1.8)$$

Since GNSS signals pass through different ionosphere regions depending on satellite elevation, the measured TEC along the signal path is called slant TEC (sTEC). To understand how TEC affects GNSS signal propagation, it is useful to start with the simplified Appleton-Hartree equation (Equation 1.2). The effects of the Earth's magnetic field and collisions are negligible at GNSS frequencies, as they are much higher than the plasma frequency. Under these assumptions, the refractive index n reduces to:

$$n = \sqrt{1 - X} \text{ with } X = \frac{\omega_0^2}{\omega^2} \quad (1.9)$$

As for GNSS frequencies $X \ll 1$ (since $\omega \gg \omega_0$), one can expand the square root using the first-order Taylor expansion:

$$\sqrt{1 - X} \approx 1 - \frac{1}{2}X \quad (1.10)$$

The first-order approximation of the refractive index thus becomes:

$$n \approx 1 - \frac{1}{2} \frac{\omega_0^2}{\omega^2} \quad (1.11)$$

or, in terms of electron density:

$$n \approx 1 - \frac{N_e q_e^2}{2\epsilon_0 m_e \omega^2} \quad (1.12)$$

By switching to frequency (using f in Hz) and merging all constant terms, Equation 1.12 simplifies to:

$$n \approx 1 - \frac{40.3 N_e}{f^2} \quad (1.13)$$

Equation 1.13 shows that the refractive index of the ionosphere is slightly less than one (roughly $1-10^{-5}$), and the deviation depends on both the local electron density and the signal frequency. Higher-frequency signals are thus less affected by the ionosphere, which is why GNSS operates in the L-band well above the typical plasma frequencies. As a signal travels through the ionosphere, its optical path is affected by the local refractive index, and thus the total optical path length Λ from the satellite to the receiver is:

$$\Lambda = \int n \, ds = \int \left(1 - \frac{40.3N_e}{f^2}\right) ds = S - \frac{40.3}{f^2} TEC \quad (1.14)$$

where S is the geometric distance between the satellite and receiver. The second term of Equation 1.14 represents the ionospheric effect, which is dispersive as it depends on the frequency. Moreover, as the ionosphere is a dispersive medium, the group and phase velocities differ, affecting the two main GNSS observables in opposite ways. The code signal travels with the group velocity, slower than light in a dispersive medium, thus resulting in a delay in time. On the other hand, the carrier phase travels with the phase velocity, faster than light in plasma, leading to a phase advance.

The phase observable is normally expressed in cycles (Φ) and is obtained by dividing the optical path (Λ) by the signal wavelength (λ):

$$\Phi[cycle] = \frac{\Lambda}{\lambda} = \frac{1}{\lambda} \left(S - \frac{40.3}{f^2} sTEC \right) = \frac{fS}{c} - \frac{40.3}{cf} sTEC \quad (1.15)$$

Considering the signal wavelength λ , Equation 1.15 can be rewritten to obtain the phase in meters:

$$L[m] = S - \frac{40.3}{f^2} sTEC \quad (1.16)$$

where L represents the phase observable in meters. The code/group delay (δ), expressed in seconds, is derived from the frequency derivative of Equation 1.16:

$$\delta[s] = \frac{dL}{df} = \frac{S}{c} + \frac{40.3}{cf^2} sTEC \quad (1.17)$$

By multiplying by the speed of light, one can obtain this as a distance:

$$P[m] = \frac{dL}{df} = S + \frac{40.3}{f^2} sTEC \quad (1.18)$$

where P is the code observable in meters. Equations 1.16 and 1.18 show how the ionospheric contribution appears with opposite signs for the two observables. From a practical perspective, code and carrier observables are complementary. Code measurements are unambiguous but relatively noisy, with a typical precision of around 0.5 meters for high-end receivers. Carrier phase measurements, in contrast, offer millimeter-level precision but are ambiguous due to the unknown number of full phase cycles at the start of the observation. In addition, dual-frequency measurements can be combined to obtain new observables designed to isolate different contributions in the measured GNSS signal. One of these is the geometry-free linear combination (GFLC). The ability of the GFLC to isolate the ionospheric component from geometric and clock terms makes it ideal for capturing the time evolution of the electron content, which is critical in identifying and characterizing ionospheric disturbances. The GFLC of phase measurements is equal to the difference between the signal phase measured for two different carrier frequencies:

$$GFLC = L_1 - L_2 = 40.3 \left(\frac{1}{f_2^2} - \frac{1}{f_1^2} \right) sTEC \quad (1.19)$$

This combination cancels out the geometric distance and receiver and satellite clock errors, leaving a residual dominated by frequency-dependent effects, such as ionospheric delay,

instrumental biases, and ambiguities. The ambiguous sTEC from carrier phase measurements is then estimated as:

$$sTEC = \frac{1}{40.3} \frac{f_1^2 f_2^2}{f_1^2 - f_2^2} (L_1 - L_2) \quad (1.20)$$

where L_1 and L_2 are the carrier phase measurements in meters at frequencies f_1 and f_2 . The retrieved TEC in this case includes a constant bias term that originates from phase ambiguities. As such, it is referred to as an uncalibrated sTEC. While the absolute value of TEC remains unknown without calibration, the time evolution of this estimate is accurate and sufficient for studying ionospheric dynamics.

As sTEC quantifies the total number of free electrons integrated along an oblique path between a GNSS satellite and a ground-based receiver, it is often useful to project this quantity to the vertical direction to allow comparisons between different observational geometries. The verticalized version of sTEC is called Vertical Total Electron Content (vTEC). Figure 1.8 shows the geometrical relation between the two measurements and their relative angles. In most studies, the verticalization process is performed assuming a 2D ionosphere at a fixed altitude (Hipp in Figure 1.8), typically fixed equal to the height of the ionization maximum (hmF2). The conversion is done using a mapping function based on the elevation angle of the satellite:

$$vTEC = sTEC \cdot F(\alpha_{elev}) = sTEC \cdot \sqrt{1 - \left[\frac{R_e \cos(\alpha_{elev})}{R_e + H_{ipp}} \right]^2} \quad (1.21)$$

where R_e is the Earth radius, H_{ipp} is the height of the ionospheric thin shell layer and α_{elev} is the elevation of the satellite. H_{ipp} is usually between 200 and 400 km, depending on solar, geomagnetic, seasonal and diurnal conditions and based on geographic location. As shown in Figure 1.8, the intersection between the line of sight (LOS) and the 2D ionosphere is called the ionospheric piercing point (IPP). The projection of the IPP on the Earth's surface is called the sub-ionospheric point (SIP). Note here that changing the height of the ionospheric layer will change the location of the IPP and SIP. This verticalization is essential when comparing measurements from different satellites or constructing ionospheric maps.




Figure 1.8: Geometry of a single-layer ionospheric shell model at altitude H_{ipp} . Figure taken from Janssen (2012).

1.5 Travelling ionospheric disturbances

The dynamic interaction between atmospheric waves and the ionosphere, described in the previous section, leads to various ionospheric disturbances. Among these, Traveling Ionospheric Disturbances (TIDs) are one of the most studied. These periodic electron density fluctuations arise from the coupling processes driven by wave propagation through the atmosphere and can be triggered both from above and below.

TIDs are typically classified based on their horizontal spatial scale, period, and propagation speed. LSTIDs have horizontal wavelengths exceeding 1,000 kilometers, periods ranging from 30 minutes to several hours, and phase velocities typically between 400 and 1,000 m/s. They are usually associated with geomagnetic disturbances and auroral processes. Medium-Scale TIDs (MSTIDs) have smaller horizontal wavelengths, typically between 100 and 1,000 kilometers, periods of 15 to 60 minutes, and normally propagate slower than LSTIDs, with phase velocities in the range of 100 to 300 m/s (Hocke & Schlegel, 1996; Borries et al., 2009; Katamzi et al., 2012; Thaganyana et al., 2022). MSTIDs are more commonly linked to sources in the lower atmosphere, such as tropospheric weather or seismic activity, and can also arise from ionospheric instabilities. Small-Scale TIDs have horizontal scales of tens of kilometers and shorter periods. SSTIDs remain poorly characterized due to observational limitations but are thought to result from localized instabilities or secondary interactions and are an active area of ongoing research.

1.5.1 Forcing from above

Although the full chain of physical mechanisms behind LSTID generation remains complex and not fully understood, the primary cause is likely ascribed to atmospheric gravity waves (AGWs) produced by thermospheric heating associated with Joule heating, particle precipitation and Lorentz forces (Banks, 1977). Oyama & Watkins (2012) demonstrated that these forces impact wind velocity and induce temperature fluctuations with different magnitudes. Lorentz forces are responsible for increased wind velocities, while Joule heating causes a rise in thermospheric temperature. As the neutral waves interact with the ionosphere, they produce coherent variations in electron density that propagate along field lines across vast distances. Background atmospheric and ionospheric conditions strongly modulate the propagation of LSTIDs. Factors such as neutral wind patterns and geomagnetic field geometry influence their direction, amplitude, and dissipation. LSTIDs often exhibit a preferred southward or southwestward propagation direction in the northern hemisphere, consistent with known thermospheric circulation patterns during disturbed conditions and due to the plasma movement at F-layer altitudes being mainly restricted along magnetic field lines. However, LSTID during quiet periods and propagating poleward rather than equatorward have also been observed (Habarulema et al., 2015; Habarulema et al., 2016). In addition to background characteristics, LSTID parameters can vary significantly depending on the timing and intensity of the geomagnetic disturbance and the presence of concurrent processes such as substorm activity or polar cap expansion.




Figure 1.9: Two-dimensional TEC perturbation maps over South America at different times (a–i) between 14:40 and 16:50 UT on 17 March 2015. Each panel corresponds to a 10-minute map, shown at 20-minute intervals. Numbers mark the LSTIDs propagating equatorward from both North and South America. The solid black line indicates the geomagnetic equator. Figure taken from Zakharenkova et al. (2016).

LSTID signatures are identifiable in GNSS-derived TEC maps, appearing as large-scale, coherent bands of increased/decreased TEC propagating over time. LSTIDs can be observed in Figure 1.9, which shows detrended TEC (dTEC) maps over South America for the geomagnetic storm of the 17th of March 2015. Many LSTIDs were generated at auroral latitudes during this storm and propagated toward the equator, where they interacted. These wavefronts can also be detected through ionosondes, which reveal synchronized variations in ionospheric parameters such as MUF. Regarding geophysical coupling, LSTIDs are a direct link in the magnetosphere-ionosphere–thermosphere system as they carry significant amounts of energy and momentum and contribute to the redistribution of ionospheric plasma.

1.5.2 Forcing from below

Medium-Scale Traveling Ionospheric Disturbances (MSTIDs) represent one of the most dynamic and diverse categories of TIDs. Their spatial scales, typically ranging from 100 to 1,000 km in horizontal wavelength and 15 to 60 minutes in period, place them at the interface between global-scale ionospheric disturbances and finer-scale structures. Unlike LSTIDs, whose occurrence is generally tied to geomagnetic activity and high-latitude energy inputs,

MSTIDs are predominantly observed under quiet geomagnetic conditions and are often associated with coupling from the lower atmosphere or electrodynamic instabilities internal to the ionosphere.

One of the most common MSTID sources is atmospheric gravity waves generated by tropospheric weather systems, including frontal boundaries, convection, orographic forcing, and jet streams. When AGWs reach ionospheric altitudes, they modulate plasma density and generate coherent TEC fluctuations. However, MSTIDs are not limited to meteorological forcing. Natural hazards such as tsunamis, earthquakes, and volcanic eruptions can produce strong acoustic and gravity waves that trigger MSTIDs with distinctive spatial and temporal patterns. Figure 1.10, for example, shows the acoustic-gravity wave field generated by the seafloor uplift during the Tohoku-Oki earthquake of the 11th of March 2011. On the right, it is possible to see the acoustic response appearing ~ 10 minutes after the earthquake as short-period oscillations denoted by high propagation speed. The right panel of Figure 1.10 instead shows the gravity wave field generated by the tsunami, visible in both the travel distance vs time plot and the dTEC maps below. As introduced in previous sections, gravity waves are characterized by a slower propagation speed, which is translated to a smaller inclination in travel-time plots.

In addition to mechanically driven MSTIDs, electrodynamic processes play a key role, particularly during nighttime at midlatitudes. The Perkins instability, long considered a dominant driver, arises from a plasma density gradient aligned with the geomagnetic field and is often enhanced by coupling between the E and F regions. This coupling, especially when sporadic E layers are present, can create conditions favorable for wave growth and horizontal propagation. The resulting disturbances often exhibit northwest-southeast alignments in the northern hemisphere and have a clear seasonal and local time dependence.




Figure 1.10. Evolution of TID signatures in the 4 hours following the earthquake/tsunami. Detrending was performed using the double-difference method with time separations of 60 s (left) and 300 s (right). The keogram panels display disturbances as a function of distance from the epicenter and time (04:00–10:00 UT) using GEONET GNSS data for all GPS satellites on day 70 of 2011. The lower row shows the corresponding detrended VTEC maps, aligned with the keograms to highlight their relationship. Figure taken from Yang et al. (2019).

1.5.3 TID Detection

The detection and characterization of TIDs rely on identifying wave-like fluctuations in ionospheric measurements, such as TEC or MUF, by extracting periodic features using detrending and filtering techniques applied both in post-processing and in near real-time (Belehaki et al., 2020; Borries et al., 2023; Guerra et al., 2025; Maletckii & Astafyeva, 2022; Segarra et al., 2025). NRT detection is particularly critical for natural hazard early warning systems, where high reliability and low latency are essential, and several GNSS-based approaches have been developed to meet this need. The VARION (Ravanelli et al., 2021; Savastano et al., 2017) algorithm, initially designed for tsunami monitoring, computes the time-differenced GFLC of GNSS phase measurements to generate $d\text{TEC}/dt$ observables in real time, removing constant biases and allowing subsequent bandpass or polynomial filtering to isolate TID signatures. The GUARDIAN (Martire et al., 2022) system operated by NASA-JPL leverages high-pass Butterworth filtering of uncalibrated GFLC data from a Pacific GNSS network to produce continuous TEC estimates for hazard monitoring, with future upgrades aimed at fully automated detection. Additional NRT detection solutions are based on Networked Transport of RTCM via Internet Protocol (NTRIP), which allows the construction of TEC derivatives that mitigate boundary effects and further reduce latency, and all of these methods have demonstrated performance comparable to post-processed products, confirming their suitability for operational applications with minimal delay (Maletckii et al., 2023; Maletckii & Astafyeva, 2022).

Beyond these NRT systems, automatic TID detection has been pursued using statistical and machine learning methods such as random forests (Brissaud & Astafyeva, 2022), long short-term memory networks (Luhmann et al., 2022), gradient boosting (Fuso et al., 2024), and convolutional neural networks (Ravanelli et al., 2024), which can discriminate between noise and event-related disturbances, classify TID signatures, and estimate arrival times; however, their use has so far been constrained largely to post-processing due to computational and operational limitations.

Although GNSS TEC has become the dominant resource for TID detection thanks to its dense global coverage and continuous availability, its line-of-sight integration prevents vertical resolution, making ionosondes and other instruments valuable complementary tools. Several studies have addressed the automatic detection of TIDs using ionosonde data. For example, the TechTIDE project developed various methodologies for the real-time detection, characterization, and tracking of TIDs based on ionosonde and other data sources (Belehaki et al., 2020). One of these methods specifically leverages ionosonde measurements to identify and track LSTIDs in real-time, as Altadill et al. (2020) demonstrated. More recently, Segarra et al. (2025) presented a climatological analysis of LSTID activity over Europe from 2014 to 2023 using the HF-INT method, which enables near-real-time detection from a European ionosonde network. Their Catalogue of LSTID events includes onset times, durations, dominant periods, and propagation vectors, highlighting the dependence of LSTID activity on local time, season, and geomagnetic conditions. These studies confirm that ionosonde-based automatic detection systems can be exploited in NRT and demonstrate their relevance in operational space weather services. Other independent TEC-based methods not originally designed for LSTID detection have also shown the capability to capture such disturbances under specific conditions.

Chapter 2

Examples of coupling processes

In order to move from general detection principles to concrete applications, this chapter presents a series of case studies that demonstrate how TIDs can be identified and characterized under different geophysical conditions. These examples serve a dual purpose: they highlight the versatility of GNSS-based techniques for extracting ionospheric disturbances, and they illustrate how different sources, whether lithospheric, atmospheric, or magnetospheric, can imprint distinct signatures on the ionosphere. The three sections focus on my contributions to published studies, addressing TIDs triggered by the Hunga Tonga–Hunga Ha‘apai eruption, the Türkiye–Syria earthquake, and tropospheric weather systems. Together, these cases provide a comparative framework that underscores the challenges and the opportunities of TID detection across diverse geophysical drivers.

2.1 Hunga Tonga Volcanic Eruption

After a series of smaller disturbances starting at the end of 2021, on January 15, 2022, at 04:15 UT, the Hunga volcano in Tonga (20.54°S, 175.38°E) (Cronin et al., 2017) produced a powerful eruption. In terms of the total energy released, this was the largest volcanic eruption since the 1991 eruption of Mount Pinatubo. Large eruptions are known to produce wave-like disturbances up to ionospheric altitudes. TIDs of volcanic origin have been observed, for instance, after the Pinatubo eruption of 1991 (Igarashi et al., 1994), as well as the 2004 eruption of the Asama volcano in Japan (Heki, 2006), the 2003 eruption of Soufrière Hills on Montserrat (Dautermann et al., 2009), and even earlier in the context of the 1980 eruption of Mount St. Helens (Roberts et al., 1982). Nevertheless, due to the infrequent occurrence of large volcanic eruptions, fewer studies have been done on their ionospheric effects than, for instance, on the effects of earthquakes and tsunamis (Astafyeva, 2019).

Two different physical mechanisms can cause TIDs to appear after a volcanic eruption: disturbances can be produced directly in the ionosphere at the location of the eruption and travel radially outwards at ionospheric altitudes, or ionospheric disturbances can result from the upward propagation of secondary/tertiary waves excited by waves propagating in the lower atmosphere. A volcanic eruption can produce waves of all types, propagating outward from the eruption at different velocities and distances before dissipating. A Lamb wave was detected as the leading wavefront for this event, followed by various disturbances of different natures (Burt, 2022; Kubota et al., 2022; Kulichkov et al., 2022; Saito, 2022). MSTIDs propagating through the ionosphere from the site of the eruption are not expected to proceed to such distances (Haaser et al., 2017; Astafyeva, 2019). Themens et al. (2022) indeed find that LSTIDs produced directly in the thermosphere above the eruption dissipate after a few thousand kilometres. However, medium-scale TIDs travelling with the Lamb wavefront seem to propagate further (Themens et al., 2022; Zhang et al., 2022), and clear signatures of waves travelling around the globe in the troposphere have been detected (Burt, 2022; Matoza et al., 2022; Wright et al., 2022).

2.1.1 Data and methods




Figure 2.1: Location of the GNSS receivers used for detrended TEC analysis.

In this study, I studied the ionospheric response to the volcanic eruption by leveraging GNSS and in situ electron density data from satellite observations. First, the GNSS data from different networks covering the European sector, shown in Figure 2.1, have been used to investigate the impact of the Hunga eruption on the ionosphere in terms of TEC. TEC at 30-second resolution has been obtained and detrended using the VARION algorithm to highlight the effect of the wave-like structure induced by the explosion. In situ measurements from the Swarm Charlie (C) satellite (Friis-Christensen et al., 2008) have also been considered to investigate the signature of the ionospheric variations related to the explosion. Specifically, the electron density provided by the Langmuir Probe on board the spacecraft and available at a 1 Hz rate was used (Jin et al., 2022). To highlight the possible signatures of the wave-like structure induced by the volcanic eruption in the Swarm C measurements, satellite ground tracks passing in the European sector and in time intervals compatible with the theoretical arrival time of the Lamb wave were selected. Because the speed of the swarm satellite (about 7.2 km/s) is much larger than the speed of the ionospheric waves (about 300 m/s), the ionospheric waves are considered fixed with respect to the satellite (Kil & Paxton, 2017; Urbar et al., 2022). To identify wave-like signatures in Swarm Ne data, I used the Fast Iterative Filtering (FIF) (Cicone, 2020; Cicone & Zhou, 2021). This algorithm can decompose a non-stationary, non-linear signal into simple oscillatory components, called intrinsic mode components or intrinsic mode functions. By summing intrinsic mode functions characterized by periodicities in the range of interest, it is possible to use FIF as a bandpass filter.




Figure 2.2: Keogram reporting the latitudinal distribution of the dTEC from 30°N to 70°N for the time interval 18:00 UT to 24:00 UT on January 15, 2022. To guide the reader's eye, the solid black lines indicate the wavefronts and the related velocity as inferred from the dTEC in the figure.

2.1.2 Results and discussion

Figure 2.2 presents the latitudinal distribution of the dTEC in a range from 30°N to 70°N in the time interval from 18:00 UT to 24:00 UT on January 15, 2022, by considering all the IPPs in a longitudinal range from 12.5°E to 17.5°E. Only values of dTEC > 0.1 TECu are shown. Solid black lines represent wavefronts propagating at 310 m/s, highlighting the presence of waves propagating southward from about 19:00 UT to 21.30 UT and northward after 22:00 UT. This double appearance is explained by the nearly isotropic propagation of Lamb waves, which radiate outward from the source in all horizontal directions. On a spherical Earth, an observer located close to the antipode of the eruption site can detect two distinct wave trains. The first corresponds to the direct path of the disturbance from the source region, while the second arises from the portion of the wavefront that propagated in the opposite direction and reached the observation area after propagating in the opposite direction over the globe and through the antipodal point. Consequently, the same eruptive event can give rise to two apparent propagation directions in the ionospheric response. The width of the blue and red bands showing the wave-like perturbations suggests that the southward wave has a smaller wavelength than the northward one. The presence of waves returning from the antipode confirms that the identified TIDs are not of auroral origin, as the volcanic eruption occurred during a minor geomagnetic storm.




Figure 2.3: Hodocrone reporting dTEC over Italy from 16,800 km to 18,500 km from the Hunga volcano as a function of time since the eruption (13–19 h). The solid and dotted black lines represent the theoretical time of arrival calculated by considering waves propagating at the velocity reported in the legend northward and southward, respectively. Dashed and dot-dashed black lines represent the corresponding error band of $\pm 5\%$.

Figure 2.3 shows the hodocrone (i.e., the travel time diagram) of dTEC over Italy, from 16,800 km to 18,500 km from the Hunga volcano as a function of time after the explosion. The solid and dotted black lines represent the theoretical time of arrival calculated by considering waves propagating at the velocity shown next to the lines. Dashed black lines represent the corresponding error band of $\pm 5\%$. The southward wavefronts of Figure 2.2 travel away from the eruption site when passing over Italy, while the northward wavefronts travel towards the eruption when passing over Italy. This behaviour is reflected in the opposite slope of the theoretical curves represented by the various black lines in Figure 14. In the case of the first wavefronts, the measured arrival time represented by the red and blue bands fits with the theoretical arrival time within the 5% error band. For the second wavefront, the measured arrival time is shorter than the calculated one.

Figure 2.4 compares the Swarm C detrended electron density measurements (dNe) with the dTEC map for the satellite track time. Specifically, Figure 2.4a shows the time profile of dNe (blue) from Swarm C, dTEC along the track (dTEC_{att}, black dots), and the corresponding spline fitting curve (red line) between 22:24:53 and 22:29:35 UT on January 15, 2022. The values of dTEC_{att} are the along-the-track (att) values of dTEC obtained by considering the dTEC values having the shorter spatio-temporal distance with each Swarm C measurement. Figure 2.4b shows dNe (black line) on top of the dTEC map for the same time interval as Figure 2.4a. Deviations from the dot-dashed line representing the satellite track (dNe = 0 cm⁻³) are proportional to the positive (right) and negative (left) dNe values. Notwithstanding the different geometry and nature of the observations, dTEC and dNe agree after 22:28 UT, i.e., the time after which the wave-like perturbations are enhanced.




Figure 2.4: (a) Time profile of dNe (blue) from Swarm C, dTEC_{att} (black dots), and corresponding spline fitting curve (red line) between 22:24:53 and 22:29:35 UT on January 15, 2022. (b) Map of dNe (blue line) from Swarm C on top of the dTEC map for the same time interval as (a). Deviations from the dot-dashed line are proportional to dNe values.

The analysis of GNSS-derived dTEC and Swarm C in situ electron density confirms that the disturbances observed over Europe were induced by the atmospheric waves generated by the Hunga eruption. The propagation velocities derived from dTEC are consistent with Lamb waves (~ 310 m/s), and the observation of both southward and northward propagating wavefronts reflects the expected geometry of isotropic propagation on a spherical Earth, with the second wave train corresponding to the return from the antipodal point. The different wavelengths inferred for the two propagation directions suggest structural modifications of the wavefront during its global travel.

The comparison between Swarm C dNe and dTEC along the satellite track further supports this interpretation. Despite the distinct measurement geometries, the two datasets converge after 22:28 UT, when the perturbations become more pronounced. This agreement highlights the complementarity of ground-based GNSS and in situ satellite measurements in capturing the spatio-temporal evolution of medium-scale travelling ionospheric disturbances.

Overall, the combined observations demonstrate that the ionospheric response to the Hunga eruption was shaped primarily by Lamb waves propagating globally in the lower atmosphere and generating secondary atmospheric waves coupled to the ionosphere. The consistency between dTEC and Swarm C strengthens confidence in attributing the detected MSTIDs to the eruption, showing that joint exploitation of GNSS and satellite in situ data provides a robust framework for characterising the ionospheric impact of large impulsive events.

2.2 Türkiye Earthquake

On 6 February 2023, two earthquakes with magnitude $MW > 7$ occurred in Turkey. The first shock was recorded at 01:17 UT with a magnitude of 7.8, while the second shock was at 10:24 UT with $MW = 7.7$ (US Geological Survey, 2023). Many aftershocks with magnitudes lower than 7 followed these primary shocks. Both major earthquakes happened in the region of the East Anatolian Fault, with the epicenters separated by about 95 km. The first event was located at 37.20°N , 37.13°E , and the second at 38.05°N , 37.25°E ; both events took place at a depth of around 10 km (Bondár & Storchak, 2011; Heki, 2021). The work of Leonard and

Barnes (1965) and Davies and Baker (1965) concerning the great Alaskan earthquake of 1964, has already demonstrated that major earthquakes can cause disturbances in the ionosphere. Since then, it has been established that these ionospheric disturbances are manifested as different types of earthquake-induced TIDs, propagating through different mechanisms (Astafyeva, 2019; Brissaud & Astafyeva, 2022; M. Li et al., 2023; S. Li et al., 2023; Meng et al., 2019). Co-seismic ionospheric disturbances are generated by waves traveling vertically up to the upper atmosphere in the vicinity of the epicenter and propagate at the speed of sound at ionospheric heights (Afraimovich et al., 2001; Afraimovich et al., 2006). On the other hand, seismic waves propagating out from the epicenter can also generate acoustic-gravity waves propagating up to the ionosphere (Astafyeva et al., 2009; Komjathy et al., 2016; Rolland et al., 2011). These disturbances are expected to propagate at the speed of the Rayleigh waves, between 2000 and 5,000 m/s, but with a delay of around 10 minutes required for the vertical propagation of disturbances from the ground to ionospheric altitude (Astafyeva, 2019; Lognonné et al., 2006a; Lognonné et al., 2006b). Since seismic waves on the ground can reach long distances, this mechanism can produce disturbances in the ionosphere beyond the range where the shock-acoustic waves traveling through the ionosphere are attenuated (e.g., Maruyama et al., 2016a, 2016b). The amplitude of ionospheric disturbances and the distance from the source at which they can be detected depend on the magnitude of the event (Heki, 2021). In addition to the earthquake magnitude, the depth and focal mechanism (Astafyeva & Heki, 2009) also affect the excitation and propagation of TIDs. On top of these primary earthquake attributes, additional factors such as atmospheric conditions (Rolland et al., 2011a; Rolland et al., 2011b), the orientation of the geomagnetic field (Astafyeva & Heki, 2009; Zettergren & Snively, 2019), and the observational geometry can also define the characteristics of the ionospheric response to tectonic events. Thus, a complex view of a superposition of different types of TIDs is observed after an earthquake, which differs significantly from one event to another.




Figure 2.5: Map showing locations of instruments used for this study. The earthquake epicenter is also shown (red asterisk).

2.2.1 Data and methods

To investigate the ionospheric signatures in TEC, I used a collection of GNSS networks spanning different distances and azimuthal directions relative to the epicenter (shown as yellow inverted triangles in Figure 2.5). Data with a time resolution of 1 and 30s were used, with 1s as the preferred time resolution when both were available due to the relatively short period expected from co-seismic TID (Astafyeva, 2019). The GNSS stations used belong to many different institutions and networks, specifically INGV (INGV RING Working Group, 2016), TUGASA-Aktif (Yildirim et al., 2011), CYPOS (Danezis et al., 2019), NOA (Chousianitis et al., 2021), IGS (Dow et al., 2009), and EUREF (Torres et al., 2009). Initially, the GNSS network covered the whole European region. Through visual inspection, it was noticed that no TEC arcs belonging to stations located further than around 1,500 km showed clear signatures related to the earthquake. Due to this, only the stations closer than such a distance were kept. The only exception to this criterion was applied to stations gathering data at 1s located in Graz, Matera, Prague and Dourbes. I used the dual-frequency GFLC of carrier-phase measurements to extract TEC perturbations. TEC was calibrated using NeQuick 2 following the findings of Chapter 3 of this thesis. To investigate the spatial behavior of the co-seismic TID, I rely on the widely used thin-layer ionospheric approximation (Mannucci et al., 1998), with the shell height set to 250 km. To extract the TID signature from the vTEC, a bandpass filter based on the FIF technique was used (Cicone & Zhou, 2021), with the passband set from 10 to 240s for 1s data and from 30 to 240s for 30s data.

2.1.2 Results and discussion

Figures 2.6, 2.7, and 2.8 show the ionospheric response and the TEC arcs for PRN17, PRN58, and PRN49, respectively. Note that the satellites considered were not the only ones showing clear TID signatures, but were chosen because they show the signature of both the shock acoustic wave (Afraimovich et al., 2001; Astafyeva et al., 2009; Heki & Ping, 2005) and the Rayleigh wave-induced TID (Ducic et al., 2003; Rolland et al., 2011a; Rolland et al., 2011b). Panels (b) show the TTD, with the X and Y axes representing the distance in time and space to the earthquake. Panels (a) show the spatial distribution of IPPs, the epicenter location, and its isodistances. Panels (b) show the TTD for the stations of Panels (a). As one might notice, the GNSS stations shown in Panels (a) of Figures 2.6, 2.7, and 2.8 are only a subset of the ones visible in Figure 2.5. This difference arises because only TEC arcs characterized by a dTEC amplitude higher than 0.05 TECu are shown to increase readability. By looking at the station distributions (Panels (a)), it is possible to investigate the different azimuthal behavior for different PRN links. In addition, the IPP tracks (Panels (a)) and dTEC (Panels(b)) are colored according to the initial arc azimuth relative to the epicenter to highlight the different ionospheric responses for different arcs.




Figure 2.6: Ionospheric response to the earthquake for satellite PRN17. (a) Spatial distribution of IPPs and the epicenter location. (b) TTD for the stations shown in panel (a). (c) TTD divided according to the TEC arcs azimuthal range.

Finally, Panels (c) show a slightly different TTD, where blue and red points correspond to negative and positive TEC perturbations. Moreover, the TTD was expanded further by dividing it into four subpanels, each corresponding to different azimuthal ranges for the first point of the IPP track relative to the earthquake epicenter. The four different azimuthal ranges are North (315° – 45°), South (135° – 225°), East (45° – 135°), and West (225° – 315°). This modification facilitates the investigation of the anisotropies in the TID propagation and parameters expected for co-seismic TIDs (Zettergren & Snively, 2019). Another difference between Panels (b) and Panels (c) is that the distance shown in Panels (b) is the distance of the given IPP at the time of maximum dTEC, while Panels (c) show its time evolution. Thanks to

the combination of the two TTDs, it is possible to investigate the waveform, amplitude, and propagation velocities for different azimuthal ranges.




Figure 2.7: Ionospheric response to the earthquake for satellite PRN58. (a) Spatial distribution of IPPs and the epicenter location. (b) TTD for the stations shown in panel (a). (c) TTD divided according to the TEC arcs azimuthal range.

First, Figure 2.6b shows how a narrow azimuthal range presents an unmatched amplitude. The corresponding IPP tracks of Figure 2.6a make it possible to discern which geographical area these azimuths correspond to. Such waveforms are related to GNSS stations located in Cyprus, and a possible explanation for such large amplitudes is the observational geometry (IPPs for PRN17 are pointing toward the epicenter at elevations around 40°) and the Rayleigh wave-induced TID characteristics. Looking at the modeling results presented by Rolland et al. (2013), it is possible to confirm that the lines of sight defined by low elevation and pointing to

the epicenter are expected to show the highest amplitude due to the integrated nature of TEC. All the waveforms visible in Figure 2.6b, 2.7b, and 2.8b resemble the typical acoustic N shape. Figures 2.6c, 2.7c, and 2.8c show the two types of ionospheric waves typically related to earthquakes. The first TID type, the co-seismic disturbance generated over the epicenter, is visible in the South and West subpanels of Figures 2.6c and 2.7c and the West subpanel of Figure 2.8c.




Figure 2.8: Ionospheric response to the earthquake for satellite PRN49. (a) Spatial distribution of IPPs and the epicenter location. (b) TTD for the stations shown in panel (a). (c) TTD divided according to the TEC arcs azimuthal range.

The near-field wave showed a large range of different amplitudes (from 0.1 to 1 TECu) and periods in the range of 2–3 min (this difference is explained by the 30s sampling rate and

by the Doppler effect induced by the IPP movement) and a propagation speed of around 1,200 km/s. Moreover, the difference in the TID velocity in the different figures is also due to the small distance covered by the first shock acoustic wave, making it difficult to manually identify such waves in a reliable and precise way. Moreover, the near-field TID shows almost no signature for those stations located North of the epicenter, which is expected due to the adverse geometry of the wave vector and magnetic field lines (MFLs). The lack of signatures East of the epicenter in Figures 2.6 and 2.7 is instead due to the scarcity of GNSS data accessible for those regions and due to the adverse observational geometry. A similar reasoning applies to the South panel of Figure 2.8c, where no clear signatures are visible even if the mutual orientation of the wave vector and MFLs is optimal. The few eastern stations available showed clear TID signatures East of the epicenter for PRN49, where the mutual orientation of the wave vector and observational link was favorable. This azimuthal anisotropy is in good agreement with previous studies (Bagiya et al., 2019; Rolland et al., 2013; Zettergren & Snively, 2019), where they used physical modeling to show how, for mid-latitude regions, the mutual orientation of acoustic wave vector and MFLs is optimal for equatorward regions with respect to the epicenter, while the opposite applies to poleward areas. Accordingly, the earthquake-induced signatures are clear only in the South, East, and West.

The second type of TID is generated by propagating seismic waves. Figure A.3 of the appendix confirms that the ionosonde network revealed the presence of an ionospheric disturbance appearing shortly after the passage of the Rayleigh wave at the ground, as measured by the seismic stations. The West subpanel of Figure 2.6c shows a clear signature of such Rayleigh wave-induced TIDs. Specifically, it propagated at around 3 km/s (in good agreement with the speed shown in Figure A.3) with a period of ~ 150 s, and the first signature was visible ~ 11 minutes after the earthquake. Since the expected delay is normally 8–9 min, the longer delay is due to the time necessary for the Rayleigh to propagate from the epicenter to the projection on the earth's surface of the first IPP that shows the TID (around 200 km, which corresponds to around 1 minute). Another interesting feature is that two waves overlap in the South quadrant of Figure 2.6c. The first TID signature (1,191 m/s) is interpreted as the co-seismic TID propagating from the epicenter, while for IPPs further than ~ 500 km, it looks as if a faster wave appeared before the near-field one and overlapped with it. This pattern could be explained by Rayleigh waves propagating through the ground at speeds around three times higher than the co-seismic TID, which propagates at the speed of sound in the F-layer. Therefore, the Rayleigh wave overcoming the slower near-field TID can explain the mode splitting at around ~ 500 km in the South quadrant. A similar behavior is also visible in the East quadrant of Figure 2.7c, where two TIDs appeared in the same observation arcs. The first one, with a speed of 2,623 m/s, is the Rayleigh wave signature, while the slower one is the near-field one. The arcs showing such signatures are all further than 600 km, which is consistent with Figure 2.6c, where the splitting of the two modes happens around 500 km of distance. This two-mode splitting behavior is typical of earthquake-induced TIDs, and many examples are available in the literature (e.g., Astafyeva, 2019).

Figure 2.9 shows the dTEC time evolution for GNSS stations in Athens, Prague, Matera, Dourbes, Graz, and Nicosia for six different PRNs that showed traveling ionospheric disturbance signatures. As one can notice from Figure 2.9, the only clear wavy signatures belong to the station located in Nicosia for PRN17 and to both Nicosia and Athens for PRN58. The TID signature belonging to the Nicosia station in PRN17 shows an unmatched amplitude of almost 1 TECu and appears around 10:35 UT, which is in accordance with ionosonde measurements (Appendix A.1). Looking at PRN 58 instead, two periodic features around 10:38 UT are visible, with an amplitude of 0.1 TECu. Here, there is again accordance in the timing between the Athens ionosonde and co-located GNSS station measurements. The GNSS station in Nicosia shows disturbances in PRN24 and PRN49, but due to the poor observational

geometry, the amplitude is not as strong as for the two aforementioned PRNs. The timing of such faint disturbances is again in accordance with the ionosonde timing.

In contrast with Athens and Nicosia, the other co-located GNSS stations show no clear periodic signatures in any PRN. This discordance between GNSS and ionosonde measurements could be explained by the fact that TEC, as an integrated quantity, mostly captures density fluctuations in the F layer, which gives the highest contribution in terms of ion density. On the other hand, ionosondes can detect waves propagating in the lower altitude layers. Thus, one possible explanation is that further than $\sim 1,500$ km, the seismic waves are not powerful enough to induce atmospheric waves capable of propagating up to F heights, while they still induce density perturbations at lower altitudes, which are detected by ionosondes.




Figure 2.9. dTEC time evolution for six GNSS stations collecting data at 1 s located in Athens, Prague, Matera, Dourbes, Graz, and Nicosia.

The Turkish earthquake provided another clear example of the complex but characteristic patterns of ionospheric disturbances triggered by seismic activity. The results highlight (i) the coexistence of acoustic and Rayleigh-induced TIDs, (ii) the strong anisotropy of their propagation, and (iii) the importance of multi-satellite GNSS observations to disentangle the spatio-temporal variability of these phenomena. These findings further support the view that coseismic TIDs are reliable indicators of seismic processes in the atmosphere–ionosphere system, and that their detailed properties can be used to probe the coupling mechanisms between the solid Earth and the upper atmosphere.

2.3 Severe Convection Linked to Jet-Front over Europe

Between 25 and 29 August 2023, continental Europe experienced a powerful jet-front system and intense convective activity. A meridional jet stream of unusually high velocity (~ 200 km/h) persisted during this period, producing strong vertical wind shear throughout the troposphere. The presence of a jet-front boundary and severe convection provides the conditions for generating AGWs that can propagate upward into the ionosphere. In this case study, during the passage of the jet-front system, geomagnetic activity remained relatively quiet, which helped isolate the contribution from the lower atmosphere rather than from external forcing such as geomagnetic storms.

The study employs a multi-instrument observational approach (GNSS, Ionosondes and Doppler sounders) to detect and characterize MSTIDs, with special emphasis on GNSS measurements of TEC. Using networks of GNSS receivers co-located with ionosondes, one can isolate temporal and spatial fluctuations in dTEC that correspond to MSTIDs.

2.3.1 Data and Methods

As for the previously presented studies, the GNSS-derived GFLC of phase measurements is the starting point of the analysis. For this study, the NeQuick 2 model (Nava et al., 2008) was exploited to estimate the bias, allowing a reduction of the vertical mapping of the calibration bias. More discussion on this observable is available in Section 3.4, where different ionospheric observables for TID studies are introduced and compared. Typically, TID studies merge a large number of observations to remove the complex effect of GNSS satellite movement. However, this study used the BeiDou C05 satellite, a geosynchronous orbit (GEO) satellite visible by GNSS receivers in the European sector.

GEO satellites are observed at a fixed longitude, meaning their IPPs remain fixed in space. Therefore, using GEO satellites removes the influences of satellite-ionosphere relative motion, elevation changing, and TEC spatial gradients on the vTEC time series. Once the GEO vTEC time series is obtained, one needs to remove the background variations unrelated to the passage of the MSTID, such as day-to-day variability and diurnal variations. To do this, Savitsky-Golay filtering was leveraged (Schafer, 2011), a polynomial detrending performed on a sliding window. In this study, the detrending window was set to 1 h and the polynomial order to three. To better highlight the different frequency components related to the MSTID passage, a time-frequency analysis by leveraging the wavelet synchrosqueezed transform was performed (Daubechies et al., 2011).

The GNSS stations considered in this study were closely located to the Průhonice, San Vito, Sopron, and Ebro ionosonde sites. Specifically, the four GNSS stations are named EBRE, GOPE, GRAZ, and MATE and belong to the EUREF permanent GNSS network (EPN; Bruyninx et al., 2019). The data used have a time resolution of 30s, enough to study MSTIDs, usually characterised by periods of tens of minutes.

2.3.2 Results and discussion




Figure 2.10: The top panel shows the GNSS single station TEC, while the middle panel shows the dTEC and its moving standard deviation calculated on a 45-minute window. Bottom panel shows the wavelet spectra at the stations co-located with Průhonice from 25 to 29 August 2023.

The GNSS receiver collocated with the Ebro ionosondes did not show clear MSTID signatures in the dTEC time series during the investigated period. Similarly, the one in San Vito showed signatures only during the daytime on 27 August. Due to the scarcity of signatures related to the tropospheric storm for the EBRE and MATE receivers, their results are not shown here. However, they can be found in Appendix A.2.

Figure 2.10 and 2.11 show the GEO vTEC, detrended dTEC along with its moving standard deviation calculated on a 45-minute window, and the wavelet analysis of vTEC for the GNSS receivers collocated with the Průhonice and Sopron ionosondes, respectively.

Looking at the whole time series, it is possible to note how the 27th of August is characterised by dTEC variations stronger than the ones for the previous and following days. Looking at the dTEC panel, the morning hours (4:00–13:00 UT, highlighted in red) of the 27th of August show the strongest TID signatures, with amplitudes as high as 0.9 TECu for both Sopron and Průhonice receivers. Similarly, the moving standard deviation peaks at 0.46 and 0.38 TECu for Průhonice and Sopron, respectively. By comparing the moving standard deviation with the previous and following days, the maximum value during noon hours almost doubles with respect to background values (always smaller or around 0.25 TECu). Weaker TID activity is also visible in the evening-night hours between the 27th and 28th (shaded blue area

in the dTEC panel), with lower amplitudes than daytime TIDs. This lower amplitude is explained by the fact that TIDs are driven by collisions between the underlying neutral wave and ionospheric particles. Thus, a lower background ionisation will cause a smaller amplitude due to the lower amount of plasma displaced.




Figure 2.11: The top panel shows the GNSS single station TEC, while the middle panel shows the dTEC and its moving standard deviation calculated on a 45-minute window. Bottom panel shows the wavelet spectra at the stations co-located with Sopron (lower part) from 25 to 29 August 2023.

Nevertheless, even if the amplitude is smaller (0.3 TECu), comparing it with nighttime values for the other days considered shows how the amplitude is twice as high for the nighttime between the 27th and 28th of August. The same pattern is shown by the moving standard deviations, with peaks reaching 0.23 and 0.17 TECu for Průhonice and Sopron, respectively. The ionosonde results summarized in Table A.1 of Appendix A.2 show very similar behaviour, where signatures of MSTID (distorted ionograms) are present for the morning hour of the 27th and around midnight of the same day for the colocated Pruhonice and Sopron ionosondes.

The wavelet analysis shown in the bottom panels helps investigate the specific periods activated by the passage of the MSTIDs, as the dTEC does not provide such detailed temporal information. For the Průhonice GNSS receiver, the wavelet spectrogram shows increased activity defined by periods around 40 and 25 min during the daytime (red patch in the dTEC panel). During nighttime (blue patch), the main period is around 30–35 minutes. By comparing the wavelet spectrogram with the RMS values of Doppler shift in different frequencies shown in Figure A.6 in Appendix A.2, it is possible to note how the periods visible in TEC are in

accordance with the Doppler sounder measurements (CDSS). Sopron GNSS TEC shows a similar behaviour, with the daytime showing fluctuations defined by periods around 40 and 25 min and the nighttime between 30 and 40 min.

As MSTIDs are a common phenomenon even during quiet times, observing coherence in both the time and amplitude of the increased TID activity between two different sites (Průhonice and Sopron) and between different instruments confirms that an underlying common phenomenon causes this deviation from the typical day-to-day variability. Moreover, as the front system was mainly localised over Central Europe, no increased MSTID activity for Ebro GNSS stations in Spain confirms that such behaviour is likely connected to the tropospheric event (see Figure A.7).

Chapter 3

Errors in TID detection

Although GNSS-based techniques have proven highly effective for monitoring Traveling Ionospheric Disturbances through the analysis of TEC-based quantities, their reliability depends strongly on data processing choices and observational geometry. In practice, several sources of error can bias the retrieved wave parameters, affecting both scientific interpretation and operational use. The most critical factors are the detrending methods employed to remove TEC background variability, the impact of the chosen elevation cutoff and shell height on the apparent wave period, and the accuracy of NRT observables in reproducing wave amplitudes. Each of these aspects introduces systematic limitations that must be carefully considered when designing detection pipelines and when comparing results across studies. To evaluate these limitations in a controlled but realistic framework, I developed a synthetic-arc methodology in which modeled TID wave fields are superimposed on calibrated GNSS TEC measurements. The methodology and results presented in Chapter 3 have been adapted from the published work Guerra et al. (2024).

3.1 The synthetic arc database

The choice of a synthetic database is driven by the necessity of having a known ground truth (the reconstructed wave field shown in Equation 3.1). That is why having a smoothed real ionospheric background with a superimposed modelled wave has been chosen as a compromise to obtain two synthetic ionospheric datasets (one for the MSTID and one for the LSTID case). Since real measurements are used as the background, I considered two years of data (2009, solar min, and 2014, solar max). The two periods defined by different solar activity are chosen to ensure the presence of different background ionization and day-to-day variability. As the amplitude of the waves is fixed, having different values of background ionization will better represent the real ionosphere. It will thus test detrending techniques in a more realistic scenario.

All the raw GNSS data were gathered by a geodetic station belonging to the INGV Rete Integrata Nazionale GNSS (RING, ring.gm.ingv.it; RING WORKING GROUP, 2016), which is located close to Altamura (AMUR: 40.91°N, 16.60°E, see Figure 3.1), Apulia, Southern Italy. The RING network has been proven to provide high-quality RINEX data to monitor TEC over Italy (Cesaroni et al., 2021). The data used were in the form of 30s RINEX files, and calibrated vTEC was retrieved through the Ciruolo calibration technique (Ciruolo et al., 2007; Cesaroni et al., 2021; Tornatore et al., 2021). Once the calibration and mapping process is completed, a 15000 real vTEC arcs database is available.

As both filtering and detrending techniques are compared, the high-frequency components are removed from the arc database by performing a Gaussian-weighted moving average. Filtering techniques can remove the high-frequency components, while the detrending techniques cannot: this pre-processing should ensure that the lower frequencies induce no bias in the amplitude error.




Figure 3.1: Reconstructed LSTID wave field calculated over Europe on a 0.5×0.5 lat/lon grid. The left panel represents the wave field calculated at t equal to 0, while the right panel represents the wave field at t equal to a half period. The black dot represents the location of the RING station AMUR.

In this study, two TID wave fields are considered; thus, I will have two different arc databases related to LSTID and MSTID. The two wave field parameters are set according to climatological values for the European sector for both MSTIDs (Hernández-Pajares et al., 2006, 2012) and LSTIDs (Hocke & Schlegel, 1996; Borries et al., 2009; Thaganyana et al., 2022), as shown in Table 3.1.

Table 3.1: Wave field parameters (amplitude, wavelength, azimuth, speed, and period) for the two different TID scenarios.

	MSTID	LSTID
Amplitude (A) [TECu]	0.2	0.4
Wavelength (λ) [km]	200	1800
Azimuth (θ) [$^\circ$]	150	195
Speed (v) [m/s]	200	400
Period [min]	16.7	75

Figure 3.1 shows a graphical representation of the LSTID wave model calculated on a $0.5^\circ \times 0.5^\circ$ latitude-longitude grid. The blue and yellow bands represent the TEC depletion/increase caused by the wave field. The two panels show phase values equal to 0 and π . As the considered TIDs are defined by different periods, the Gaussian-weighted moving window size used for pre-processing is set according to 1.33 times the given TID period (22.5 and 100 min for the MSTID and LSTID scenarios, respectively). A 2D plane wave field, presented in Equation 3.1, is used to add the TID to the smoothed background arc databases. I assume that the wave field propagates at a given height equal to the IPP height (HIPP). This assumption is justified by the ionospheric plasma conditions at mid-latitudes, where most of the ionized matter is localized in a thin layer; thus, I assume that the TID propagates in the thin ionospheric layer (Mannucci et al., 1998). Moreover, HIPP is usually set by the user according to realistic values of the height of the F2 ionospheric layer (Astafyeva et al., 2022), which further justifies my assumption. I then model the TEC variation (dTEC) over space (x, y) and time (t) induced by a wave like the following:

$$dTEC(x, y, t) = A \sin\left(\frac{2\pi x}{\lambda} \cos\theta + \frac{2\pi y}{\lambda} \sin\theta + \frac{2\pi vt}{\lambda}\right) \quad (3.1)$$

where A is the wave amplitude, λ is its wavelength, θ is the azimuth of propagation, v is the TID speed, x and y are latitude and longitude converted to km of arc over the Earth's surface.

3.2 Detrending techniques comparison

Once the arc database is generated, the following step is to define the detrending techniques that are applied to all the wavy TEC arcs obtained. To extract wavy features from the vTEC signal, six different detrending/filtering techniques were applied to the reconstructed arcs:

1. Moving average (MA).
2. The multi-order numerical difference (DD).
3. Third-order Savitzky-Golay filter (SGOLAY).
4. Polynomial detrending (POLY).
5. Finite Impulse Response Band-Pass Filter (BUTF).
6. Fast Iterative Filtering band-pass filter (FIF).

While techniques 1–5 are well-known and widely applied in many scientific fields, the FIF-based band-pass filter requires further explanation. FIF is a novel signal decomposition technique capable of decomposing non-stationary and non-linear signals into simple oscillatory (around zero) components called intrinsic mode functions. To leverage FIF as a band-pass filter, one needs to sum all modes defined by frequencies inside the range of interest at every time step. FIF has proven to be particularly suitable for analyzing ionospheric data (see, e.g., Ghobadi et al., 2020; Spogli et al., 2021, 2019; Urbar et al., 2022; Verhulst et al., 2022). As FIF is based on the Fast Fourier Transform, it assumes periodicity at the boundaries of the signal under study. Stallone et al. (2020) proposed an extension approach that drastically reduces boundary errors without introducing spurious modes in the decomposition. Thus, the signal extension increases the signal length with a half-sample antisymmetric extension to mitigate the boundary effects.




Figure 3.2: The top panels (a, b, c) show the different arcs of TEC considered in the reconstruction process under the MSTID scenario. Specifically, panel (a) shows the output of the GG calibration algorithm (vTEC), panel (b) shows the moving average smoothed vTEC arc and panel (c) shows the smoothed arc summed with the TID wave model (blue line of panel (d)). The bottom panel (d) instead shows the detrended (FIF-based band-pass) version of the synthetic arc (orange line) along with the TID wave model calculated at the IPP locations (blue line).

A preliminary study was conducted specifically for the SGOLAY and POLY techniques to determine their main parameters, as these require more tuning than simply fixing the period band or sliding window size. In this analysis, shown in Appendix B, I evaluated the ability of different polynomial orders to reproduce wave-like features in terms of amplitude and correlation coefficient errors. Based on these results, the POLY degree was set to 5th for the LSTID scenario and 10th for the MSTID one, as these provided the best compromise between amplitude and correlation errors. For SGOLAY, the best performance in both scenarios was obtained with a 2nd-order polynomial. Regarding the sliding window, the optimal settings were three times the period for LSTIDs and twice the period for MSTIDs. Nevertheless, since the true TID period cannot be known a priori in real applications, multiples of the wave model period were not used. Instead, the sliding window size was set relative to the MA one: twice the MA window for LSTIDs and 2.5 times the MA window for MSTIDs. This approach ensures consistency across methods and prevents an unfair advantage from assuming prior knowledge of the TID characteristics.

For the other techniques, parameters were set according to standard LSTID and MSTID characteristics, with the period being the most impactful factor. The BUTF and FIF bands were defined as 10–40 min for MSTIDs and 45–90 min for LSTIDs. The MA was applied using a 60-min window for LSTIDs and a 30-min window for MSTIDs. The DD time separation for MSTIDs was fixed at 300 s, following values commonly used in climatological studies exploiting GNSS-derived TEC (Hernández-Pajares et al., 2006, 2012). For LSTIDs, where no established DD setting was available, a time separation of 30 min was used, equal to half the MA window size. Figure 2 shows the reconstruction and detrending process for a single arc under the MSTID propagation scenario. Specifically, Figure 3.2 shows an example of the different steps of the reconstruction process for the MSTID scenario. Calibrated $vTEC$, obtained from the Ciralo calibration (Figure 3.2a), is the real ionospheric background. Instead, Figures 3.2b and 3.2c are the smoothed background $vTEC$ before applying the wave field, and the sum between the latter and the TID model is calculated at each IPP by applying Equation 3.1. Figure 3.2d shows the detrended (FIF-based band-pass) version of the synthetic arc along with the TID model calculated at the IPP locations.

Once the two synthetic arcs databases had been detrended with the six techniques mentioned in the previous section, it was possible to investigate the amplitude error (AME) and time domain error (TDE) induced by the wave extraction technique. Note that TDE refers to time-shifting or contraction/extension in time of the decomposed signal, an artefact that prevents accurately deriving the TID speed from interferometry-like techniques. For the time domain error (TDE), I considered one minus the normalized cross-correlation at lag 0. According to this, AME and TDE are expressed as:

$$AME = x_r - x \quad (3.2)$$

$$TDE = 1 - \frac{x_r - x}{\sqrt{(x_r \cdot x_r)(x \cdot x)}} \quad (3.3)$$

where x represents the detrended arc and x_r the reconstructed wave field. Because of the different nature of the two errors, while the number of AMEs is equal to the sample number, there is only one TDE for each arc. Finally, I looked at the error distribution and statistics of AME and TDE to better investigate the general picture of the error.

First, the amplitude error induced by the different detrending techniques is investigated. Figure 3.3 shows the overall AME statistics for the six techniques. From the Cumulative Distribution Functions (CDFs) in the top panels, one can quantify how likely it is to expect an AME error smaller than the one of interest. For example, it is possible to know that 80% of the

time, the AME is smaller than the corresponding x value for the technique of interest by fixing a cumulative $P(x)$ of 80%. Looking at the CDF under the LSTID scenario and fixing a cumulative probability of 80%, one can observe that FIF and SGOLAY are the best techniques AME-wise, with an x value of 0.125 TECu (around 35% of the original amplitude). In contrast, the numerical difference technique (DD) performs the worst (almost 0.5 TECu, more than 100% of the LSTID amplitude). This pattern is confirmed by the bottom-left panel, where one can see that FIF and SGOLAY show the smallest percentiles. Now, focusing on the MSTID scenario (top-right panel), one can note that the general picture is similar to the LSTID one, with FIF and SGOLAY being again the most reliable techniques (80% of the AME are around 0.05 TECu, 25% of the original MSTID amplitude). Instead, DD is the least accurate (0.1 TECu, 50% of initial amplitude). Regarding the other three techniques, one can remark that none clearly outperforms the others. In particular, the blue and purple lines for the LSTID scenario of Figure 3.3 almost perfectly overlap. Before looking into the correlation error, one has to stress that the DD technique, which is the one performing worst, is not meant to extract the amplitude reliably since it is supposed to amplify up to $2\times$ the wavy feature defined by the period set as the main technique parameter (time separation).




Figure 3.3: Amplitude error statistics for the six techniques considered. MSTID (right column) and LSTID (left column) parameters correspond to the ones presented in Table 1; the ionospheric thin shell height was set to 350 km. The top panels show the CDFs of AME for the different techniques (each line corresponds to the CDF of a given technique). X-axis shows the AME error in TECu, while the Y-axis corresponds to the probability $P(\text{AME})$ of having an error equal to or smaller than the corresponding X value. The bottom panels, instead, present the median along with their related percentiles (5, 16, 84, and 95) for the different filtering/detrending techniques.

The violin plots shown in Figure 3.4 are used to investigate the correlation error statistics. The violin plots effectively present the probability distribution and the overall statistics (mean, median, and interquartile) in a compact way. First, focus on the LSTID scenario (left panel). As for the AME, FIF and SGOLAY perform best in preserving the original wavy feature in the frequency domain, with FIF showing slightly better performance. The mean and median for the FIF distribution (cyan) are the smallest, and the same applies to the interquartiles. In addition, thanks to the violin plot, it is possible to observe that FIF has the highest amount of

correlation errors close to zero. In addition, FIF also shows the thinnest distribution for high correlation error values, meaning it is the least prone to large frequency distortions. Before investigating the MSTID case, one has to stress that every distribution is highly asymmetrical, as confirmed by the significant difference between the mean and the median. Focusing on numerical values, FIF, MA, and DD show a comparable correlation error mean under the MSTID scenario. In contrast, in the LSTID case, FIF and SGOLAY outperform the other techniques. Considering both TID scenarios, I can conclude that FIF and SGOLAY are the best techniques again, ranking first and second for both the LSTID and MSTID scenarios.




Figure 3.4: Violin plot for the correlation coefficient error. The X-axis shows the correlation error, while each violin (colour-coded according to different techniques) is a box plot where the shape represents the probability distribution of the correlation error for the given technique. The black diamond and line are, respectively, the median and mean of the correlation error distribution, while the shaded area represents the interquartile range. The left panel represents the LSTID scenario, while the right one is for MSTID.

To conclude, considering the findings derived from Figures 3.3 and 3.4, one can conclude that, among the techniques considered, FIF and SGOLAY have proven to be the best at extracting wavy features from post-processing observables, assuring the best performance in estimating the amplitude while reliably preserving the wave frequency. In addition, concerning the computational speed of the techniques, SGOLAY would be the technique of choice since it is around 300 times faster than FIF. Nevertheless, some considerations about the DD technique must be made. As previously stated, the DD technique is designed to amplify features up to $2\times$, with the highest degree of amplification for features defined by a frequency equal to the one corresponding to the moving window. Thus, including the DD technique in the AME ranking is somehow unfair because it is specifically designed to act on the amplitude of the feature of interest.

3.3 Elevation cutoff impact

Once AME and TDE caused by the different detrending algorithms have been ranked, the focus is moved to the effect on the wave period induced by the underlying nature of the

measurement. As said, the TID is assumed to propagate in a thin shell layer; thus, different heights of such a shell will be related to different IPP speeds (Hernández-Pajares et al., 2006). Since the IPPs are moving across the TID, the measured wave will suffer a Doppler effect directly related to the IPP speed and, thus, to the satellite elevation and height of the ionospheric shell (e.g., IPP velocity increases for lower elevation angles and higher shell heights). Savastano et al. (2017) show that, by setting an incorrect ionospheric height, the TID estimated speed differs from the real one. Despite that, TID speed is usually derived from interferometry techniques (Afraimovich et al., 1998; Maletckii and Astafyeva, 2021), and thus, the main issue would be the difficulty in correlating waves that show different periods due to the Doppler effect. Detrending techniques are unnecessary to investigate this phenomenon because the Doppler effect depends only on the relative velocity between the IPP and the TID. Therefore, since the TID velocity vector is set and remembering that the TID is assumed to propagate precisely at the height of the ionospheric shell, I only need to derive the IPP velocity vector for each data sample. Once the IPP velocity vector is evaluated, starting from the consecutive IPP coordinates, it is possible to calculate the Doppler effect thanks to the following equation:

$$f' = f \left(1 + \frac{v_{rel}}{|\vec{v}_{TID}|} \right) \quad \text{where:} \quad v_{REL} = \frac{\vec{v}_{TID} \cdot \vec{v}_{IPP}}{|\vec{v}_{TID}|} \quad (3.4)$$

where v_{REL} is the TID-IPP relative velocity, $\vec{v}_{TID}$ is the TID velocity vector, f is the original wave frequency, and f' is the apparent frequency. As the IPP velocity strongly depends on the satellite elevation and IPP height, I investigated the Doppler-induced error statistics for different values of the two parameters. Finally, it is possible to infer the expected error given the ionospheric shell height and the elevation cut-off angle through statistics.




Figure 3.5: Cumulative distribution functions (CDF) of the Doppler error for different TID scenarios, heights of the IPP, and elevation masks. Plots on the right/left correspond to MSTID/LSTID, while top/bottom plots are for an IPP height of 250/350 km. The differently coloured lines show the CDF for different elevation masks, where blue, red, yellow, purple, and green are 20, 30, 40, 50, and 60 degrees, respectively. The X-axis represents the percentage period error, while the Y-axis shows the cumulative probability.

Figure 3.5 shows that the CDFs of the period error change depending on the TID scenario, IPP height, and elevation cut-off. The CDFs of period error for the LSTID scenario and an IPP height of 250 km are shown on the top left. The first clear result is that the lower the elevation cut-off, the higher the expected period error on average. Specifically, the 80% cumulative $P(x)$ corresponds to a period error of 9, 11, and 15% for 60-, 40- and 20-degree elevation cut-off, respectively. Now, one can compare the scenario (LSTID, 250 km) with the bottom-left panel, which is the same plot but for an increased IPP height (350 km). Thanks to this comparison, one can quantify how much the increased IPP height influences the extracted period.

Specifically, considering the 80% cumulative $P(x)$, I obtain period errors of 20, 14, and 12% of the original period for 20, 40, and 60 degrees, respectively. As expected from theory and confirmed by the statistics, a higher IPP height causes the period error to increase on average.

Looking at the MSTID results, the first clear observation is that the period error almost doubles on average compared to the LSTID case. This behaviour was expected since the Doppler-induced period error is directly connected to the TID speed, as seen in Equation 3.4; therefore, the slower the TID, the higher the Doppler effect. Numerically speaking, the same 80% cumulative $P(x)$ related to the same elevation cut-offs as for LSTID are 21, 22.5, and 33% for an IPP height of 250 km and 32, 34, and 53% for 350 km under the propagating MSTID scenario. Thus, the slower MSTID shows period errors that are doubled with respect to the faster TID, a ratio that corresponds to the two TIDs speed ratio, 2 (see Table 3.1). Moreover, looking at the MSTID scenario, one can observe that the CDFs show a flex point, a feature that is not visible in the LSTID ones. This feature is likely due to the nature of the IPP trajectories, which are directly related to the GNSS orbital parameters. Due to the MSTID speed being comparable to the IPP one and because of the azimuth being fixed, there might be a significant difference in the error statistics for different orbital planes. Thus, if one consider the CDF as a sum of the error CDF for each orbital plane, one might be largely shifted towards larger errors, causing such flex points. The same does not apply to LSTID because the wave speed is always larger than the IPP one, limiting the degree of error induced.

3.4 Near-Real-Time Observables

In this section, I present two different NRT-derivable observables, and I propose an evolution of one of the two that should ensure a higher level of accuracy in the estimation of the extracted signal amplitude. As of now, the first two observables have been used extensively for NRT and post-processing applications, and many examples are available in the literature (Astafyeva et al., 2022; Brissaud & Astafyeva, 2022; Ravanelli et al., 2021; Maletckii & Astafyeva, 2021; Martire et al., 2022). To evaluate the performance of such observables, a statistical analysis is performed by comparing the signal extracted from the three NRT observables with the one derived from a widely used post-processing technique proposed by Ciruolo et al. (2007). The first NRT-derivable observable considered is the sTEC derived from phase measurements, which can be retrieved in real-time through the Networked Transport of RTCM via Internet Protocol (NTRIP). sTEC can be computed starting from phase measurement by applying Equation 1.20. Once the phase-derived sTEC is obtained, the TEC value of the first data point is subtracted from the entire time series to limit the unknown bias characteristic of phase measurements. By detrending this phase-derived sTEC, it is possible to extract the wavy features correctly in the time domain. However, its reliability in extracting amplitudes is not comparable due to the slanted nature of the observable. Therefore, one must verticalize it to extract the amplitude reliably (especially for low elevations where there is the least accordance in the sTEC and vTEC magnitude). The second observable is thus obtained by converting $sTEC_{PH}$ to $vTEC_{PH}$ through the mapping function shown in Equation 1.21. Note

here that when applying the mapping function to $s\text{TEC}_{\text{PH}}$, the bias will not be constant anymore due to the dependence on the elevation angle. By looking at the mapping function, it is possible to see that the induced amplitude error will be minimal for high elevation. That is why I propose a third and novel approach, where an initial guess is performed by estimating the initial value of the $s\text{TEC}_{\text{PH}}$ time series through an ionospheric model. One of the main constraints of the ionospheric model is that it must be extremely quick (hundredths of seconds per TEC arc) in computing the initial bias. This constraint is extremely important for NRT applications. Once the initial guess is gathered from the model, the $s\text{TEC}$ is verticalized by applying Equation 1.21; in principle, thanks to this approach, the amount of error induced by the verticalization of the initial bias should be lower, with a minimum for high elevations and a maximum for low ones (this observable will be called $v\text{TEC}_{\text{Ne}}$). The ionospheric model used is NeQuick2 (Nava et al., 2008), which was selected based on its computational lightness and speed at calculating the TEC (a few milliseconds per arc), given the receiver and satellite coordinates.




Figure 3.6: The top panel (a) shows the different NRT-like observables studied and the satellite elevation. The solid blue line is an example of a reconstructed arc, just like the one in Figure 2c. The orange line instead is the satellite elevation. The bottom panel (b) focuses instead on the detrended version of the four observables presented in panel (a), which share the same colour code. dTEC peaks corresponding to low, medium, and high elevations are zoomed to better visualize the slight differences between the detrended observables.

Figure 3.6 shows four different observables belonging to the same arc, three of which have been obtained from the three different NRT techniques under the MSTID scenario, and the fourth is the calibrated $v\text{TEC}$, which, due to its post-processing nature, is considered the ground truth. Focusing on the NRT observables, the blue dashed line represents $s\text{TEC}_{\text{PH}}$. Following, the blue dotted line is $v\text{TEC}_{\text{PH}}$ (so, the blue dashed one times the mapping function), while the dash-dotted blue line is $v\text{TEC}_{\text{Ne}}$. At first glance, one can notice that the detrended arcs (Figure 3.6b) are almost identical for the values close to the middle arc time. This divergence is due to the elevation (orange line in Figure 3.6a) decreasing while reaching the arc boundaries since every continuous observational arc (no loss of lock) starts when the satellite rises and ends when the satellite disappears below the horizon. Moreover, it is worth

noticing that Figure 3.6b shows that the wave period changes throughout the arc duration, with the first half decreasing while the second half increases. As for the Doppler error section, here, the reliability of the six detrending techniques is not of interest. The objective is to identify which NRT observables best resemble the results obtainable in a post-processing study. Therefore, the detrending technique is fixed for both scenarios. The technique used is SGOLAY because it proved to be the best trade-off between computational cost and error performance. Once the NRT arcs are detrended, the AME is computed. Note here that the different NRT approaches do not cause different effects on the TID period. However, they amplify or attenuate the signal amplitude depending on the sign of the mapped bias. That is why the NRT error analysis only covers the AME while disregarding the TDE.




Figure 3.7: Amplitude error statistics for the dTEC derived from NRT observables. The top panels show the CDF of the amplitude error, with each line representing a different observable (blue, red, and yellow for sTEC_{PH}, vTEC_{PH}, and vTEC_{Ne}, respectively). The X-axis shows the amplitude error, while the Y-axis shows the cumulative probability. The bottom panels show the median and the 18–84 and 5–95 percentiles for the different techniques. Here, the Y-axis represents the amplitude error in TECu. The plots on the left and right correspond to LSTID and MSTID, respectively.

Figure 3.7 shows the median/percentiles of the three observable CDFs for the two TID cases. Looking at the CDFs, it is possible to remark that vTEC_{Ne} outperforms the other two observables in both the MSTID and LSTID scenarios. Focusing only on the LSTID CDFs, one can notice that sTEC_{PH} is way less accurate than the others. Specifically, the 80% cumulative P(x) amplitude error for vTEC_{Ne} is almost negligible (0.015 TECu, comparable to the noise level for phase measurements), while for sTEC_{PH} and vTEC_{PH}, it is respectively around 0.16 and 0.065. This difference between the three observables is confirmed further by the median and percentiles shown in the bottom-left panel of Figure 3.7. Contrarily, the performance of the two verticalized observables is similar under the MSTID scenario. Numerically speaking, fixing at 80% the cumulative probability, one gets an AME of 0.075, 0.005, and 0.0025 for sTEC_{PH}, vTEC_{PH}, and vTEC_{Ne}, respectively. Furthermore, by looking at percentiles, it is possible to notice that the 16–84 percentiles for vTEC_{PH} and vTEC_{Ne} are almost identical (both

around 0.001 TECu, symmetric distribution), while the ones for $sTEC_{PH}$ are clearly higher (0.05 TECu).

This difference means that $sTEC_{PH}$ is way more prone to large error values, while $vTEC_{Ne}$ and $vTEC_{PH}$ manage to reduce their impact/number. To conclude, considering both scenarios, one can confidently affirm that $vTEC_{Ne}$ is a reliable choice for extracting the TID amplitude in an NRT application. In addition, by comparing the $vTEC_{Ne}$ amplitude error performance with the one related to the best technique, FIF, it is possible to see that the error induced by working in an NRT scenario is smaller than the AME induced by the detrending process. Specifically, the $vTEC_{Ne}$ AME is 10% and 5% of the FIF error for the LSTID and MSTID scenarios, respectively.

Chapter 4

LSTID Nowcasting

While case studies such as the ones shown in Chapter 2 highlight the diverse sources and manifestations of TIDs, operational applications require methods capable of identifying and characterizing them systematically over long periods or in real time. LSTIDs, in particular, represent a significant challenge for space weather monitoring: they are recurrent during geomagnetic disturbances, affect satellite navigation and communication systems, and their climatological behavior reflects the coupling between high-latitude energy input and mid-latitude ionospheric variability. Building a reliable database of LSTID events and tools for autonomous recognition is thus a prerequisite for scientific analysis and near-real-time applications. This section introduces an autonomous detection algorithm designed to identify LSTID signatures in TEC data derived from GNSS observations. The method operates indistinctly in both geomagnetically quiet and disturbed conditions, and it is optimized to extract key LSTID parameters such as the horizontal apparent speed (along the N-S direction), period, and wavelength from dTEC keograms (i.e., representation of the dTEC as a function of time and latitude). To assess its performance, the algorithm was applied to the period from January 2021 to December 2024 over the European area. This enabled a comprehensive statistical analysis of LSTID occurrence, allowing comparison with known climatological patterns reported in the literature.

Tsugawa et al. (2004) visually inspected dTEC maps over Japan to identify LSTID signatures from April 1999 to December 2002. Ding et al. (2008) performed a similar study by visually inspecting dTEC maps of North America for major storms from 2003 to 2005. Lastly, Borries et al. (2009) performed an automated analysis of LSTIDs derived from dTEC maps and keograms of the European area for 37 geomagnetic storms. The results of the three studies are summarized in Table 1 in terms of LSTID speed and period.

Table 4.1: Summary of the characteristics and results of statistical analysis performed in previous studies.

	Geographical Area	Period considered	LSTID Speed [m/s]	LSTID Period [min]
Tsugawa et al., 2004	Japan	April 1999 – December 2002	475 ± 171	80 ± 29
Ding et al., 2008	North America	Major magnetic storms from 2003 to 2005	300	108
Borries et al., 2009	Europe	Major magnetic storms from 2001 to 2007	684 ± 195	59 ± 11

Previous studies that leveraged TEC required visual data inspection or focused only on major geomagnetic storms. In contrast, the technique presented here no longer requires visual inspection and was applied to identify LSTIDs under quiet and disturbed conditions. The automatization of the technique raises some difficulties, such as the robustness against false positives; on the other hand, it does not include the visual inspection bias and also allows identifying LSTIDs that are not directly related to geomagnetic storms, as it was applied to all the days in the considered period.

4.1 The detection process layout




Figure 4.1: Flowchart of the LSTIDs detection method.

Figure 4.1 outlines the automated LSTID detection process. It begins with Data Gathering and Station Selection, where dual-frequency GNSS data are retrieved from an optimized subselection of stations belonging to the EUREF network (Bruyninx et al., 2019). Next, in Detrended TEC and Daily Keograms generation, phase measurements are used to derive slant TEC, which is then verticalized and detrended with a Savitzky-Golay filter. These dTEC values are organized into latitude-time keograms for further analysis. The Region of Interest (ROI) Identification Process identifies potential LSTID features using image processing. Binary masking and segmentation techniques, including Bayesian Gaussian Mixture Models (BGMM), isolate coherent positive and negative deviations. The ROI Feature Extraction step estimates propagation speed and other parameters using the Random Sample Consensus (RANSAC) algorithm to fit lines to the ROIs, providing metrics like start time, latitude range, and velocity. Following this, the algorithm joins opposite ROIs to retrieve LSTID occurrences, filtering out false detections based on criteria like expected speed, period, and latitudinal overlap.

4.2 Data gathering and station selection with K-Means

The detection technique leverages dTEC measurements to identify LSTID features in latitude-time images. Image retrieval starts from uncalibrated verticalized TEC derived from GNSS phase measurements. Thus, the first necessary step is to retrieve dual-frequency phase measurements from the area of interest. Multi-constellation RINEX files in 30s format are gathered from the EUREF Permanent GNSS Network receivers (Bruyninx et al., 2019). Keograms are made by selecting observations pertaining to a certain longitudinal range, and data are collected from receivers located between 10°E and 25°E and in the latitudinal range 30°N to 60°N.




Figure 4.2: EUREF GNSS permanent network (black dots) and the selection of GNSS stations performed by applying K-Means clustering (red dots). The blue shaded area represents the geographical region considered for generating the keograms.

To reduce computational time and manage data size when deriving keograms, the original EUREF GNSS network is downsized by selecting a subset of stations. Instead of randomly choosing stations, a clustering algorithm was employed to ensure adequate spatial coverage using the K-Means algorithm (MacQueen, 1967; Lloyd, 1982). K-Means is an unsupervised learning algorithm that groups data into a predefined number of clusters, 100 in this case. A 100 stations provide a good compromise between the IPP coverage of the area of interest and the density of IPP per pixel of the keogram image. More details are provided in Appendix C.1, where IPP density keograms are shown for different numbers of stations. The station selection process begins by randomly selecting 100 stations as initial cluster centroids. Each GNSS station from the full network is assigned to the nearest centroid, forming clusters. The centroids are iteratively recalculated as the mean position of the stations in their respective clusters. This process continues until the assignments stabilize, meaning no stations change clusters during an iteration. The primary goal of this clustering was to eliminate redundant data from densely covered regions while preserving the spatial distribution necessary to capture ionospheric variations. To mitigate the impact of randomly initialized centroids, the K-Means algorithm is

repeated 5000 times with different starting points. The solution minimizing the average distance between stations and their cluster centroids is selected among these iterations. This ensures a robust and well-optimized distribution of selected stations.

Additionally, each year, the station selection is repeated to account for temporal changes in the network configuration. For example, the selected stations for 2024 are shown in Figure 4.2. The keogram generation process operates by analyzing data one day at a time. Thus, to remove the impacts of measurements gathered around 00:00 and 24:00 UT of the given day, the last/first 2 hours of the previous/following day are included in the process.

To analyze the performance of the detection algorithm outside Europe, it was also applied to two different geographical regions for the major geomagnetic storm of 11-12 August 2024. The data used belong to GEONET (Sagiya, 2004) for Japan and CORS (National Geodetic Survey, 2025) for North America. As GEONET and CORS comprise thousands of GNSS stations, the same station selection process was applied. Figure C.3 shows the spatial distribution of the 100 stations selected for the two regions.




Figure 4.3: GNSS stations considered for the performance analysis of the detection algorithm for two different geographical regions. Red dots show the 100 stations selected from the CORS network, while blue dots are the GEONET ones.

4.3 Keogram generation

Once the RINEX files for the optimized network have been collected, the GFLC shown in Equation 1.20 is calculated for GPS, GLONASS, BeiDou and Galileo. As for the previously presented studies, the ionosphere is considered a 2D layer at a height of 300 km. Then, the GFLC is verticalized by applying Equation 1.21 after setting the initial observation point to 0 TECu. According to the study presented in section 3.4, it is proven that this allows the best compromise between amplitude error after the observable is detrended and computational speed. Only observations with an elevation higher than 20° are considered at this stage to reduce the impact of cycle slips related to multipath. Once the TEC observable is obtained, a detrending technique is applied to highlight periodic features ascribable to LSTID. Again, following the findings presented in section 3.2, a Savitzky-Golay filter was selected (Schafer, 2011), as it provides the best compromise between detrending errors and computational weight.

Specifically, a second-order filter with a time window of 95 minutes was used. The rationale behind choosing the 95-minute window is as follows: previously, I demonstrated that optimal performance was achieved with a 120-minute window when identifying a modeled LSTID with a period of 75 minutes. According to Borries et al. (2009), one can expect European LSTIDs to have, on average, a period of approximately 60 minutes. Therefore, by maintaining the same window-to-period ratio, I deduce that the optimal window would be 95 minutes. This choice of window size imposes a limitation on the LSTIDs that the technique can identify: a larger time window would enable the identification of LSTIDs with longer periods. However, since one can expect most LSTIDs to have periods shorter than approximately 90 minutes, I accept this trade-off. Moreover, the technique could be applied with different parameters (such as adjusting the detrending window) to allow targeted studies of longer-period LSTIDs. In addition to Savitzky-Golay, a moving average filter is applied on a time window of 15 minutes to reduce the high-frequency noise. As detrending techniques are known to suffer from border effects, a second elevation cutoff of 30° is applied.

Starting from the dTEC measurements, it is possible to create the keogram by using the data falling in the latitude and longitude range of 37.5°N to 57.5°N and 10°E to 25°E . After this, the median of all measurements falling in each $0.05^\circ/30\text{s}$ bin is calculated and assigned to the relative latitude/time pixel. The data from the last/first two hours of the previous/following day are removed, thus obtaining a keogram free of boundary effects. For example, Figure 4.4a shows the keogram image obtained for the 15th of October 2024. The final step before the ROI identification process is an interpolation of dTEC on a grid equal to the keogram image size. By doing this, the dTEC values for pixels not covered by observations are filled with linear interpolation, obtaining a processable 2D image. Once the database of keogram images is generated for January 2021 to December 2024, it is possible to proceed with the ROI identification process.

4.4 LSTID autonomous identification algorithm

Being wave-like structures, LSTIDs can be identified as oblique bands of aligned positive or negative dTEC. Therefore, to identify an LSTID in a keogram, the first step would be identifying connected regions characterized by alternating negative/positive amplitude. The regions that are potential LSTID features are hereafter referred to as ROIs. Before applying the ROI identification algorithm, some preprocessing is performed.

First, a local average downsampling of size 4 in time is performed. As the original keogram image is 401×2880 pixels, the downsampled one will be 401×720 . This preprocessing has two major advantages: the first is that ROI identification is faster due to the reduced image size, and the second is that the averaging step reduces image noise, which is detrimental for the instance segmentation algorithm that will be later described and applied. A step-by-step process is used to identify the ROIs in the preprocessed keogram images. First, all the values with an absolute amplitude higher than 0.015 TECu are used to create one negative and one positive binary mask. The two masks are defined as follows:

$$B_+(x, y) = \begin{cases} 1, & M(x, y) > T \\ 0, & M(x, y) \leq T \end{cases} \quad (4.1)$$

$$B_-(x, y) = \begin{cases} 1, & M(x, y) < T \\ 0, & M(x, y) \geq T \end{cases} \quad (4.2)$$

here B_+ and B_- are the two binary masks, M is the preprocessed matrix, and T is the amplitude threshold set to 0.015 TECu . The impact of setting a different dTEC threshold is shown in

Appendix C.2. ROIs are clusters of spatially connected pixels within the binary masks that satisfy certain conditions. Two pixels are considered connected if they satisfy a 2-connected adjacency criterion, meaning they can be linked through a path consisting of one or two orthogonal (horizontal or vertical) steps. The connected regions with a latitudinal range smaller than 10 degrees and a solidity lower than 0.35 are discarded. Here, solidity refers to the ratio between the pixels belonging to the connected region and the total number of pixels inside the convex hull. The impact of the solidity parameter is shown in Appendix C.3. The connected regions characterized by a minor axis of the fitted ellipse longer than 25 pixels (50 minutes) are checked for the size of their largest concavity. A covariance matrix is built from the pixel coordinates to fit the ellipse. This matrix captures how the pixel positions vary along the horizontal, vertical, and diagonal directions. The ellipse parameters are then derived by performing an eigenvalue decomposition of this covariance matrix. The eigenvectors of the matrix indicate the orientation of the ellipse, while the eigenvalues correspond to the squared lengths of the major and minor axes. The axis with the larger eigenvalue represents the direction in which the region is most elongated (the major axis), and the smaller eigenvalue corresponds to the minor axis. To evaluate the size of the largest concavity, the convex hull is compared with the alphashape-derived concave boundary (Edelsbrunner et al., 1983) to retrieve the region's concavities. If the largest concave area is bigger than 700 pixels, the connected region undergoes a second step of instance segmentation using the BGMM clustering algorithm applied to the normalized coordinates of points. Other clustering algorithms have been evaluated, with BGMM showing the best performance. However, the comparison between different techniques is not included in this thesis. The BGMM (Roberts et al., 1998) is a probabilistic model that assumes data points come from a mixture of K Gaussian distributions. The parameters of the Gaussian distributions (such as means and covariances) are treated as random variables in BGMM, and Bayesian inference methods are used to estimate their values. The number of clusters K is gradually increased, and the algorithm stops when the performance improvement from increasing K by one is lower than in the previous step. This unsupervised approach ensures that the segmentation is not overly fragmented and helps avoid creating too many clusters when the additional subdivision does not improve the model significantly. The irregular region is then divided into K subregions by disconnecting the pixels belonging to the intersection of each couple of the K -connected subregions. The new K subregions are checked again based on the solidity and latitudinal elongation criteria. This secondary segmentation step tackles the scenario typically related to two consecutive LSTID bands of the same magnitude connected by noise at higher latitudes that appear as a single region. More details on this process and on how the area of the largest concavity parameter impacts the technique detections are provided in Appendix C.5. All the connected regions obtained at this stage are now filtered based on the inclination of the major axis of the fitted ellipse. The connected regions are discarded if the inclination is smaller than 75° (corresponding to a meridional speed of ~ 150 m/s). All the remaining connected regions are now considered ROIs. The output of the ROI identification process is shown in Figure 4.4b. Figure 4.4 shows how the technique can detect most of the band-like structures for the 15th of October 2024. Figure 4.4b shows a large number of ROIs in different colors; nevertheless, some are not ascribable to LSTID occurrences. For instance, feature 5 is likely related to the passage of the solar terminator, while features 6, 7, 8 and 9 are likely connected to daytime MSTIDs. These are discarded in the following steps of the detection algorithm, where ROI pairing is performed.




Figure 4.4: An example of the result of the ROI identification step for the 15th of October 2024. Panel (a) shows the keogram image colored according to dTEC amplitude, with red/blue area representing positive/negative deviations from the background trend. Panel (a) also shows the output of the RANSAC fitting algorithm in gray, a straight line that models the North-South propagation of TEC disturbances. Panel (b) shows the different ROIs with their pixels in different colors.

Once the ROIs of a given image are identified, it is possible to derive some of their parameters. Since a keogram is a space-time representation of dTEC, the slope of a given ROI corresponds to the meridional component of the horizontal velocity. To derive the slope of ROIs, a straight line is fitted to the ROI pixels with the RANSAC algorithm (Fischler & Bolles, 1981), which excludes outliers in the fitting process to ensure a higher level of robustness with respect to other fitting techniques.

RANSAC is an iterative algorithm employed to derive the parameters of a given model. In this case, the model considered is a weighted linear regression fit (Draper & Smith, 2014). The weights employed in the fitting process are set according to the pixel dTEC intensity. First, a random subselection of the points S is chosen and the model parameters θ_S are estimated:

$$\theta_S = \operatorname{argmin}_{\theta} \sum_{x \in S} \operatorname{loss}(x, \theta) \quad (4.3)$$

Then, the consensus is evaluated by calculating the number of points x_i of the full set X characterized by an error lower than ϵ :

$$C_s = \{x_i \in X : \text{loss}(x_i, \theta_s) < \epsilon\} \quad (4.4)$$

After the consensus is evaluated for a given number of subselections, the set of θ_s that maximizes consensus is chosen. Thanks to RANSAC, it is possible to derive the ROI start time, end time, latitude range, slope/velocity and the fitting quality metrics (mean absolute error, root mean squared error, and R2 score). As shown in Figure 4.4a, RANSAC can correctly fit the middle line of all the band-like features related to LSTIDs.

According to my assumptions, LSTIDs are wave-like structures propagating at a velocity above 300 m/s with periods longer than 30 minutes. Thus, ROIs cannot be considered LSTIDs, as a single ROI represents only one positive or negative variation. Therefore, the ROIs obtained from daily keograms are merged into a single database and undergo further filtering processes to ensure that they are related to LSTID occurrences and not to noise or other ionospheric processes that are not of interest. A given ROI is considered an LSTID ROI if an ROI of the opposite sign exists and meets the following requirements:

- The difference in time propagated at 47.5°N is smaller than 45 minutes
- The speed is in the range of 300 to 1500 m/s
- The slope angle difference is smaller than 2°
- The latitudinal overlap is at least 7.5°

All the ROIs that meet the requirements are considered LSTID signatures. The high threshold for the velocity range (1500 m/s) was chosen due to the high sensitivity of the slope-based velocity retrieval, as working with angles close to 90°, the tangent grows dramatically for small angle differences.

Moreover, I acknowledge the fact that all potential LSTIDs defined by a period longer than 90 minutes are discarded. However, as specified in the previous sections, the technique parameters are not considered to be fixed. In contrast, they allow users to tailor the methodology to different periodicities or velocities. How the different parameters impact the methodology results on the 4-year database is presented in the Appendix sections C.2, C.3, C.4 and C.5.

4.5 Validation on case events

To validate the reliability of the detection technique, the results for two different case events are presented. The first example of LSTID ROI identification corresponds to the extreme geomagnetic storm that occurred on 11–12 August 2024. Figure 4.5 highlights that the generation of numerous LSTIDs characterized this storm.




Figure 4.5: latitude-time evolution of dTEC disturbances for the 11-12 August 2024 storm, with lines representing the RANSAC-fitted line. White dashed lines indicate detected ROIs that are not considered LSTID ROIs, while white solid lines are ROIs that satisfy the LSTID conditions and are plotted along with their corresponding speed in m/s, shown in the white boxes.

The first notable LSTID activity is a low-amplitude wave train propagating at 500–600 m/s, observed around 21:00 on the 11th of August. This train consists of two negative and one positive deviation from the background. Following this, after 22:00, many continuous fluctuations of varying amplitudes and speeds propagate consistently until 03:00 on the 12th of August, and the technique correctly detects these. However, at approximately 02:00, the method misidentifies a positive ROI due to high noise levels between latitudes 45° and 50°. This results in the ROI being incorrectly split into two parts: the lower segment is too short in latitude, and the upper segment fails to overlap with adjacent positive ROIs. Consequently, both parts are discarded. A similar issue arises with the broad negative ROI around 04:00, which is discarded due to a failure in the BGMM step of the segmentation process. Here, the model fails to separate two consecutive overlapping negative disturbances. Around 05:00, an ambiguous detection occurs with a negative dTEC feature propagating at a speed of 568 m/s. Additionally, before the correctly detected LSTID ROIs between 12:00 and 17:00 on the 12th of August, several northward-propagating features are identified between 11:00 and 12:00. These are discarded due to their negative speed values. Although these features are likely LSTIDs crossing the equator, they were excluded based on the velocity criteria. Adjusting this criterion would allow the detection of poleward-propagating LSTIDs, which could be useful for studying such phenomena. Later, around 19:00, two consecutive features (a negative deviation followed by a positive one) appear to form an LSTID but are not fully detected. This is because only the negative feature is identified as a connected region, likely due to lower noise levels (as indicated by the white dashed line around 19:00), especially at latitudes above 45°. The identified negative feature is discarded because no corresponding positive feature is detected within the preceding or following 45 minutes. Lastly, the white dashed line visible around 00:00 highlights a limitation of the technique. Because the method processes daily keogram images, ROIs that extend across two keograms are split and identified in both images with incorrect velocities. This limitation explains the high-speed white solid line just before 00:00 and the subsequent line, which is discarded for unrealistically high velocity. Here, one

must specify that Figure 4.5 was obtained by the concatenation of two daily keograms available in the database.




Figure 4.6: keogram belonging to the weak geomagnetic storm of 1-2 April 2023 with lines representing the RANSAC fitted line. White solid lines are ROI that satisfy the LSTID conditions and are plotted along with their corresponding speed in m/s, which is shown in the white boxes.

Figure 6 presents the LSTID ROI identification process for the moderately disturbed period on 1–2 April 2024. As anticipated, in contrast to the strong storm depicted in Figure 4.5, the weaker geomagnetic activity and reduced energy injection at high latitudes resulted in fewer visible LSTIDs. This observation is particularly significant, as previous statistical studies on LSTID for the European region (e.g., Borries et al. 2009) have primarily focused on major storms, whereas moderately disturbed periods occur more frequently and are thus critical for this study. In Figure 4.6, no white dashed lines are visible, indicating that all identified ROIs satisfied the LSTID conditions. A solitary LSTID propagating at approximately 650 m/s appears around 15:00 and is followed by intermittent LSTID activity between 18:00 and 02:00 on the morning of the 2nd of April. Notably, no major misdetections are evident in this figure. However, two negative features with speeds of 903 and 365 m/s represent false detections of LSTIDs. These false detections highlight the trade-off involved in setting the filtering parameters for the ROI identification technique. Increasing the strictness of the criteria reduces the number of false detections at the cost of more missed detections. The rationale behind the choice of the technique parameters is shown in Appendix C.2, C.3, C.4, and C.5. Since no ground truth for LSTIDs is available, it was not feasible to conduct a more rigorous quantitative evaluation of the impact of filtering parameters on ROI identification performance in terms of false positives and negatives.

After evaluating the detection algorithm's performance in two case studies, the results for the same geomagnetic storm of Figure 4.5 but for different geographical regions are presented. Figure 4.7 presents the detection output for the GNSS stations shown in the left panel of Figure 4.3. The keogram spans a 20° latitudinal range, from 28° to 48° North, maintaining the same extent as EUREF but with different upper and lower limits. Additionally, the longitudinal slice used to generate the keograms is 5° wider. A key challenge in these regions is the limited availability of satellite observations. The CORS and GEONET networks do not provide

BeiDou data, and only a small number of stations collect Galileo signals. As a result, the total number of available line-of-sight measurements is significantly reduced, even though the number of stations remains the same. To mitigate this issue, the longitudinal extent was expanded for both North America and Japan, including more data in the keogram generation process. This broader longitudinal coverage helps lower noise levels and ensures sufficient data coverage across all latitudes and periods. At the same time, enlarging the keogram longitudinal sector also shows that the technique is flexible to changes in latitudinal and longitudinal ranges.




Figure 4.7: dTEC keogram for the 11-12 August 2024 storm over North America. The latitudinal extent goes from 28° to 48° North. The longitudinal range considered to generate the keogram is from 85° to 105° West. White lines represent the RANSAC fitted line. White dashed lines indicate detected ROIs that are not considered LSTID ROIs, while white solid lines satisfy the LSTID conditions and are plotted along with their corresponding speed in m/s, shown in the white boxes.

In the case of North America, the keogram reveals a high noise level, making LSTID detection more irregular and challenging. The first wave trains appear between 22:00 and 04:30, but increased ionospheric activity introduces overlapping disturbances beyond this period, complicating the detection. The lower latitudinal boundary also increases the likelihood of interhemispheric LSTIDs, which can interfere with image segmentation. Around 11:30, two regions of interest (ROI) are filtered out (marked with dashed white lines) because the detected red feature exceeds the maximum velocity threshold (1550 m/s). A similar issue arises between 19:00 and 20:00, where two more ROIs are discarded due to their speeds exceeding 2 km/s. Lastly, at 15:00, an ROI is missed by the identification process, causing the adjacent positive and negative ROIs to be discarded as well, since they lack a properly paired opposite-amplitude feature with matching physical characteristics. Despite these challenges stemming from high noise levels and overlapping ionospheric disturbances, the algorithm maintains an acceptable level of performance.

Figure 4.8 presents the results for the Japanese region, where the keogram spans a 15° latitudinal range, narrower than the European and North American cases. Overall, when comparing the detection results shown in Figure 4.7 with those of Figure 4.8, it is possible to note how Figure 4.8 shows more coherent parallel structures and thus are more easily identified LSTID ROIs over Japan. However, two notable misdetections occur around 15:00 and 17:00.

The first misdetection arises from underestimating the negative band's inclination due to noise at low latitudes, which prevents the required speed to match with the preceding positive ROI (both fitted lines appear dashed). The second misdetection occurs due to an incorrect segmentation of the irregular region during the BGMM processing step. This results in an underestimated speed for the positive ROI, preventing it from merging with the following correctly classified LSTID ROIs. Despite these minor misclassifications, the algorithm demonstrates strong detection capabilities, particularly in the Japanese sector, where noise levels are comparatively lower than in the North America case, even if the latitudinal extension is 5 degrees smaller.




Figure 4.8: dTEC keogram for the 11-12 August 2024 storm over Japan. The latitudinal extent goes from 30° to 45° North. The longitudinal range considered to generate the keogram is from 125° to 145° East. White lines represent the RANSAC fitted line. White dashed lines indicate detected ROIs that are not considered LSTID ROIs, while white solid lines satisfy the LSTID conditions and are plotted along with their corresponding speed in m/s, shown in the white boxes.

4.6 Statistical analysis

Now that the detection results have been presented for two case studies, a statistical analysis of LSTID occurrences is performed to investigate, on a broader level, the reliability of the proposed algorithm. LSTID ROI occurrences are investigated in relation to the hour of the day and the month to highlight potential influences of the background ionization and relative solar wind/geomagnetic field geometry. Moreover, the relationship between geomagnetic indices and the LSTID detection occurrences and rates is investigated.

The magnetic indices considered are Hp30, SymH and IMAGE auroral electrojet indices (IE). The geomagnetic Hp30 index, described by Yamazaki et al. (2022), is similar to the widely adopted Kp index, but offers higher time resolutions, specifically 30 minutes. Moreover, unlike the Kp index, which is capped at a maximum value of 9, the Hp30 index is open-ended, allowing it to better represent the intensity of severe geomagnetic storms in greater detail than the three-hourly Kp index. The two indices, Hp30 and Kp, are strongly correlated. For example, for the 11-12 August 2024 storm, the maximum Kp reached 8 between 12 and 15 UT, while Hp30 had an average value for the same period of 7.88.

The Symmetric H-component (SYM_H) index is a key indicator of ring current disturbances during geomagnetic storms, utilizing data from a longitudinal array of low-latitude ground-based magnetometers (Iyemori, 1990). SYM_H is similar to the Disturbance Storm Time (Dst) index as it characterizes the ring current disturbances; however, it offers a refined 1-minute temporal resolution, making it particularly valuable for investigating short-term fluctuations in geomagnetic disturbances.

Lastly, IE is an alternative to the classic Kyoto AE index (Tanskanen, 2009) provided by the Finnish Meteorological Institute. The main difference is that the IE encompasses a larger network of magnetometers distributed across the northern European area, thus allowing a more significant representation of the area where European LSTIDs are normally generated.

In addition to the statistical study and geomagnetic activity influences, thanks to the RANSAC fitting procedure, it is also possible to estimate the velocity and period distribution along with their dependence on geomagnetic conditions. Before moving to the statistical results, the definition of the LSTID ROI speed and period must be clarified. Any given LSTID ROI has a related period and speed. While the speed is derived by converting the slope to a meridional velocity, the LSTID period is twice the distance between the given ROI and the adjacent and opposite amplitude ROIs. For instance, ROI 10 and 11 of Figure 4.4 will have different periods since ROI 10 will be characterized by a period twice the time distance between ROI 10 and 11. In contrast, the ROI 11 period will be twice the average time distance between ROI 11 and ROIs 10 and 12.




Figure 4.9: LSTID ROIs parameters statistics. Panels (a) and (b) show the monthly and hourly distribution of LSTID ROI occurrences. Moreover, (a) also shows the monthly mean of the Hp30 index in orange. Similarly, (b) shows the hourly mean of the 1-hour shifted IE index in orange. Panel (c) shows the period distribution, with the mean and standard deviation on top of the panel. Panel (d) shows the mean and standard deviation for the LSTID ROIs speed.

Having demonstrated the technique's capability to reconstruct LSTIDs in two distinct scenarios, the focus moves on its statistical characterization when applied to the entire database of keogram images.

As LSTIDs need time to propagate from the auroral zone to the area where they are detected, a time shift was introduced to the indices to understand better the relationship between geomagnetic parameters and LSTID ROI occurrences: 2 hours for Hp30 and SymH and 1 hour for IE. The time shift values were selected to maximize the detection rates.

Additionally, the hourly and monthly distributions of LSTID ROI occurrences are shown in Figure 4.9. Figure 4.9a shows that LSTIDs are notably more frequent around the equinoxes, with a preference for the autumn equinox. The distribution exhibits a clear bimodal pattern. This behavior can be attributed to the increased strength and frequency of geomagnetic storms during equinoxes (Gonzalez et al., 1994; Tsugawa et al., 2004), driven by the geomagnetic field's orientation relative to the solar wind. Moreover, the monthly distribution shows that LSTIDs are more likely to happen in autumn than in spring in Europe; this difference is partially explained by the higher value of average Hp30 during autumn, as shown by the orange curve. However, some other phenomena must play a crucial role, as the difference in the number of occurrences for autumn and spring is not completely justified by the difference in Hp30.

Furthermore, as LSTIDs represent displaced plasma, their higher amplitude during elevated background plasma density periods facilitates the detection. A similar explanation applies to the hourly distribution in Figure 4.9b, which shows that LSTIDs are primarily detected between 12:00 UT and 01:00 UT. The orange line in Figure 4.9b represents the hourly mean of the one-hour shifted IE index, revealing similarities between the two distributions. LSTIDs are thought to be generated by plasma displacement caused by neutral waves. However, post-midnight discrepancies arise due to the reduced background plasma density, which limits the detectability of LSTIDs, even if electrojet current enhancements are still generating neutral waves. Turning to LSTID ROI periods, Figure 4.9c shows a mean period of 62 minutes with a standard deviation of 14 minutes, closely matching the results of Borries et al. (2009), who reported an average period of 59 minutes with a similar standard deviation. The speed distribution, presented in Figure 4.9d, has a mean of 629 m/s and a standard deviation of 215 m/s, aligning well with Borries et al. (2009), who reported values of 684 m/s and 195 m/s, respectively. It is worth noting, however, that the distributions for period and speed differ in one key aspect. While the mean period corresponds closely to the mode, the mode of the speed distribution is lower than its mean (525 m/s compared to 629 m/s). This discrepancy arises from the velocity estimation process, which relies on the slope of ROIs in the keogram. The technique is susceptible to overestimating the average velocity because small errors of different signs in the slope of quasi-vertical lines result in different errors in the calculated velocities. In particular, positive deviations produce larger errors due to the tangential relationship between slope and speed, which can lead to significant overestimations of the calculated speed.




Figure 4.10: violin plot of the period (a), speed (b) and mean LSTID ROI amplitude (c) with respect to the one-hour shifted IE index. The black diamonds represent the median, the grey horizontal lines are the mean, the darker shaded area represents the interquartile range, and the shape of the half-violin represents the distribution shape. Every violin shows the statistics for the LSTID ROIs related to a 1-hour shifted IE corresponding to the value shown on the x-axis.

As LSTID generation is closely linked to Joule heating, a relationship between the LSTID wave parameters and the IMAGE IE index is expected, as it measures the intensification of ionospheric currents responsible for Joule heating. Figure 4.10 illustrates how the LSTID ROI period, speed, and mean LSTID ROI amplitude distribution vary with the shifted IE index. Here, the mean LSTID ROI amplitude refers to the average value of dTEC for the pixels belonging to the considered LSTID ROI. Figure 4.10a reveals an almost linear decrease in the LSTID period as IE increases. For an IE of 0 to 100 nT, the median period is approximately 67 minutes, but it decreases to around 54 minutes for IE values exceeding 1000 nT. While the period exhibits a clear inverse relationship with IE, the relation to speed, shown in Figure 4.10b, is more complex. Speed retrieval, being the most error-prone aspect of the technique, has a higher variability, particularly for low IE values. Nevertheless, a slight upward trend in speed is visible, particularly for IE values above 200 nT. The median speed rises from approximately 566 m/s at an IE of 300 nT to around 628 m/s for IE values greater than 1000 nT. Finally, the median of the mean LSTID ROI amplitude, shown in Figure 4.10c, demonstrates a direct relationship with the IE index. The amplitude represents the average absolute dTEC value for pixels within each ROI. The amplitude increases from about 0.05 TECu for an IE of 0 to 100 nT to approximately 0.12 TECu for IE values exceeding 1000 nT. This relationship is consistent with findings by Borries et al. (2009). In summary, a stronger energy injection into the auroral oval leads to LSTIDs with shorter wavelengths and higher amplitudes. These results suggest that higher IE values generate stronger atmospheric disturbances mapped to the ionosphere as faster, shorter-period, and higher-amplitude LSTIDs. This relation highlights the link between increased IE and intensified wind, temperature, and pressure disturbances in the thermosphere.




Figure 4.11: LSTID ROIs detection occurrences and rates with respect to shifted geomagnetic parameters along with their probability distribution. Panels (a), (b) and (c) show in blue the number of LSTID ROIs detected when the corresponding shifted Hp30, SymH and IE index corresponded to the values shown on the x-axis. The orange dots and line represent the LSTID ROI detection rate in percentage, normalized with respect to the distribution of occurrences of the corresponding shifted geomagnetic indices shown on the x-axis. The green squares and lines instead represent the percentage probability distribution of the geomagnetic indices shown on the x-axis.

After examining the parameters of LSTID ROIs and their hourly and monthly dependencies, the focus is moved to their occurrence and detection rate relative to shifted geomagnetic indices. Figure 4.11 presents the results, where the blue bars show the distribution of events with respect to Hp30 (Figure 4.11a), SymH (Figure 4.11b), and IE (Figure 4.11c). Since geomagnetic disturbances are less frequent with higher magnitudes, the raw occurrence distribution of LSTID ROIs alone does not provide a complete picture of their relationship to geomagnetic activity. To address this, the green squares in Figure 4.11 depict the probability distribution of geomagnetic indices, while the orange dots show the detection rate expressed as a percentage. The detection rate is the number of LSTID ROIs normalized by the corresponding probability distribution of the geomagnetic index values along the x-axis. Figure 4.11a indicates that the highest occurrence of LSTID ROIs (blue histogram) corresponds to Hp30 values of 3. However, the probability distribution of Hp30 (green line) shows that the most frequent values are 1 and 2, with Hp30 values above five occurring each in less than 1% of cases. This implies that most LSTIDs are generated during mildly disturbed conditions (Hp30 between 3 and 5). Notably, the detection rate increases almost linearly with Hp30, reaching a maximum of 46% at Hp30 = 7 before decreasing. This behavior aligns with findings by Tsugawa et al. (2004), who observed a maximum occurrence rate of 75% for Kp = 9. The lower detection rate in this study is due to the finer temporal resolution of Hp30 (30 minutes versus 3 hours for Kp), which increases the likelihood of time bins without detectable LSTIDs. Additionally, the geographic coverage of this study introduces higher noise levels compared to studies where denser receiver networks were used, such as in Japan (Tsugawa et al., 2004).

For extreme storms, interference between equatorward and poleward LSTIDs at midlatitudes and noise from auroral oval displacement and other phenomena likely reduce detection rates. Another possible explanation is that during extreme storms, the assumption of a thin-layer ionosphere might fail due to the increase in slab thickness (Pignalberi et al., 2022), causing an increased error in the projection of IPP and in the integration of disturbances along the line of sight. Figure 4.11b shows the distribution of LSTID ROIs with respect to the SymH index. Similar to Hp30, most LSTIDs are associated with quiet to mildly disturbed periods, with the highest occurrences for SymH values between -25 and -15 nT. The detection rate increases steadily as SymH decreases, peaking at 58% for SymH $\approx$ -120 nT. However, a saturation effect is observed as detection rates decline for extreme disturbances (SymH < -120

nT). An intriguing feature in Figure 4.11b is the spike in detection rate to 87.5% for $\text{SymH} \approx 60$ nT. Despite the low number of occurrences, this anomaly suggests a unique generation mechanism for LSTIDs under such conditions, as enhanced auroral currents are unlikely for highly positive SymH values. Investigating the specific events responsible for this anomaly and their physical causes is beyond the scope of this study and will be addressed in future work. Figure 4.11c examines the relationship between LSTID ROIs and the IMAGE IE index, which directly measures the phenomena responsible for gravity wave excitation. Most LSTID occurrences correspond to IE values between 175 and 225 nT, representing mildly disturbed conditions (7% probability distribution of IE). As with Hp30 and SymH , the detection rate increases with IE, reaching 40% for values between 875 and 925 nT before saturating or declining to around 35% at higher IE values. In summary, the highest number of LSTID ROIs are associated with medium-low geomagnetic activity, as extreme storms are relatively rare. However, the detection probability increases dramatically during highly disturbed periods, reaching 46%. It is important to note that the detection rate in this study is based on 30-minute time bins. Larger time bins, such as the 3-hour bins characteristic of the K_p index, would yield higher detection rates. For example, under 3-hour binning, the detection rate for a K_p of 7 would increase to approximately 75%.

Chapter 5

LSTID Forecasting

As discussed in previous chapters, LSTIDs are one of the space weather phenomena affecting the ionosphere on continental scales, typically generated by enhanced energy input at high latitudes during geomagnetic storms. Their equatorward propagation disturbs HF communication, GNSS positioning, and radar operations, which makes reliable forecasting of their occurrence a priority for ionospheric monitoring systems. Unlike detection, which has been addressed through various instruments and methodologies (Belehaki et al., 2020), forecasting remains challenging due to the nonlinear coupling between solar wind forcing, magnetospheric response, and ionosphere–thermosphere dynamics.

In this thesis, LSTID forecasting is addressed through two complementary approaches, each developed in the framework of a different project in which I was involved and supported by independent event catalogues. The first model, developed within ARISTIDE, builds on the autonomous GNSS-based catalogue of LSTIDs generated from the keogram-based image processing presented in Chapter 4. Instead, the second prototype, developed in the T-FORS project, relies on an ionosonde-based catalogue validated by experts (Segarra et al., 2025).

The following sections present these two forecasting strategies in parallel: first outlining their data sources and event catalogues, then describing the construction of forecasting datasets, the modelling approaches, and the role of explainable AI in interpreting results. Finally, validation and operational perspectives are discussed, highlighting complementarities and differences between the GNSS-based ARISTIDE and the ionosonde-driven T-FORS models.

5.1 From LSTID catalogues to target variables

Any forecasting framework requires a well-defined target variable against which predictive models can be trained and evaluated. In the case of LSTIDs, the target must be derived from detection catalogues composed of distinct events characterised by onset time and duration. These catalogues serve a dual role: they provide the ground truth for supervised learning and impose the temporal and spatial definitions of an LSTID event.

For machine learning, these catalogues are typically converted into binary time series: each time bin is labelled as containing an LSTID. This process ensures alignment between the input drivers and the event occurrence, and enables models to learn statistical associations across large datasets. The following subsections present the two event catalogues used in this thesis as target variables for forecasting: the HF-Interferometry (HF-INT) based catalogue from the T-FORS project, and the GNSS keogram-based catalogue developed within ARISTIDE.

5.1.1 T-FORS Catalogue

For the T-FORS model, a refined version of the TechTIDE catalogue generated by the HF-INT method was exploited (Altadill et al., 2020; Segarra et al., 2025). This catalogue covering the years from 2014 to 2022 is based on the analysis of MUF(3000)F2 data from ten European ionosondes, where the HF-INT method searches for coherent periodicities across multiple stations. An event is declared if at least three ionosondes show consistent periods,

spectral features, and cross-correlations. The HF-INT method is explained in detail in Altadill et al. (2020).

To increase reliability, the raw HF-INT outputs were refined through expert validation (Segarra et al., 2025). During this process, specialists checked for consistency in azimuth propagation across the ionosonde network and filtered out spurious detections. The resulting HF-INT LSTID catalogue provides averaged characteristics of each event, including period, amplitude, spectral energy contribution (SEC), propagation velocity, onset time, duration, and a quality indicator linked to data availability (Segarra et al., 2025).

This event catalogue was converted into a binary time series at a 30-minute cadence for forecasting. Each interval overlapping with an LSTID event was labelled positive, and all others negative. To align with the forecasting task, the labelling was shifted such that a bin was considered positive if an event was expected to occur within the following three hours. In other words, the model is trained to forecast if there will be an LSTID in the following three hours. This explicitly defines the forecasting horizon and ensures that the classifier learns precursory signatures in the external drivers that precede event onset.

Despite this refinement, limitations remain. The catalogue is geographically coarse-grained: events are treated as if they affect the entire European sector, whereas in reality, LSTIDs are more localised. This approximation reduces spatial specificity and introduces uncertainty, particularly during weaker disturbances. The catalogue could be affected by missing events: the presence of concurrent ionospheric disturbances (e.g., sporadic E layer, spread F) limits the capability of the ionosonde-based measurements in highlighting the occurrence of LSTID. Moreover, the reliance on expert validation adds an element of subjectivity, making the catalogue robust but not fully free from bias.

Finally, it should be noted that the typical propagation time from auroral latitudes to the European mid-latitudes is between ~ 3.5 h and ~ 1.3 h for phase speeds of 300–800 m/s. This physical timescale supports the choice of a three-hour forecasting horizon, as it corresponds to the time needed for disturbances generated at high latitudes to reach the mid-latitude ionospheric sector covered by the T-FORS model.

5.1.2 ARISTIDE Catalogue

As anticipated, the target variable for the ARISTIDE forecasting model is derived from the GNSS-based catalogue of LSTID ROIs obtained through the keogram analysis pipeline described in Chapter 4. In this catalogue, each LSTID ROI is defined by a single line, corresponding to either a positive or a negative amplitude in the keogram. For forecasting purposes, the full catalogue spanning 2021–2024 was transformed into a binary time series and aligned with the resolution of the external drivers.

Constructing the binary target series involves several preprocessing steps to guarantee robustness and physical interpretability. First, a baseline series was generated at 30-second intervals, covering the full study period. For each LSTID line in the catalogue, the time interval between its starting and ending point was expanded by a tolerance of ± 600 s, and the corresponding bins were set to one. This temporal padding accounts for uncertainties in the automatic line detection and ensures that closely spaced tracks are merged into a single coherent interval.

Next, the binary sequence was cleaned of short gaps. Consecutive windows of zeros shorter than 7200 s (two hours) were filled with ones, effectively preventing the fragmentation of disturbances that should be considered part of a single event. This step reflects the physical expectation that large-scale disturbances occur as sustained activity rather than isolated lines.

Finally, quality control was applied to discard spurious detections. Since each ROI corresponds to only one line, positive and negative ROIs were considered together within a

candidate event interval. A minimum of four LSTID ROI was required to validate an event; intervals failing this criterion were reset to zero. The resulting sequence was then resampled to match the 30-minute cadence of the external drivers.

The forecasting task, however, does not operate directly on this binary series. Instead, model inputs are organised into fixed-length training windows of driver data, and each window is paired with an output label defined over a subsequent control window. The label is positive if at least one positive value is present in the control window, and negative otherwise. In this way, the problem is explicitly formulated as an early-warning task: the model must learn to associate patterns in the drivers with the probability of an LSTID appearing in the near future. Importantly, the training and control window length are not fixed a priori but treated as a hyperparameter of the optimisation framework, allowing the forecasting horizon to adapt to the temporal scale most informative for the system.

5.2 From physical drivers to input features

The occurrence of LSTIDs at midlatitudes is mostly caused by energy input from the solar wind into the magnetosphere–ionosphere system (Ferreira et al., 2020). This input modifies currents and conductivities in the auroral zone, which in turn excite thermospheric gravity waves that propagate equatorward as LSTIDs. To capture these processes in a forecasting framework, external drivers are leveraged to quantify the strength, timing, and variability of this energy deposition in the auroral oval.

From a physical perspective, different classes of drivers represent distinct stages in the solar-terrestrial chain. First, solar-wind parameters (velocity, density, and IMF) provide information on the upstream conditions at the Lagrangian point L1. Their fluctuations determine whether energy will be injected into the magnetosphere, with coupling functions such as the Newell index offering a measure of this energy transfer efficiency (Newell et al., 2007). The following step in the solar-terrestrial chain is the geomagnetic field response to the solar wind-magnetosphere-ionosphere coupling, which is directly measured by geomagnetic indices. In this work, regional indices representing auroral, sub-auroral, and equatorial regions were considered. Hp30 and ap30 describe subauroral activity, while Dst and SymH measure the strength of the storm-time ring current, characteristic of equatorial regions. The IMAGE auroral electrojet indices (IU, IL, IE) quantify the strength of the high-latitude electrojets over the Scandinavian region, which are directly linked to Joule heating and thus to the excitation of atmospheric gravity waves over the European sector.

Long-term variability of the ionosphere is set by solar radiation, for which the F10.7 index is used as a proxy for EUV flux. In addition, the solar zenith angle (SZA) is included as a driver that controls the daily baseline ionisation level and conductivity. Daytime low SZA values favour stronger ionospheric responses and clearer signatures, while high SZA values during post-midnight conditions reduce background electron density, often limiting detectability.

5.2.1 T-FORS input features

In addition to the external drivers described above, the T-FORS model integrates HF-INT outputs obtained from the TechTIDE Application Interface Program (API). This NRT outputs include a regional activity index (HF-INT_{EU}), which summarises coherent LSTID activity across the European ionosonde network, and perturbation parameters derived from individual stations belonging to the TechTIDE network. Among these are the SEC, the azimuth of propagation, and the velocity of the disturbance, all of which provide physically meaningful parameters characterising ongoing ionospheric wave-like activity. The full set of input features used in T-FORS is listed, along with a brief description, in Table 5.1.

Several pre-processing steps were applied to engineer the full set of input features. Time series data often contain short-term fluctuations that obscure the underlying behaviour of the system. To address this, moving averages were employed as smoothing operators, enabling the model to focus on persistent patterns rather than noise. Different averaging windows (2, 3, and 12 hours) were used in parallel to capture short-term variability and longer-term persistence. Specifically, moving averages were applied to the auroral electrojet indices (IE, IU) and the HF-INT_{EU} index, where averaging enhanced the correlation with the forecasting target.

Beyond smoothing, trend information was extracted by decomposing the auroral indices into ascending, stationary, and descending phases. This was achieved by applying an exponential moving average (EMA) to emphasise longer-term variations, and then clustering the results using k-means (Ikotun et al., 2023). This transformation allows the model to encode the magnitude of a driver and its current trajectory, which is often more informative for anticipating disturbances.

Finally, all variables were resampled to a uniform 30-minute cadence, matching the target labels. Where outliers could bias the average, median aggregation was preferred over mean aggregation. This standardisation ensured consistency across the feature space while reducing computational complexity and minimising the influence of spurious spikes.

The resulting input set thus combines upstream solar-wind and geomagnetic indices with ionosonde-based activity descriptors, enriched by smoothed and trend-oriented transformations. This comprehensive feature space provides the T-FORS forecasting model with the physical context of energy input and direct ionospheric signatures, increasing its ability to detect precursor conditions leading to LSTID onset.

Table 5.1: Physical drivers of the T-FORS model, along with a brief description.

Feature Name	Description
IE, IU	Auroral electrojet indices
HF-INT _{EU}	High-frequency interferometry index, identifying activity over the European network of ionosondes
SZA	Solar zenith angle, the angle between the local zenith and the line of sight from that point to the Sun
F10.7	Solar radio flux at 10.7 cm (2800 MHz)
ρ , vx, Bz	Solar wind and interplanetary magnetic field parameters, measured by ACE and/or DSCOVR spacecraft at the Lagrange point L1 and projected to the bow shock nose.
Dst	Disturbance storm time index
Newell	Coupling function that quantifies the energy transfer from the magnetosphere to the ionosphere
Hp30	Half-hourly sub-auroral geomagnetic index in logarithmic form
SEC (Athens, Juliusruh, Pruhonice, Rome, San Vito)	Spectral energy contribution (%) of the perturbation (one feature per ionosonde)
Azimuth (Athens, Juliusruh, Pruhonice, Rome, San Vito)	Azimuth of the perturbation (one feature per ionosonde)
Velocity (Athens, Juliusruh, Pruhonice, Rome, San Vito)	Velocity of the perturbation (one feature per ionosonde)

5.2.2 ARISTIDE input features

As for the T-FORS case, the ARISTIDE model leverages a set of external drivers that describe both upstream solar wind conditions and the magnetospheric-ionospheric response, complemented by the solar zenith angle to account for diurnal modulation. The full list of parameters used for the ARISTIDE model is shown in Table 5.2. All the input series were resampled to a 30-minute cadence, smoothed with an exponential moving average, and normalised to [0, 1] with min-maxing.

Table 5.2: Physical drivers of the ARISTIDE model, along with a brief description.

Feature Name	Description
IE, IU, IL	Auroral electrojet indices
SymH	Disturbance storm time index
Hp30, ap30	Half-hourly sub-auroral geomagnetic index in logarithmic form
SZA	Solar zenith angle, the angle between the local zenith and the line of sight from that point to the Sun
v , ρ	Solar wind speed magnitude and particle density at L1.
B_{tot} , B_x , B_y , B_z	Interplanetary magnetic field at L1. Magnitude and GSM components.
F10.7	Solar radio flux at 10.7 cm (2800 MHz)

A preliminary study on the predictive potential of the input features was performed in the form of a cross-correlation analysis. The cross-correlation coefficient R_{xy} at time lag τ is calculated as follows:

$$R_{xy}(\tau) = \frac{1}{N} \sum_{t=0}^{N-1} x(t) \cdot y(t + \tau) \quad (5.1)$$

where N is the time series length and $x(t)$ and $y(t + \tau)$ are the values of the two time series at times t and t shifted by τ . The cross-correlation is then normalized with respect to the standard deviation of the two series to ensure that its value is in the $[-1,1]$ range.

$$\rho_{xy}(\tau) = \frac{R_{xy}(\tau)}{\sigma_x \sigma_y} \quad (5.2)$$

where $\rho_{xy}(\tau)$ is the normalized cross-correlation at time shift τ and σ_x , σ_y are the standard deviations of the time series x and y .

Specifically, the cross-correlation analysis highlights the different levels of coupling between the input drivers and the binary target series used in ARISTIDE. Figures 5.1, 5.2, and 5.3 display the correlation coefficients as a function of progressive time shifts, corresponding to 30 minutes for each step. The time-lag corresponding to the highest correlation highlights the typical delay between the input and output features, while the decay of correlation with increasing shift indicates the temporal persistence of the predictive power of each parameter.




Figure 5.1: Correlation coefficient analysis for the ap30 and Hp30 indices. The y-axis shows the correlation coefficient relative to the time shift of the x-axis. The x-axis shows the number of time steps in units of 30 min.

Geomagnetic indices (ap30, Hp30) exhibit a strong correlation already at zero steps, peaking at 0.34 and 0.42 at ~ 3 time steps (1.5 hours). This behaviour shows their high sensitivity to short-term disturbances, as the correlation rapidly decreases beyond ~ 6 shifts (3 hours). Their usefulness lies primarily in the immediate response to geomagnetic activity, making them effective for real-time forecasting.

Solar wind parameters (Figure 5.2) present a more diverse picture. The total interplanetary magnetic field strength (B_{tot}) shows the highest initial correlation (~ 0.27) with a decay similar to ap30 and Hp30, suggesting that coupling processes with the magnetosphere-ionosphere system occur with time scales of a few hours. Density and velocity show weaker correlations, peaking below 0.2 and diminishing gradually, which implies a weaker role in driving ionospheric disturbances. The IMF components behave differently: B_z shows a clear negative correlation at short lags, consistent with its geoeffective southward (and negative) orientation. At the same time, B_y and B_x contribute less significantly due to their lower impact on the efficiency of the solar wind-magnetospheric coupling. As expected, the time lags for solar wind parameters are slightly higher (~ 6 , 3 hours) than those related to sub-auroral indices (ap30, Hp30). This discrepancy is explained by the time the solar wind needs (measured at L1) to reach the Earth's magnetosphere (~ 1 hour, depending on solar wind speed).



Figure 5.2: Correlation coefficient analysis for the solar wind parameters B_{tot} , B_x , B_y , B_z , density, and speed. The y-axis shows the correlation coefficient relative to the time shift of the x-axis. The x-axis shows the number of time steps in units of 30 min.

Auroral electrojet indices (IE, IU, IL) demonstrate correlations of similar order of magnitude to ap30 and Hp30, but with oscillatory behaviour at longer lags (Figure 5.3). This behaviour reflects a degree of periodicity in the complex magnetospheric-ionospheric coupling processes at high latitudes, where substorm dynamics and current systems can lead to non-monotonic correlation patterns. The auroral electrojet is known to be stronger during nighttime due to reconnection in the geomagnetic tail, leading to a recurrent intensification. This behavior is visible in Figure 4.9b, where the IE distribution shows a higher average value for nighttime hours.

Overall, the results confirm that geomagnetic indices and the orientation and strength of the solar wind magnetic field are the most informative predictors at short lead times. Solar wind density and speed provide additional context but have a lower direct correlation. Moreover, the time lags of maximum correlation are in accordance with the delays expected for LSTID generation and propagation towards middle latitudes. These findings justify the choice of including the broad set of drivers listed in Table 5.2, since each parameter contributes to capturing different aspects of the solar wind–magnetosphere–ionosphere system.

Following these results, lagged versions of selected drivers were added to the feature space whenever the correlation exceeded 0.3 at non-zero lags. This enrichment step captures

precursory information more effectively and reflects the underlying physical delay in energy transfer through the coupled system.




Figure 5.3: Correlation coefficient analysis for the IMAGE auroral electrojet indices IE, IU, and IL. The y-axis shows the correlation coefficient relative to the time shift of the x-axis. The x-axis shows the number of time steps in units of 30 min.

5.3 The Machine Learning Models

The forecasting problem is addressed through two distinct machine learning approaches, each reflecting the project context in which it was developed. Both aim to predict the occurrence of LSTIDs in the near future, but differ in how temporal information is represented and exploited.

In the T-FORS project, the task is framed as a classification problem, where each 30-minute step is treated as an independent instance labelled according to the expected occurrence of an event. CatBoost (Prokhorenkova et al., 2019), a gradient boosting on decision trees model, was chosen for its ability to handle heterogeneous predictors and imbalanced data, while offering built-in interpretability through SHAP analysis (Lundberg & Lee, 2017). Validation relied on time-aware cross-validation to respect chronological order between training and testing segments.

In ARISTIDE, the task is formulated as time-series classification. Here, the QUANT algorithm (Dempster et al., 2023) processes sliding windows of input drivers, representing them through quantile features that capture temporal structure efficiently. This approach is computationally lightweight yet competitive with more complex models. Validation was based on sliding-window partitions, with performance assessed in terms of both classification skill and predictive lead time.

While the modelling paradigms differ, boosted decision trees in T-FORS and interval-based quantile classification in ARISTIDE, both systems incorporate explainable AI to ensure that model behaviour can be related to physical processes.

An additional aspect shared by both frameworks is their probabilistic formulation. Instead of only producing binary yes/no outputs, the models produce a calibrated probability that an LSTID will begin within the next three hours. Probabilistic forecasting is crucial for early-warning systems because it quantifies the confidence of each prediction, enabling adaptive decision-making under uncertainty. Users can adjust operational responses depending on the probability level rather than a fixed threshold, balancing the trade-off between missed detections and false alarms according to context.

A model calibration step is performed to ensure that these probabilities are meaningful. Calibration aligns the predicted probabilities with the actual observed frequencies of events; e.g., among all predictions with 0.7 probability, approximately 70 % should correspond to real LSTIDs. Techniques such as Venn–Abers predictors or isotonic regression are employed to correct systematic biases in the raw model outputs. Properly calibrated probabilities provide reliable risk information and make the forecasting system suitable for operational early-warning use, where quantified uncertainty is as important as detection accuracy.

5.3.1 T-FORS – CatBoost

Within the T-FORS project, the LSTID forecasting task was cast as a classification problem, where each 30-minute interval is labelled according to whether an event is expected within the following three hours. This formulation transforms the forecasting challenge into a supervised learning setting in which the model must map multivariate inputs into a binary output.




Figure 5.4: CatBoost working principle, which learns subsequent trees to reduce the error made by the previous ones, thus minimizing the loss function in Equation 5.2.

The proposed model stems from an efficient, fast, and scalable gradient-boosting on decision trees framework, CatBoost, which uses oblivious (or symmetric) trees as the base learner (Dorogush et al., 2018). An oblivious tree is limited to a single symmetrical split, with either side being balanced by the algorithm’s built-in feature importance measurement. This feature, compared to classic trees, leads to an efficient implementation, ensuring efficient training and inherent regularisation. The model works by training a series of learners serially using the boosting method and accumulating the outputs of all learners as a result (Figure 5.4), thereby improving the accuracy and applicability of the algorithm. The loss function used for training is a log-loss function, which for an N-class problem is defined as:

$$L = \frac{\sum_{i=1}^N w_i (c_i \log(p_i) + (1-c_i) \log(1-p_i))}{\sum_{i=1}^N w_i} \quad (5.3)$$

with p_i being the probability estimate relative to the class c_i , and w_i being class weights.

The benefits of using the CatBoost ML framework imply that the required data preprocessing is quite minimal. Indeed, the model natively handles categorical variables, missing values, and automatic feature scaling. To yield sound estimates of the model generalisation error, a cross-validation approach was employed during model training, appropriately tailored for time series whose samples are not independent and identically distributed. This approach provides train/test indices to split data samples into train/test sets; in each split, test indices are higher than the previous ones, and shuffling of the instances is not applied. In the k-th split, one gets the first k folds as a train set, and the (k + 1)-th fold as a test set. Hence, unlike standard cross-validation methods, successive training sets are supersets of those that come before them. Model hyperparameters are fine-tuned within an efficient optimisation framework, Optuna (Akiba et al., 2019). The latter allows for sampling hyperparameters and pruning unpromising trials via a Treestructured Parzen Estimator (TPE) algorithm (Watanabe, 2025). This sampling algorithm pinpoints an optimal configuration using

a Gaussian mixture model to learn model hyperparameters. Precision, recall and F1-score, and the confusion matrix were used to monitor the optimization process and quantify the model's performance. Recall (or sensitivity) defines the fraction of true positive samples being predicted as positive among the original samples, namely:

$$R = \frac{T_P}{T_P + F_N} \quad (5.4)$$

where T_P and F_N denote true positives and false negatives, respectively. Precision is the fraction of true positive samples among the samples predicted as positives, namely:

$$P = \frac{T_P}{T_P + F_P} \quad (5.5)$$

with F_P being the number of false positives. F1-score considers both precision and recall, finding a trade-off between the two in the form of a harmonic mean, or:

$$F_1 = \frac{2}{R^{-1} + P^{-1}} = \frac{2PR}{P + R} \quad (5.6)$$

Out of true/false positive/negative, one can then provide an overview of the model's classifications. The confusion matrix for a binary classifier consists of a 2x2 matrix comparing the actual target values with those predicted by the model. This gives a holistic view of how the classifier is performing and what kind of errors it is making. This trade-off is particularly relevant in an early-warning context: high precision reduces false alarms that may erode user trust, while high recall ensures that critical LSTID events are not missed, even at the cost of more frequent alerts.

An artificial intelligence system that cannot justify how the results have been generated from the data, i.e., a black-box model, is difficult to interpret and trust. Therefore, once the model has been trained, an explanatory framework was employed to assess the influence of the various features on the model output, both from a global and from a local, event-level perspective. A global explanation, i.e., the effect of features on the overall model, can serve as feature importance and help select features that actually impact the model performance, considering not only the effect of individual variables, but also the synergistic interplay between them. On the other hand, a local explanation, i.e., the explanation of prediction for an individual sample, allows one to infer and visualise the primary features influencing the prediction results in a single sample as well as the degree of influence of the corresponding features.

Explainable Artificial Intelligence (XAI) aims to help understand how models make decisions in the prediction process and provide guidance for feature selection and model optimisation by improving model transparency. The SHapley Additive exPlanation (SHAP) method (Lundberg & Lee, 2017), derived from coalitional game theory, was used to interpret the classification model. This explanatory framework is based on Shapley values, which assume that each feature is a player in a game where the prediction is the payout, allowing us to attribute the payout among features fairly. According to SHAP, the predicted values of the model are interpreted as a linear function g of the variables, namely:

$$g(z) = \varphi_0 + \sum_{i=1}^N \varphi_i z_i \quad (5.7)$$

where z is the n -dimensional feature vector, φ_0 denotes the model baseline value and φ_i is the SHAP value of the feature z_i which represents the average marginal contribution to the model

output. The larger the absolute SHAP value, the greater the influence of the feature on the predicted value of the model, and its positivity or negativity represents the direction of influence. From their properties, it follows that the SHAP values of all the input features sum up to the difference between the baseline (expected) model output and the current model output for the prediction being explained.

As anticipated, machine learning models leverage an output threshold that transforms a raw model score into a binary output. That is why probabilistic forecasting methods become relevant whenever a certain level of uncertainty in a system cannot be reduced, which is mostly the case when complex systems are involved. Forecasts based on probabilistic models can ensure robust decisions are made during alert management operations (Krzysztofowicz, 2024).

Once the model is calibrated, the output score can be interpreted as a conditional probability estimate $P(y|X)$ for each class (i.e., the probability of occurrence of an LSTID). It can therefore be provided directly as a model output, allowing the user to choose the action to undertake, based on the level of risk they wish to accept. A classifier is said to be calibrated if a class prediction with a score $p \in [0, 1]$ is correct $(100 \cdot p)\%$ of the time. Formally, the classifier $f: X \rightarrow Y$ is calibrated if:

$$P(\hat{Y} = y | \hat{P} = p) = p, \forall p \in [0, 1] \quad (5.8)$$

where $\hat{Y}$ denotes the class prediction for X , to which corresponds a predicted probability $\hat{P} = f(X)$, being $(x, y) \in (X, Y)$ a couple of random variables from the input and output domains (Silva Filho et al., 2023).

From this point of view, it becomes obvious that it is essential that, before performing classification, one should have the best possible probabilities to work with, especially for cost-sensitive decisions. Unfortunately, standard estimators do not necessarily provide this, often returning a ranking/score rather than a probability, and calibration should be performed. To improve model calibration, in T-FORS, Venn-ABERS predictors were leveraged, a form of conformal prediction with the ability to output not only better-calibrated probability estimates, but also a measure of confidence in these estimates (Vovk & Petej, 2014). Venn-ABERS predictors perform two isotonic regressions (using the greatest convex minorant of the cumulative sum diagram): one for the negative class, leading to scores p_0 , the other for the positive class, leading to scores p_1 . These two form an interval within which the true probability is deemed to be located. A single-valued probability can be obtained by combining these two results. Calibration is then assessed by means of the Brier score and reliability diagrams, which plot the observed frequency of an event against the forecast probability of an event. The Brier score is defined as:

$$B = \frac{1}{N} \sum_{i=1}^N (p_i - o_i)^2 \quad (5.9)$$

where $p_i \in [0, 1]$ is the prediction probability of the occurrence of the i -th event, and o_i is equal to 1 if the event occurred, 0 otherwise.

5.3.2 ARISTIDE – QUANT

The main difference between the T-FORS and ARISTIDE models is that while the former is considered a classification problem, the latter is framed as a time series classification task. In ARISTIDE, the input is the time evolution of the input features, rather than single values. The QUANT algorithm was applied to the input time series before being ingested into the classification algorithm. This choice was motivated by its competitive performance and

computational simplicity. The core observation behind QUANT is that it is possible to achieve accuracy comparable to more advanced algorithms simply by ordering the values in each interval and using them as features (Dempster et al., 2023). To reduce the dimensionality of the feature space, these ordered values can be sub-sampled, i.e., represented by quantiles computed within each interval. Figure 5.5 illustrates this principle, showing how an interval is subdivided into sub-intervals and how quantiles are extracted from each, down to the limiting case of a single value per sub-interval.




Figure 5.5: Illustration of the QUANT method: subdivision of an interval into sub-intervals and extraction of quantiles at each level, down to the limiting case of a single value per sub-interval.

The key aspect of the QUANT framework lies in its use of a single feature type (quantiles) computed over a hierarchy of dyadic intervals, obtained by recursively dividing each time window into two equal parts. These quantile features approximate the empirical distribution of values within each interval and efficiently encode the temporal structure of the input signals. One of the main advantages of this approach is that, once the quantile feature space has been generated, any standard classifier for tabular data can be applied. In the original QUANT formulation, extremely randomized trees with 200 estimators were used; however, for ARISTIDE, several classifiers were tested on top of the quantile features, including:

- CatBoost (Prokhorenkova et al., 2019)
- k-Nearest Neighbors (Guo et al., 2003)
- Random Forest (Breiman, 2001)
- Extremely randomized trees (Geurts et al., 2006)
- Histogram-based Gradient Boosting (Guryanov, 2019)
- XGBoost (Chen & Guestrin, 2016)
- AdaBoost with SAMME (Hastie et al., 2009)

The operational early-warning system is based on the model achieving the highest F-score during validation, as this metric balances precision and recall. Nonetheless, the multi-model design of ARISTIDE allows other classifiers to be substituted in production if needed, for

instance, to optimise for alternative metrics depending on operational priorities or user requirements.

The dataset was divided following a 3–1–1 temporal split, preserving chronological order. Specifically, three consecutive months were used for training, followed by one for validation and one for testing. Over the full database (2021–2024), this corresponds to approximately 30 months of training data and 9 months each for validation and testing. Given the natural imbalance of the LSTID catalogue, with less than 10% positive samples, random undersampling was applied to the training set to achieve a balanced ratio of 2 negative cases per positive one.

A complementary study was carried out to determine the optimal combination of training and control window lengths, the former defining the input history available to the model, and the latter the forecasting horizon. This analysis, performed using the Extremely randomized trees (ExtraTrees) classifier with 500 estimators, evaluated performance metrics across multiple window configurations to identify the combination yielding the best predictive skill. Once these parameters were fixed, final model hyperparameters were optimised using Optuna on the validation dataset (Akiba et al., 2019).

As in the T-FORS case, explainable AI techniques were integrated into ARISTIDE to interpret model behaviour. The TimeSHAP algorithm (Bento et al., 2021) extends the SHAP framework to sequential data, attributing importance not only to which drivers contributed most to a forecast, but also to when these contributions occurred within the input window. This provides both feature-level and temporal attributions, establishing a transparent link between the model’s statistical learning and the underlying physical mechanisms responsible for LSTID generation and propagation.

Finally, ARISTIDE also adopts a probabilistic forecasting formulation, in which the model outputs the probability that an LSTID will occur within the control window. To ensure that these probabilities are reliable, the model’s raw scores are post-processed using isotonic regression, a non-parametric calibration method that aligns predicted probabilities with the actual observed event frequencies (Berta et al., 2023). Unlike parametric approaches, isotonic regression does not assume the functional form of this mapping. Instead, it fits a monotonic, piecewise-constant function that preserves the rank ordering of the predictions while correcting distortions in the probability scale. Calibration is performed by comparing predicted and observed frequencies on an independent dataset and applying the derived mapping to all forecasts. This process guarantees that a forecasted probability (for instance, 0.8) corresponds to an empirical likelihood of roughly 80%. By converting nominal classification scores into statistically meaningful probabilities, isotonic calibration transforms ARISTIDE’s outputs into trustworthy early-warning indicators, suitable for operational use where quantified uncertainty is essential.

5.4 Results and comparisons

This section presents the main forecasting results for the T-FORS and ARISTIDE models and compares their performance across common evaluation criteria. The analysis covers predictive skill and interpretability, examining how each model captures the statistical and physical relationships between external drivers and LSTID occurrence.

Model performance is assessed using standard classification metrics (precision, recall, F1-score) to quantify the models’ ability to identify true events while limiting false alarms. Reliability is evaluated through probability calibration curves and the Brier score, which jointly measure how well the predicted probabilities correspond to observed event frequencies.

To interpret model behaviour, SHAP value analyses identify the most influential drivers and their temporal contributions to positive forecasts. This allows a direct comparison between

the dominant predictors in the ionosonde-based T-FORS model and the GNSS-based ARISTIDE system, highlighting how each responds to different manifestations of geomagnetic forcing.

Finally, the two models are compared: their differences primarily reflect the observational nature of their input catalogues (ionosonde versus GNSS) and their complementary sensitivity to local ionospheric conditions.

5.4.1a T-FORS – Hyperparameter Optimization

First, as mentioned in Section 5.3.1, an optimisation framework was included in the model development loop, acting on some of the model's hyperparameters. Those subjects to optimisation are:

- The maximum number of trees (iterations) that can be built – more trees should lead to greater accuracy, but also to longer training and inference time.
- The learning rate used for reducing the gradient step – a smaller (or larger) rate signifies that each tree offers a smaller (or larger) update to the model, resulting in gradual learning (or speeding up the learning process).
- The depth of the trees – the greater the depth, the greater the complexity.
- The minimum number of training samples in a tree leaf – the model will not search for new leaf splits when the sample count is less than this value.
- The weight for the positive class is used as a multiplier in the loss function in Equation 5.3 to penalise mistakes in positive samples more – a higher class weight means putting more emphasis on a specific class.
- The strength of the L2 regularisation term of the loss function serves as an additional handle to control the model's bias towards a specific class.

The values of these hyperparameters upon optimisation are shown in Table 5.3. Optuna sets a very reasonable configuration for the hyperparameters, choosing around 5,000 weak learners with a relatively shallow structure to avoid overfitting. Another expedient in this sense is selecting a minimum number of instances per leaf in the order of a few tens. The framework then fine-tunes the model by playing with a weight around double for the positive class with respect to the negative one, while attempting to balance false positives by providing a moderate L2 regularisation.

Table 5.3: Hyperparameters and their values after optimisation.

Hyperparameter	Optimised value
Iterations	4,997
Learning rate	0.0344
Depth	5
Minimum training samples in a leaf	44
Weights (neg. class, pos. class)	(1, 2.1874)
L^2 regularization	4.9454

5.4.1b T-FORS – Metrics, Explainability, and Calibration




Figure 5.6: Model output (Balanced prediction), target variable (derived from catalogue), and time series of some physical drivers for 11–15 March 2022.

Once the iterative optimisation process for model hyperparameters is finished, it is possible to visualise its output, relating it to the physical drivers and target variable that were the subject of the model’s learning process. An example of what is obtained is shown in Figure 5.6, referring to the days from 11 to 15 March 2022, in which one can see the model output in relation to the target variable and the time series of some of the physical drivers. The target variable is constructed from the HF-INT catalogue. Indeed, from the latter it derives the start and end time of the LSTID event, but extends it by 3 hours before the onset time, so that the algorithm learns the purpose of the modelling: estimating the occurrence of an LSTID in the next 3 hours.

First, a preliminary validation of the model is performed, evaluating the confusion matrix. Figure 5.7 shows the confusion matrix for the T-FORS model normalized by rows. The model correctly identifies 51.3% of the LSTIDs reported in the HF-INT catalogue, while ensuring a moderate incidence of false positives, amounting to 4.5% of the non-LSTID events.




Figure 5.7: Confusion matrix for the T-FORS model.

Beyond that, the SHAP values, introduced in Section 5.3.1, help understand the model's decision-making process and validate its operation from a physical point of view. To obtain a global measure, the mean absolute SHAP value of each feature was computed (Figure 5.8). This highlights whether the features deemed most relevant by the model align with physical expectations. Among the most impactful variables (on average) on the model's decisions, there are:

- The HF-INT_{EU} index and its 2-hour moving average are an important proxy for the presence of HF activity on the European sector related to LSTIDs.
- Solar zenith angle controls the intensity of solar radiation reaching the ionosphere and its ionisation.
- IU and IE indices, along with their moving averages and variations, measure geomagnetic activity over the European sector due to the solar wind's injection of energy into the magnetosphere-ionosphere system.
- F10.7, which is a proxy for solar activity.
- Hp30 is a Kp-like index that measures the global geomagnetic activity due to interactions between solar wind and the Earth's magnetosphere.
- Solar wind density can increase particle flux, enhancing geomagnetic activity and transferring more energy into the ionosphere. Higher solar wind densities can generate and sustain LSTIDs.

The picture returned by the model using XAI methods is non-trivial, especially given the non-linear nature of the phenomenon and the model itself. As such, the SHAP formalism is an important tool in debugging ML methods with sensible domain knowledge – i.e., known physics principles.




Figure 5.8: Ranked feature importance as the average SHAP value, representing the input features' impact on the model output.

The cross-validated performance of the model in terms of Precision (Equation 5.6), Recall (Equation 5.4), and F1-score (Equation 5.6) is shown in Table 5.4. As a baseline to interpret the goodness of metrics in Table 5.4, one can consider a naïve, rule-based model: if the HF-INT_{EU} index rises above 1.75, the model predicts the occurrence of an LSTID in the next 3 hours; no LSTIDs are predicted otherwise. This simple model achieves a precision of 0.79 and a recall of 0.17, thus a F1-score of 0.29. Precision is quite high since the catalogue is a human-validated version of the HF-INT_{EU} index; however, recall is very low (only one in six predicted events occurs), thus pushing the F1-score down. It becomes apparent how the model trades precision for recall, yielding a F1-score gain of +93%.

Table 5.4: T-FORS and Aristide model performance.

	F1-Score	Precision	Recall
T-FORS	0.56	0.62	0.51
ARISTIDE	0.63	0.57	0.70

As mentioned in previous sections, classification is a post hoc decision layer on top of a probabilistic prediction; thus, it involves making a definitive decision that can be applied incorrectly. From the point of view of space weather forecasting, calibrated class probabilities are essential. However, classifiers do not generally yield class probabilities; at best, they offer classification scores that approximate true probabilities (Johansson & Gabrielsson, 2019). In potentially high-risk settings, such miscalibrations can become hazardous and invalidate decision-making. However, incorporating a calibration phase, one can output not only calibrated probability estimates but also a measure of confidence in these estimates, leading to conceivably safer outcomes.

The reliability plot of the uncalibrated model is reported in Figure 5.9, where the horizontal axis reports the predicted probability, while the vertical axis shows the observed relative frequency of an event.

It becomes apparent how the uncalibrated model shows a significant imbalance from the point of view of the positive class. For instance, a predicted mean probability of 0.75 corresponds to an actual fraction of positives of only 30% (instead of 75%). In general, the model tends to strongly overestimate the probability of occurrence of LSTIDs, especially in the range of estimated probabilities between 0.1 and 0.9 (approximately). The achieved Brier score is 0.120.

The Brier score achieved after Venn-ABERS calibration is 0.078. The calibration is still imperfect, but the model exhibits a more predictable and safe behaviour.




Figure 5.9: Calibration curves for the raw and calibrated T-FORS forecasting model. The dashed line represents a perfectly calibrated model, while the dotted line shows the actual calibration curve, illustrating how predicted probabilities correspond to true outcomes.

5.4.2a ARISTIDE – Hyperparameter optimization

The first step in evaluating the ARISTIDE forecasting framework concerns the selection of the optimal combination of training and control window lengths. The training window defines the duration of the input history provided to the model, while the control window specifies the prediction horizon. These parameters jointly determine the temporal context available for the model to infer input–output relationships.

Table 5.5 reports the classification metrics (F1-score, precision, and recall) obtained for different training and control window size combinations using the Extra Trees classifier with 500 estimators. This experiment aimed to identify the configuration that maximises overall predictive skill, as quantified by the F1-score.

Results show that all configurations produce comparable performance, with F1-scores ranging between 0.613 and 0.626. A 4-hour control window and a 4-hour training window achieve the highest F1-score (0.626), slightly outperforming the neighbouring configurations. This suggests that the model benefits from a sufficiently long input history to capture the variability of the external drivers, while a control horizon of 4 hours remains well aligned with the typical propagation times of LSTIDs from high to mid latitudes.

The minor trade-offs observed between precision and recall across different window sizes highlight that the chosen temporal context influences the model’s ability to distinguish precursor activity from background variability but does not drastically alter its overall skill.

Based on these results, the configuration 4h/4h was adopted to optimize and evaluate the ARISTIDE model. This setting should provide the best compromise between lead time,

stability, and physical interpretability, ensuring that the forecasting window remains consistent with the temporal scales of LSTID generation and propagation.

Table 5.5: classification metrics for different training and control windows. Numbers in bold show the best value for each column. The model considered is the ExtraTrees classifier.

Window Size		F1	Precision	Recall
Control	Training			
4h	4h	0.626	0.595	0.661
3h	4h	0.620	0.577	0.669
2h	4h	0.623	0.568	0.689
4h	3h	0.622	0.581	0.668
3h	3h	0.613	0.561	0.676
2h	3h	0.617	0.549	0.705
4h	2h	0.622	0.580	0.670
3h	2h	0.623	0.572	0.685
2h	2h	0.619	0.550	0.708

Once the optimal training and control window lengths were established, different classifiers were tested on top of the QUANT feature representation to identify the most effective modelling approach for the ARISTIDE forecasting system. Table 5.6 summarises the classification metrics obtained for each algorithm using a 4-hour training window and a 4-hour control window.

Overall, performance across models is consistent, with F1-scores ranging between 0.49 and 0.63. The CatBoost classifier achieved the highest F1-score (0.628), closely followed by the Histogram-based Gradient Boosting, Random Forest, and Extra Trees models. These results confirm the suitability of ensemble-based algorithms for capturing the non-linear and multivariate dependencies that characterise the relationship between geomagnetic drivers and LSTID occurrence.

Despite CatBoost's marginally higher performance, the Extra Trees classifier was ultimately selected for the operational configuration of ARISTIDE. Its computational simplicity, faster training time, and reduced tuning requirements make it more practical for large-scale or near-real-time forecasting applications, while maintaining nearly identical predictive skill ($F1 = 0.626$).

Simpler learners such as AdaBoost and k-Nearest Neighbors showed significantly lower F1-scores, confirming that the LSTID forecasting problem benefits from ensemble averaging and feature-randomisation strategies that enhance generalisation in the presence of noisy, correlated predictors.

Based on this analysis, the Extra Trees model with a 4-hour training and 4-hour control window was selected as the final ARISTIDE configuration for further investigation, including feature ranking through TimeSHAP, probabilistic calibration, and reliability evaluation.

After identifying the Extra Trees classifier as the most suitable algorithm for the ARISTIDE forecasting framework, a systematic hyperparameter optimization was done using Optuna to maximise the F1-score on the validation set. The optimisation explored parameters governing tree structure, ensemble size, and sampling strategy to achieve the best trade-off between accuracy, robustness, and computational efficiency.

Table 5.6: Classification metric for different machine learning models applied inside QUANT for a training and control window of 4 hours. Numbers in bold represent the highest score for each column.


Model	F1	Precision	Recall
CatBoost	0.628	0.574	0.692
AdaBoost	0.596	0.499	0.738
XGB	0.596	0.550	0.650
HistGradientBoosting	0.627	0.556	0.719
ExtraTrees	0.626	0.595	0.661
RandomForest	0.627	0.562	0.708
k-NearestNeighbors	0.490	0.421	0.584

The best-performing configuration achieved an F1-score of 0.631 on the validation dataset with the following parameters: 450 estimators, a maximum depth of 16, a minimum of 17 samples required to split a node, and one sample per leaf. All predictors were considered for each split, and bootstrap sampling was enabled to improve ensemble diversity. The splitting criterion was set to log loss (Equation 5.3), promoting probabilistic accuracy rather than purely categorical separation, while no class weighting was applied since balance between positive and negative cases had already been ensured through random undersampling.

This configuration ensures a model that is expressive yet stable. A maximum depth of 16 allows the ensemble to capture non-linear interactions among the geophysical drivers without overfitting, while requiring at least 17 samples per split reduces sensitivity to noise and rare patterns. The use of 450 estimators provides sufficient ensemble averaging to stabilise forecasts without significantly increasing computational cost, which remains a critical consideration for near-real-time applications.

5.4.2b ARISTIDE – Metrics, Explainability, and Calibration

The final optimised model was then evaluated on the independent test set to assess generalisation and to benchmark its performance against the T-FORS forecasting system. The resulting scores are summarised in Table 5.4, which reports the F1-score, precision, and recall for both models.

**Figure 5.10:** Confusion matrix for the ARISTIDE model.

As shown in Table 5.4, the ARISTIDE model achieves an F1-score of 0.63, outperforming the T-FORS model ($F1 = 0.56$) across all metrics except precision. The higher recall (0.70) highlights ARISTIDE's ability to correctly identify a greater proportion of true LSTID events, which is crucial for early-warning purposes where missing an event is more critical than issuing a false alarm. The slightly lower precision (0.57) reflects a broader sensitivity to weaker or borderline disturbances, some of which may not meet the detection thresholds of the ionosonde-based catalogue.

Conversely, the T-FORS model attains higher precision (0.62) but lower recall (0.51), consistent with its more selective, expert-validated detection strategy. This difference is also visible in the corresponding confusion matrices (Figures 5.7 and 5.10), where T-FORS shows a lower false-alarm rate but misses a larger fraction of true events, while ARISTIDE achieves a more balanced distribution between correctly detected and missed occurrences.

It is also important to note that the two catalogues differ in temporal and solar-cycle coverage: ARISTIDE spans 2021–2024, corresponding to the ascending phase of Solar Cycle 25, whereas T-FORS covers 2014–2022, nearly a full solar cycle, including several years of low solar activity. The broader temporal coverage and variable geomagnetic forcing in the T-FORS dataset likely contribute to its lower overall recall, as periods of low solar activity are characterised by weaker ionospheric variability and reduced event detectability.

These complementary characteristics suggest that the two systems capture different aspects of the same physical process: ARISTIDE provides a more responsive forecasting of mid-latitude disturbances during active solar conditions, while T-FORS offers a conservative, more generalized model grounded in expert-validated detections. Combining their outputs could yield a more robust and adaptable framework for LSTID early warning across the European sector.




Figure 5.11: Ranked feature importance as the average absolute SHAP value, representing the input features' impact on the model output.

TimeSHAP analysis was applied to the final trained ensemble to interpret the physical meaning of the ARISTIDE model predictions. This method quantifies the contribution of each input feature to the model's probabilistic output, allowing the identification of the most influential drivers and their temporal relevance.

Figure 5.11 shows the ranked feature importance expressed as the mean absolute SHAP value across the test set. The Hp30 index emerges as the dominant predictor, confirming the

strong dependence of LSTID activity on subauroral geomagnetic disturbances. The solar zenith angle (SZA) follows closely, highlighting the importance of the ionospheric illumination state in modulating both the generation and detectability of disturbances. The F10.7 solar flux index and the ap30 and IL indices also appear among the leading contributors, indicating that background solar activity and auroral electrojet intensity influence the occurrence probability. Moderate contributions from IMF Bz and total magnetic field B further emphasise the role of solar wind–magnetosphere coupling.

Figure 5.12 presents the ranked event-level SHAP importance, which describes the temporal relevance of input samples relative to the forecast time. The largest contributions correspond to $t = 0$ and $t = -1$, indicating that the model primarily bases its forecasts on the most recent 30–60 minutes of driver evolution, while retaining informative context up to three hours earlier. This temporal decay is physically consistent with the propagation times of LSTIDs from their high-latitude source regions to mid-latitude observation points. These analyses confirm that the ARISTIDE model identifies the correct physical drivers of ionospheric disturbances and attributes them on realistic temporal scales, reinforcing the model's interpretability and physical consistency with known LSTID generation mechanisms.




Figure 5.12: Ranked event importance as the average absolute SHAP value, representing the input features' impact on the model output. Event indexes represent the input time series samples with respect to the time of forecasting.




Figure 5.13: Calibration curves for the raw and calibrated ARISTIDE forecasting model. The dashed line represents a perfectly calibrated model, while the dotted line shows the actual calibration curve, illustrating how predicted probabilities correspond to true outcomes.

Calibration curves were computed before and after the isotonic regression step to assess the reliability of the probabilistic forecasts generated by the ARISTIDE model. Figure 5.13 shows the relationship between the predicted probabilities and the corresponding observed event frequencies and the Brier score.

The uncalibrated model (left panel) yields a Brier score of 0.087, indicating a generally reliable performance but with underconfidence, as the curve mostly lies below the ideal diagonal. After isotonic calibration, the Brier score decreases to 0.074, and the calibration curve aligns more closely with the one-to-one line, showing that predicted probabilities now correspond more accurately to the observed frequency of LSTID occurrence.

However, a moderate degradation in calibration is observed for predicted probabilities above ~ 0.7 , where the curve deviates from the diagonal. This behaviour can be explained primarily by data imbalance: high-probability forecasts correspond to rare cases, meaning that each probability bin in this range contains few samples and thus exhibits greater statistical uncertainty. Furthermore, these high-probability forecasts should occur during the most intense geomagnetic disturbances, conditions in which the automatic GNSS-based LSTID detection becomes less reliable, as discussed in Chapter 4. During such extreme storms, increased background noise reduces the clarity of keogram features, making true event identification more difficult.

Overall, isotonic calibration corrects the general bias of the raw model outputs, yielding a well-aligned and statistically reliable probability distribution. Despite the residual uncertainty at the highest probability range, ARISTIDE provides meaningful probabilistic forecasts suitable for operational use in early-warning and risk-assessment contexts.




Figure 5.14: Time series of SymH, Hp30, IMAGE electrojet indices, and solar wind magnetic field time series for 09–16 February 2022. The bottom panel shows the dTEC keogram from which the catalogue is created, showing positive (red) and negative (blue) deviations from the background, with the probabilistic LSTID forecasting obtained with the calibrated model superimposed.

Figure 5.14 presents an example of the calibrated ARISTIDE model output for 9–12 February 2022, alongside the main external drivers and the corresponding dTEC keogram. The panels show the evolution of the Sym-H, Hp30, and auroral electrojet indices (IU, IL) and the interplanetary magnetic field magnitude and Bz component, while the bottom panel displays the GNSS-derived keogram. The black line represents the model’s calibrated probability of LSTID occurrence.

The probabilistic forecast closely follows the temporal evolution of geomagnetic and auroral activity. The main peaks in predicted probability coincide with periods of enhanced Hp30 and electrojet intensity, depressed SymH, and strong southward Bz, all characteristic signatures of magnetospheric energy input capable of exciting large-scale travelling disturbances. Comparing the LSTID probabilistic prediction with the keogram, it is possible to

see how an increased LSTID probability corresponds to oblique bands in the dTEC keograms, visible between 16:00 and 04:00 on the 10-11 February and between 20:00 and 24:00 for the 11th. Another notable feature is that the LSTID probability starts to increase before the appearance of LSTIDs in the keograms, showing the forecasting capability of the model. Figure 5.15 helps understand which features mattered in raising the LSTID forecasting probability; specifically, for the forecasting of 10 Feb 2022 at 14:30, the highest contributions belong to Hp30, solar zenith angle, and solar wind magnetic field.

This example illustrates how the calibrated ARISTIDE model produces physically consistent probabilistic forecasts that track the development of ionospheric disturbances in real time. The model's output reflects the strength and timing of magnetospheric forcing, providing a quantitative early-warning measure of LSTID likelihood across European mid-latitude regions.



Figure 5.15: Event-level SHAP force plot for 10 February 2022 at 14:30 UT, showing each feature's contribution to the model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.

5.4.3 Model comparison on case event

A comparative analysis of the T-FORS and ARISTIDE model outputs during specific disturbed intervals is essential to evaluate both their forecast accuracy and the physical consistency of the mechanisms each model has learned. Examining for the same geomagnetic disturbances enables a qualitative assessment of differences in sensitivity, response timing, and feature relevance, offering insight into how each model captures the magnetosphere–ionosphere coupling processes that drive LSTID occurrence.

The first period analyzed, summarized in Figure 5.16, spans 9–13 February 2022, encompassing several moderate disturbances. Figure 5.16a shows the forecasted probability of LSTID occurrence from both models, revealing a systematic delay between the two outputs. This behavior is clarified by comparison with the observational catalogues in Figures 5.16f and 5.16g, which display LSTID detections from HF-INT and GNSS sources, respectively. The two catalogues differ markedly, reflecting the intrinsic differences between their observing techniques. GNSS-based detection often misses post-midnight LSTIDs during moderate activity, while HF-based methods can be disrupted by sporadic E layers or spread-F conditions that prevent reliable ionogram acquisition. Despite these limitations, each model performs well relative to the dataset it was trained on: T-FORS with HF detections and ARISTIDE with GNSS events. Both models also exhibit enhanced forecast probabilities even when no event is recorded in the corresponding catalogue. This expected behavior indicates that both systems have learned to quantify the LSTID-genic potential of the geomagnetic conditions. The coupling efficiency at high latitudes is further modulated by the solar zenith angle and F10.7 flux, which help the models account for climatological variability of the two catalogues linked to solar illumination and solar-cycle phase, factors that differently influence the performance of HF and GNSS detection systems.




Figure 5.16: Summary plot of ARISTIDE and T-FORS probabilistic forecastings, LSTID measurements, and input time-series for 9-13 February 2022. Panel (a) shows the probability of occurrence of an LSTID in the next 3/4 hours in blue/orange for T-FORS/ARISTIDE. Panels (b) and (c) show the SymH and Hp30 time series. Panel (d) shows the time evolution of the IMAGE electrojet indices IL (blue) and IU (orange). Panel (e) shows the solar wind magnetic field magnitude (blue) and vertical component (orange). Panel (f) shows the HF-INT_{EU} index value and the HF-based LSTID catalogue from which we built the T-FORS model output. Panel (g) shows the dTEC latitudinal keogram with the GNSS-based LSTID catalog superimposed.

To better understand the internal decision process, Figures 5.17 and 5.18 present event-level SHAP force plots for two representative events. The first case (Figure 5.17) corresponds to 10 February 2022 at 18:00 UT, a moderate-activity period where ARISTIDE outputs a high probability (0.93), mainly driven by Hp30, IL, and magnetic-field components (B, Bz). T-FORS, in contrast, predicts a lower probability (0.33), dominated positively by the solar zenith angle and short-term variations of IU, while the low HF index (a direct measure of LSTID) lowers the LSTID onset likelihood. The second case (Figure 5.18) illustrates 13 February 2022

00:00 UT. In this case, T-FORS produces a very confident prediction (0.97), caused by the detection of LSTID represented by a high HF-INT_{EU} index and the concurrent intensification of geomagnetic activity. ARISTIDE, on the other hand, assigns a more moderate probability (0.37), with geomagnetic indices (Hp30 and IL) pushing towards a positive LSTID probability, while solar zenith angle limits the forecast due to the low background ionization at the time of forecast (00:00).

Overall, these case studies give a first glimpse of complementary behaviors: T-FORS responds directly to near-real-time ionospheric indicators while weighting geomagnetic drivers, whereas ARISTIDE incorporates a broader magnetospheric and solar-wind context, compensating for the absence of direct LSTID observations in its inputs. Their comparison demonstrates that combining both frameworks yields a more complete depiction of the space-weather conditions underlying large-scale ionospheric disturbances, mitigating the individual limitations of HF- and GNSS-based sensing approaches.



Figure 5.17: Event-level SHAP force plot for 10 February 2022 at 18:00 UT, showing each feature's contribution to the ARISTIDE (top) and T-FORS (bottom) model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.




Figure 5.18: Event-level SHAP force plot for 13 February 2022 at 00:00 UT, showing each feature's contribution to the ARISTIDE (top) and T-FORS (bottom) model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.




Figure 5.19: Summary plot of ARISTIDE and T-FORS probabilistic forecastings, LSTID measurements, and input time-series for 11–13 July 2022. Refer to Figure 5.16 for panel description.

The second case study focuses on 11–13 July 2022, summarized in Figure 5.19, a period marked by low-moderate geomagnetic activity and multiple nighttime LSTID detections in GNSS and HF. This interval offers an additional test of the models under different background ionospheric conditions compared with the February 2022 case, particularly higher solar illumination and increased daytime coupling efficiency.

Figure 5.19a shows that both models predict a significant rise in LSTID probability between the sunset of the 11th and the midnight of the 12th, consistent with the HF-INT_{EU} index and GNSS-based detections (Figures 5.19f and 5.19g). ARISTIDE maintains a smoother and higher probability enhancement, while T-FORS exhibits sharper, burst-like increases closely synchronized with HF-INT_{EU} variations. Despite this, both models respond coherently to the geomagnetic evolution seen in Hp30, IL, IU, and Bz, indicating that their internal feature weightings are physically consistent with the external drivers and are modulated by other parameters, such as the HF-INT_{EU} for T-FORS and the solar zenith angle for ARISTIDE.




Figure 5.20: Event-level SHAP force plot for 12 July 2022 at 00:00 UT, showing each feature's contribution to the ARISTIDE (top) and T-FORS (bottom) model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.




Figure 5.21: Event-level SHAP force plot for 12 July 2022 at 04:30 UT, showing each feature's contribution to the ARISTIDE (top) and T-FORS (bottom) model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.



Figure 5.22: Event-level SHAP force plot for 12 July 2022 at 12:00 UT, showing each feature's contribution to the ARISTIDE (top) and T-FORS (bottom) model prediction. Red/blue arrows mark positive/negative contributions that increase/decrease the probability relative to the base value.

To clarify their internal decision mechanisms, Figures 5.20, 5.21, and 5.22 illustrate event-level SHAP force plots for three representative times on 12 July.

During the first event (00:00 UT, Figure 5.20), ARISTIDE predicts a fairly high probability (0.88) raised mainly by Hp30, ap30, and IL, reflecting the ongoing magnetospheric coupling and reduced by the unfavorable solar zenith angle conditions (nighttime). T-FORS returns an even higher probability (0.97), largely influenced by HF-INT and its short-term moving average. This indicates that both models identify favorable disturbance conditions, but T-FORS places most weight on direct LSTID measurement rather than geomagnetic indices.

The second event (04:30 UT, Figure 5.21) corresponds to a peak in LSTID activity. Here, T-FORS again predicts a very high probability (0.96), dominated again by HF-INT_{EU} and the IU variations, confirming its direct sensitivity to real-time detection proxies. ARISTIDE, in contrast, outputs a moderate value (0.49), primarily influenced by Hp30, IL, and ap30, with negative contributions from the solar zenith angle.

By comparison, the midday event (12:00 UT, Figure 5.22) highlights a situation where ARISTIDE and T-FORS diverge most. ARISTIDE predicts a relatively high probability (0.84), driven by Hp30, IL, and solar zenith angle, consistent with its climatological weighting of daytime favourable detection. Instead, T-FORS produces a much lower probability (0.12), as the solar zenith angle and HF-INT suppress the model's response, with the latter representing typically unfavorable conditions for ionosonde-based detections despite ongoing magnetospheric forcing.

The July 2022 analysis confirms the complementary behavior observed in February 2022. T-FORS reacts sharply to real-time ionospheric measurements and manages to reproduce the HF-INT-based catalogue, while ARISTIDE tracks the broader geomagnetic context with smoother transitions due to the absence of inputs related to LSTID activity. This behaviour was also visible in the confusion matrices shown in Figures 5.7 and 5.10, with ARISTIDE showing an amount of false positives almost twice that of T-FORS (8.6% and 4.5%). On the other hand, being T-FORS tightly related to the HF-INT_{EU}, when the index is low, the model will output a negative forecast despite the favourable geomagnetic coupling conditions, something that ARISTIDE does not suffer (see, for example, Figures 5.21 and 5.22). Again, this is confirmed by the confusion matrices, where T-FORS shows a number of false negatives 60% higher than ARISTIDE (48.8% and 30.6%). The two models describe different facets of the same physical process: T-FORS capturing the immediate ionospheric response and ARISTIDE representing the larger-scale energy transfer and coupling efficiency that precedes and sustains LSTID propagation. Moreover, the two case events highlighted the importance of not limiting the forecast to a binary output, but rather framing it in its probabilistic version.

Conclusions and future work

The primary goal of this doctoral thesis was to develop a reproducible and data-driven framework for the systematic detection and forecasting of Large-Scale Traveling Ionospheric Disturbances (LSTIDs). This research successfully addressed the necessity for automated monitoring of these wave-like perturbations that degrade GNSS and communication systems. The work achieved this through integrating physical modeling, advanced signal processing, and machine learning approaches, predominantly utilizing GNSS-based TEC observations. The methodology followed a structured progression: quantifying data processing errors (Chapter 3), establishing an autonomous detection system (Chapter 4), and culminating in the development of forecasting models (Chapter 5). The successful completion of this multi-phase pipeline has yielded a comprehensive operational system that significantly advances the field of ionospheric space weather monitoring. In fact, to the best of current knowledge, no other model for LSTID forecasting has been developed apart from the one presented in this doctoral thesis.

The foundational work focused on eliminating systematic errors inherent in TEC processing, a necessary step for a reliable and accurate TID detection. This phase involved quantifying errors using a specially constructed synthetic TEC measurements database, where known disturbance models were superimposed on what derived from real GNSS measurements. The main findings from the study are the following:

- A systematic comparison of six common signal processing methods established that the Fast Iterative Filtering and the Savitzky-Golay filter were the most reliable methods for extracting wave amplitude and preserving frequency content. This finding ensures the highest fidelity of the parameters derived for LSTIDs and MSTIDs. For example, in the LSTID scenario studied in Chapters 4 and 5, when applying a Savitzky-Golay filter to the synthetic TEC database, the amplitude error is lower than 0.15 TECu for 80% of the samples. Furthermore, acknowledging computational cost, the SGOLAY method was identified as the best operational tool in time-critical systems due to its high speed, being approximately 300 times faster than FIF.
- The research quantified how observational geometry introduces error into period retrieval, a factor previously neglected in many studies. The Doppler-induced error arises from the relative velocity between the TID wavefront and the GNSS Ionospheric Piercing Point (IPP). The error increased significantly when lowering the elevation cut-off angle or increasing the assumed ionospheric shell height (which can be considered as the impact of different heights of the F2 layer). For the critical MSTID scenario (slower waves), the worst-case configuration (20° elevation and 350 km height) resulted in the expected period error exceeding 50% in 20% of all samples. This analysis provides essential constraints for setting observational criteria and interpreting wave parameters.
- The comparison between the reliable post-processing calibrated data (used as ground-truth) and the NRT derivable observable has shown that the two verticalized observables, $vTEC_{Ne}$ and $vTEC_{PH}$ assure a high degree of precision in the amplitude estimation. Under the LSTID scenario, the two vertical observables showed negligible amplitude error (80% CDF lower than 0.01 TECu). In the MSTID case, the NeQuick-calibrated one ($vTEC_{Ne}$) showed better performance, with an 80% CDF of 0.01 TECu against the 0.035 TECu for $vTEC_{PH}$.

The methodologies developed in Chapter 3 enabled a more reliable deployment of a fully autonomous LSTID detection algorithm presented in Chapter 4. The presented algorithm provides robust identification of LSTIDs, extracting key parameters such as propagation speed and period from keogram images. The approach successfully identifies LSTIDs during quiet and disturbed periods, as demonstrated through case studies and comprehensive statistical analysis for the European sector spanning 2021–2024. Moreover, this thesis showed that the same technique can be applied to other geographical areas with comparable performances, requiring only some parameter tuning to tackle differences in the GNSS networks. The key findings of the climatological analysis of Chapter 4 are the following:

- LSTIDs with periods falling in the 30-90 minutes range are more frequently observed during equinoxes and are most commonly detected during mildly disturbed geomagnetic conditions (e.g., H_p30 values between 3 and 5 or $SymH$ values between -25 and -15 nT). The probability of detection increases significantly during periods of high geomagnetic activity, with detection rates reaching up to 50%.
- The results confirm a direct relationship between LSTID occurrence and geomagnetic indices such as H_p30 , $SymH$, and IMAGE IE. Increased geomagnetic activity leads to LSTIDs characterized by shorter wavelengths, higher speeds, and greater amplitudes, consistent with known generation mechanisms tied to Joule heating and associated atmospheric disturbances.
- The results are consistent with previous studies conducted over different regions and periods using dTEC measurements (Tsugawa et al., 2004; Ding et al., 2008; Borries et al., 2009). In particular, they closely align with those of Borries et al. (2009), who reported an average period of 59 minutes and a meridional phase velocity of 684 m/s, values that are in good agreement with the observed period of 62 minutes and meridional phase speed of 629 m/s. When compared with results from other geographic areas, some differences emerge. Tsugawa et al. (2004) identified a bimodal seasonal distribution, with a peak in LSTID occurrence during the spring equinox, whereas this analysis showed a maximum during the autumn equinox. Similarly, Ding et al. (2008) reported different average values for period and velocity. Their study also revealed a different dependence on local time. The discrepancies in speed and period across these studies can be attributed to several factors, including differences in the inclination and declination of the geomagnetic field, variations in geomagnetic latitude ranges, and, most importantly, differences in detection methodology.

Thanks to the new LSTID catalogue generated through the autonomous detection pipeline presented in Chapter 4, and by leveraging an independent event catalogue derived from the TechTIDE ionosonde network with the HF-INT method, it was possible to take a decisive step toward LSTID forecasting. Chapter 5 built upon this foundation to explore data-driven predictive modelling of ionospheric disturbances, addressing the challenge of anticipating LSTID occurrence, an objective of significant operational relevance for space-weather applications. The work presented in Chapter 5 led to several key outcomes:

- Two complementary machine-learning models were designed and validated: T-FORS and ARISTIDE. Both operate in a probabilistic forecasting mode, with horizons of 3 and 4 hours, respectively. Each model integrates geomagnetic indices describing the solar wind-magnetosphere coupling. In addition, T-FORS combines the geomagnetic context with near-real-time ionospheric information from the HF-INT_{EU} index, thereby including direct evidence of whether LSTID activity is already ongoing. ARISTIDE, by contrast, relies solely on external drivers without any ionospheric observables

among its inputs. This structural difference has important consequences. T-FORS is tightly coupled to the instantaneous ionospheric response and thus can forecast disturbances already developing in the ionosphere with high precision, but has limited sensitivity when activity is not ongoing. ARISTIDE, on the other hand, interprets the magnetospheric forcing environment independently of current ionospheric observations. It captures the potential for LSTID generation some hours in advance, providing earlier but less specific forecasts. Consequently, while T-FORS minimizes false positives and maintains high precision, ARISTIDE offers higher recall and anticipates events that have not yet produced measurable HF signatures.

- By applying SHAP and TimeSHAP analyses, the models' decisions were linked to well-known physical processes. Features such as the Hp30, ap30, and IMAGE electrojet indices, interplanetary magnetic-field components (particularly B_z), and the solar-zenith angle emerged as dominant predictors. Their temporal importance patterns revealed that enhanced geomagnetic coupling and night-time ionospheric conditions increase the probability of LSTID onset, a result fully consistent with the established understanding of sub-auroral and thermospheric wave generation. This interpretability ensures that the models are not mere statistical predictors but physically grounded diagnostic tools.
- The probabilistic nature of the models was refined through post-processing calibration techniques (isotonic and Venn-Abers methods). This calibration transformed raw classifier outputs into reliable probability estimates, a critical step for operational use. The resulting reliability diagrams showed good alignment between forecast probabilities and observed event frequencies, allowing users to interpret forecasts in terms of risk levels rather than binary predictions.
- The comparison between T-FORS and ARISTIDE revealed a clear complementarity. T-FORS produced sharper responses with lower false-positive rates but a higher false-negative rate under low-activity conditions. Conversely, ARISTIDE yielded smoother transitions with slightly more false positives but fewer missed detections. Their joint operation constitutes the first demonstration of an explainable system capable of producing consistent, physically interpretable LSTID forecasts.

The developments of Chapter 5 thus complete the methodological pipeline envisioned at the start of this thesis, from systematic data correction and signal processing (Chapter 3), through autonomous detection and catalogue generation (Chapter 4), to probabilistic, explainable forecasting (Chapter 5). This sequential structure demonstrates how data-driven approaches can reproduce and anticipate ionospheric disturbances with accuracy and interpretability when carefully constrained by physical understanding.

The results of this thesis establish a reproducible and physically interpretable framework for LSTID detection and forecasting. Nevertheless, several developments remain necessary to consolidate its scientific maturity, expand its applicability, and enable operational use. The following directions outline the main priorities for future work, focusing on the extension, real-time implementation, and integration of the proposed systems.

- The autonomous GNSS-based detection method, developed and validated in Europe, should be applied to additional geographical regions. Extending its application across different geomagnetic environments will make it possible to characterise longitudinal and hemispheric differences in LSTID occurrence, propagation direction, and seasonal variability. This transfer will require parameter tuning, but the core algorithm remains applicable, as demonstrated. The outcome would be a standardised, globally comparable dataset of LSTID activity derived from a consistent methodology.

- The current detection method outputs binary detections without an explicit confidence score. Introducing a quantitative quality metric would enable automatic detection, reliability, and uncertainty assessment. The score could incorporate factors such as the signal-to-noise ratio of the keogram ridge, coherence across latitude bands, internal consistency of phase velocity, and data coverage. This metric would allow for filtering low-confidence events, weighting training data in model development, and transparent communication of confidence levels in operational use.
- A key step toward operational readiness is to deploy the detection chain in near-real-time mode. Processing live GNSS streams and updating the keograms continuously would enable continuous ionosphere monitoring and automatically generate event alerts. This will require automated data ingestion and efficient computational implementation to handle large volumes of incoming data. A real-time detection system would form the foundation of an LSTID nowcasting service, directly supporting forecasting models and space-weather operations.
- The current LSTID catalogue spans only a few years within Solar Cycle 25. Extending it to at least one full solar cycle will ensure representation of all geomagnetic conditions, from solar minimum to maximum, and provide a statistically complete picture of long-term variability. Such a catalogue will enable broader climatological studies of solar-cycle dependence and background ionospheric conditions influencing LSTID propagation. It will also provide a richer and more balanced dataset for retraining and benchmarking forecasting models.
- The forecasting models developed in Chapter 5, T-FORS and ARISTIDE, were trained on retrospective datasets. Future implementations should be adapted to ingest real-time data streams of geomagnetic indices, solar-wind parameters, and ionospheric indicators such as the HF-INT_{EU} index. Continuous retraining or adaptive calibration on live inputs will allow the models to remain robust against evolving solar and geomagnetic conditions, ensuring consistent reliability and interpretability in operational environments.
- LSTIDs are physical phenomena independent of the detection technique, yet every instrument and method samples them with its limitations. To achieve a more faithful representation of true LSTID occurrence, independent catalogues should be merged through cross-verification. This would reduce method-dependent biases and provide a consolidated “ground truth” reference catalogue. Such an integrated product would support both physical analyses and supervised training of future hybrid forecasting systems.

In summary, the next stage of this research should focus on scaling the detection and forecasting framework from a regional, research-oriented demonstrator to a global, real-time, and quality-controlled system. By extending spatial coverage, ensuring continuous operation, retraining models on live data, and integrating multiple observational sources, the framework can evolve into a comprehensive capability for monitoring and predicting ionospheric disturbances. Such developments will deepen understanding of Sun–Earth coupling and support the operational needs of space-weather forecasting and GNSS-dependent technologies worldwide.

Bibliography

- Afraimovich, E. L., Astafieva, E. I., & Kirushkin, V. V. (2006). Localization of the source of ionospheric disturbance generated during an earthquake. *International Journal of Geomagnetism and Aeronomy*, 6(2), GI2002. <https://doi.org/10.1029/2004GI000092>
- Afraimovich, E. L., Palamartchouk, K. S., & Perevalova, N. P. (1998). GPS radio interferometry of travelling ionospheric disturbances. *Journal of Atmospheric and Solar-Terrestrial Physics*, 60(12), 1205–1223. [https://doi.org/10.1016/S1364-6826\(98\)00074-1](https://doi.org/10.1016/S1364-6826(98)00074-1)
- Afraimovich, E. L., Perevalova, N. P., Plotnikov, A. V., & Uralov, A. M. (2001). The shock-acoustic waves generated by earthquakes. *Annales Geophysicae*, 19(4), 395–409. <https://doi.org/10.5194/angeo-19-395-2001>
- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A Next-generation Hyperparameter Optimization Framework (No. arXiv:1907.10902). arXiv. <https://doi.org/10.48550/arXiv.1907.10902>
- Altadill, D., Segarra, A., Blanch, E., Juan, J. M., Paznukhov, V. V., Buresova, D., Galkin, I., Reinisch, B. W., & Belehaki, A. (2020). A method for real-time identification and tracking of traveling ionospheric disturbances using ionosonde data: First results. *Journal of Space Weather and Space Climate*, 10. <https://doi.org/10.1051/swsc/2019042>
- Aravindan, P., & Iyer, K. N. (1990). Day-to-day variability in ionospheric electron content at low latitudes. *Planetary and Space Science*. [https://doi.org/10.1016/0032-0633\(90\)90033-M](https://doi.org/10.1016/0032-0633(90)90033-M)
- Astafieva, E. (2019). Ionospheric Detection of Natural Hazards. *Reviews of Geophysics*, 57(4), 1265–1288. <https://doi.org/10.1029/2019RG000668>
- Astafieva, E., & Heki, K. (2009). Dependence of waveform of near-field coseismic ionospheric disturbances on focal mechanisms. *Earth, Planets and Space*, 61(7), 939–943. <https://doi.org/10.1186/BF03353206>
- Astafieva, E., Heki, K., Kiryushkin, V., Afraimovich, E., & Shalimov, S. (2009). Two-mode long-distance propagation of coseismic ionosphere disturbances. *Journal of Geophysical Research: Space Physics*, 114(10). <https://doi.org/10.1029/2008JA013853>
- Astafieva, E., Maletckii, B., Mikesell, T. D., Munaibari, E., Ravanelli, M., Coisson, P., Manta, F., & Rolland, L. (2022). The 15 January 2022 Hunga Tonga Eruption History as Inferred From Ionospheric Observations. *Geophysical Research Letters*, 49(10). <https://doi.org/10.1029/2022GL098827>
- Bagiya, M. S., Sunil, A. S., Rolland, L., Nayak, S., Ponraj, M., Thomas, D., & Ramesh, D. S. (2019). Mapping the Impact of Non-Tectonic Forcing mechanisms on GNSS measured Coseismic Ionospheric Perturbations. *Scientific Reports*, 9(1). <https://doi.org/10.1038/s41598-019-54354-0>
- Balan, N., Bailey, G. J., Jayachandran, B., & Jayachandran, B. (1993). Ionospheric evidence for a nonlinear relationship between the solar e.u.v. and 10.7 cm fluxes during an intense solar cycle. *Planetary and Space Science*. [https://doi.org/10.1016/0032-0633\(93\)90043-2](https://doi.org/10.1016/0032-0633(93)90043-2)
- Balan, N., Bailey, G. J., & Su, Y. Z. (1996). Variations of the ionosphere and related solar fluxes during solar cycles 21 and 22. *Advances in Space Research*. [https://doi.org/10.1016/0273-1177\(95\)00828-3](https://doi.org/10.1016/0273-1177(95)00828-3)
- Balan, N., Balan, N., Yamamoto, M., Sreeja, V., Batista, I. S., Lynn, K. J. W., Abdu, M. A., Ravindran, S., Kikuchi, T., Otsuka, Y., Shokawa, K., & Alex, S. (2011). A statistical study of the response of the dayside equatorial F2 layer to the main phase of intense geomagnetic storms as an indicator of penetration electric field. *Journal of Geophysical Research*. <https://doi.org/10.1029/2010JA016001>
- Banks, P. M. (1977). Observations of joule and particle heating in the auroral zone. *Journal of Atmospheric and Terrestrial Physics*, 39(2), 179–193. [https://doi.org/10.1016/0021-9169\(77\)90112-X](https://doi.org/10.1016/0021-9169(77)90112-X)
- Barr, R. (1982). The nocturnal D-region as seen by VLF radio waves. *Journal of Atmospheric and Solar-Terrestrial Physics*. [https://doi.org/10.1016/0021-9169\(82\)90047-2](https://doi.org/10.1016/0021-9169(82)90047-2)
- Barta, V., Potuznikova, K., Knížová, P. K., Chum, J., Mielich, J., Guerra, M., Buresova, D., Sivakandan, M., Verhulst, T. G. W., Berényi, K. A., Altadill, D., Segarra, A., Kouba, D., Mosna, Z., Buzás, A., Spogli, L., Urbár, J., Paula, V. de, Ventriglia, V., & Belehaki, A. (2025). Multi-instrument analysis of medium-scale travelling ionospheric disturbances generated by an intense tropospheric jet-front system with severe convection in Europe in August 2023. *Journal of Space Weather and Space Climate*, 15, 31. <https://doi.org/10.1051/swsc/2025025>

- Bates, H. F. (1974). Atmospheric expansion from Joule heating. *Planetary and Space Science*. [https://doi.org/10.1016/0032-0633\(74\)90162-7](https://doi.org/10.1016/0032-0633(74)90162-7)
- Belehaki, A., Tzagouri, I., Altadill, D., Blanch, E., Borries, C., Buresova, D., Chum, J., Galkin, I., Juan, J. M., Segarra, A., Timote, C. C., Tziotziou, K., Verhulst, T. G. W., & Watermann, J. (2020). An overview of methodologies for real-time detection, characterisation and tracking of traveling ionospheric disturbances developed in the TechTIDE project. *Journal of Space Weather and Space Climate*, 10. <https://doi.org/10.1051/swsc/2020043>
- Beniguel, Y., & Adam, J.-P. (2007). Effects of Scintillations in GNSS Operation. https://doi.org/10.1007/1-4020-5446-7_18
- Bento, J., Saleiro, P., Cruz, A. F., Figueiredo, M. A. T., & Bizarro, P. (2021). TimeSHAP: Explaining Recurrent Models through Sequence Perturbations. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2565–2573. <https://doi.org/10.1145/3447548.3467166>
- Berta, E., Bach, F., & Jordan, M. (2023). Classifier Calibration with ROC-Regularized Isotonic Regression (No. arXiv:2311.12436). *arXiv*. <https://doi.org/10.48550/arXiv.2311.12436>
- Bondár, I., & Storchak, D. (2011). Improved location procedures at the International Seismological Centre. *Geophysical Journal International*, 186(3), 1220–1244. <https://doi.org/10.1111/j.1365-246X.2011.05107.x>
- Borries, C., Ferreira, A. A., Nykiel, G., & Borges, R. A. (2023). A new index for statistical analyses and prediction of travelling ionospheric disturbances. *Journal of Atmospheric and Solar-Terrestrial Physics*, 247, 106069. <https://doi.org/10.1016/j.jastp.2023.106069>
- Borries, C., Jakowski, N., & Wilken, V. (2009). Storm induced large scale TIDs observed in GPS derived TEC. *Annales Geophysicae*, 27(4), 1605–1612. <https://doi.org/10.5194/angeo-27-1605-2009>
- Boška, J., & Šauli, P. (2001). Observations of gravity waves of meteorological origin in the F-Region ionosphere. *Physics and Chemistry of The Earth Part C-Solar-Terrestrial and Planetary Science*. [https://doi.org/10.1016/S1464-1917\(01\)00024-1](https://doi.org/10.1016/S1464-1917(01)00024-1)
- Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Brissaud, Q., & Astafyeva, E. (2022). Near-real-time detection of co-seismic ionospheric disturbances using machine learning. *Geophysical Journal International*, 230(3), 2117–2130. <https://doi.org/10.1093/gji/ggac167>
- Bruyninx, C., Legrand, J., Fabian, A., & Pottiaux, E. (2019). GNSS metadata and data validation in the EUREF Permanent Network. *GPS Solutions*, 23(4), 106. <https://doi.org/10.1007/s10291-019-0880-9>
- Burt, S. (2022). Multiple airwaves crossing Britain and Ireland following the eruption of Hunga Tonga–Hunga Ha’apai on 15 January 2022. *Weather*, 77(3), 76–81. <https://doi.org/10.1002/WEA.4182>
- Cesaroni, C., Spogli, L., & De Franceschi, G. (2021). IONORING: Real-time monitoring of the total electron content over Italy. *Remote Sensing*, 13(16). <https://doi.org/10.3390/rs13163290>
- Chen, Y., Chen, Y., Wang, W., Burns, A. G., Liu, S., Gong, J., Yue, X., Jiang, G., & Coster, A. J. (2015). Ionospheric response to CIR-induced recurrent geomagnetic activity during the declining phase of solar cycle 23. *Journal of Geophysical Research*. <https://doi.org/10.1002/2014JA020657>
- 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>
- Chousianitis, K., Papanikolaou, X., Drakatos, G., & Tselentis, G. -Akis. (2021). NOANET: A Continuously Operating GNSS Network for Solid-Earth Sciences in Greece. *Seismological Research Letters*, 92(3), 2050–2064. <https://doi.org/10.1785/0220200340>
- Cicone, A. (2020). Iterative filtering as a direct method for the decomposition of nonstationary signals. *Numerical Algorithms*, 85(3), 811–827. <https://doi.org/10.1007/s11075-019-00838-z>
- Cicone, A., & Zhou, H. (2021). Numerical analysis for iterative filtering with new efficient implementations based on FFT. *Numerische Mathematik*, 147(1), 1–28. <https://doi.org/10.1007/s00211-020-01165-5>
- 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(2), 111–120. <https://doi.org/10.1007/s00190-006-0093-1>
- Cronin S, Brenna M, Smith I, Barker S, Tost M, Ford M, Tonga’onevai S, Kula T, Vaiomounga R. 2017. New volcanic island unveils explosive past. *Eos* 98. <https://doi.org/10.1029/2017eo076589>.
- Crowley, G., & Meier, R. R. (2013). Disturbed O/N₂ Ratios and their Transport to Middle and Low Latitudes. <https://doi.org/10.1029/181GM20>

- Danezis, C., Chatzinikos, M., & Kotsakis, C. (2019, April 1). LINEAR AND NONLINEAR DEFORMATION EFFECTS IN THE PERMANENT GNSS NETWORK OF CYPRUS (version 2019).
- Davies, K., & Baker, D. M. (1965). Ionospheric effects observed around the time of the Alaskan earthquake of March 28, 1964. *Journal of Geophysical Research* (1896-1977), 70(9), 2251–2253. <https://doi.org/10.1029/JZ070i009p02251>
- Davis, C. (1996). https://www.ukssdc.ac.uk/ionosondes/ionogram_interpretation.html
- Daubechies I, Lu J, Wu H-T. 2011. Synchrosqueezed wavelet transforms: an empirical mode decomposition-like tool. *Appl Comput Harmonic Anal* 30(2): 243–261. <https://doi.org/10.1016/j.acha.2010.08.002>.
- Dautermann, T., Calais, E., & Mattioli, G. S. (2009). Global Positioning System detection and energy estimation of the ionospheric wave caused by the 13 July 2003 explosion of the Soufrière Hills Volcano, Montserrat. *Journal of Geophysical Research: Solid Earth*, 114(B2). <https://doi.org/10.1029/2008JB005722>
- Deminov, M. G., Shubin, V. N., & Deminov, R. G. (2021). Dependence of the Annual Asymmetry in NmF2 on Geomagnetic Latitude and Solar Activity. *Geomagnetism and Aeronomy*. <https://doi.org/10.1134/S0016793221030038>
- Dempster, A., Schmidt, D. F., & Webb, G. I. (2023). QUANT: A Minimalist Interval Method for Time Series Classification (No. arXiv:2308.00928). arXiv. <https://doi.org/10.48550/arXiv.2308.00928>
- Ding, F., Wan, W., Liu, L., Afraimovich, E. I., Voeykov, S. v., & Perevalova, N. p. (2008). A statistical study of large-scale traveling ionospheric disturbances observed by GPS TEC during major magnetic storms over the years 2003–2005. *Journal of Geophysical Research: Space Physics*, 113(A3). <https://doi.org/10.1029/2008JA013037>
- Donnelly, R. F. (1976). Empirical models of solar flare X ray and EUV emission for use in studying their E and F region effects. *Journal of Geophysical Research*. <https://doi.org/10.1029/JA081I025P04745>
- Dorogush, A. V., Ershov, V., & Gulin, A. (2018). CatBoost: Gradient boosting with categorical features support (No. arXiv:1810.11363). arXiv. <https://doi.org/10.48550/arXiv.1810.11363>
- Dow, J. M., Neilan, R. E., & Rizos, C. (2009). The International GNSS Service in a changing landscape of Global Navigation Satellite Systems. *Journal of Geodesy*, 83(3), 191–198. <https://doi.org/10.1007/s00190-008-0300-3>
- Draper, N. R., & Smith, H. (2014). *Applied regression analysis* (3rd ed.). <https://doi.org/10.1002/9781118625590.ch15>
- Ducic, V., Artru, J., & Lognonné, P. (2003). Ionospheric remote sensing of the Denali Earthquake Rayleigh surface waves. *Geophysical Research Letters*, 30(18). <https://doi.org/10.1029/2003GL017812>
- Edelsbrunner, H., Kirkpatrick, D., & Seidel, R. (1983). On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29(4), 551–559. *IEEE Transactions on Information Theory*. <https://doi.org/10.1109/TIT.1983.1056714>
- Ezquer, R. G., Ezquer, R. G., Ezquer, R. G., Mosert, M., Corbella, R., Erazu, M., Radicella, S. M., Cabrera, M. A., Cabrera, M. A., & de la Zerda, L. (2004). Day-to-day variability of ionospheric characteristics in the American sector. *Advances in Space Research*. <https://doi.org/10.1016/J.ASR.2004.03.016>
- Feldstein, Y. I., Popov, V. A., Cumnock, J. A., Prigancová, A., Blomberg, L., Kozyra, J. U., Tsurutani, B. T., Gromova, L. I., & Levitin, A. E. (2006). Auroral electrojets and boundaries of plasma domains in the magnetosphere during magnetically disturbed intervals. *Annales Geophysicae*. <https://doi.org/10.5194/ANGEO-24-2243-2006>
- Ferreira, A. A., Borries, C., Xiong, C., Borges, R. A., Mielich, J., & Kouba, D. (2020). Identification of potential precursors for the occurrence of Large-Scale Traveling Ionospheric Disturbances in a case study during September 2017. *Journal of Space Weather and Space Climate*, 10, 32. <https://doi.org/10.1051/swsc/2020029>
- Fischler, M. A., & Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6), 381–395. <https://doi.org/10.1145/358669.358692>
- Friis-Christensen, E., Lühr, H., Knudsen, D., & Haagmans, R. (2008). Swarm – An Earth Observation Mission investigating Geospace. *Advances in Space Research*, 41(1), 210–216. <https://doi.org/10.1016/j.asr.2006.10.008>
- Fuller-Rowell, T. (2011). Storm-Time Response of the Thermosphere–Ionosphere System. https://doi.org/10.1007/978-94-007-0326-1_32
- Fuller-Rowell, T., Richmond, A. D., & Maruyama, N. (2013). Global Modeling of Storm-Time Thermospheric Dynamics and Electrodynamics. <https://doi.org/10.1029/181GM18>

- Fuso, F., Crocetti, L., Ravanelli, M., & Soja, B. (2024). Machine learning-based detection of TEC signatures related to earthquakes and tsunamis: The 2015 Illapel case study. *GPS Solutions*, 28(3), 106. <https://doi.org/10.1007/s10291-024-01649-z>
- Gail, W. B., Prag, A. B., Coco, D. S., & Coker, C. (1993). A statistical characterization of local mid-latitude total electron content. *Journal of Geophysical Research*. <https://doi.org/10.1029/92JA01597>
- Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, 63(1), 3–42. <https://doi.org/10.1007/s10994-006-6226-1>
- Ghobadi, H., Spogli, L., Alfonsi, L., Cesaroni, C., Cicone, A., Linty, N., Romano, V., & Cafaro, M. (2020). Disentangling ionospheric refraction and diffraction effects in GNSS raw phase through fast iterative filtering technique. *GPS Solutions*, 24(3), 85. <https://doi.org/10.1007/s10291-020-01001-1>
- Gonzalez, W. D., Joselyn, J. A., Kamide, Y., Kroehl, H. W., Rostoker, G., Tsurutani, B. T., & Vasyliunas, V. M. (1994). What is a geomagnetic storm? *Journal of Geophysical Research: Space Physics*, 99(A4), 5771–5792. <https://doi.org/10.1029/93JA02867>
- Guerra, M., Cesaroni, C., Fiorentino, N., Tosone, F., Ventriglia, V., Pica, E., Astafyeva, E., Maletckii, B., & Spogli, L. (2025). Automatic Detection of Large-Scale Traveling Ionospheric Disturbances Using GNSS Data and Image Processing Techniques. *Space Weather*, 23(9), e2025SW004423. <https://doi.org/10.1029/2025SW004423>
- Guerra, M., Cesaroni, C., Ravanelli, M., & Spogli, L. (2024). Travelling ionospheric disturbances detection: A statistical study of detrending techniques, induced period error and near real-time observables. *Journal of Space Weather and Space Climate*, 14, 17. <https://doi.org/10.1051/swsc/2024017>
- Guo, G., Wang, H., Bell, D., Bi, Y., & Greer, K. (2003). KNN Model-Based Approach in Classification. In R. Meersman, Z. Tari, & D. C. Schmidt (Eds), *On The Move to Meaningful Internet Systems 2003: CoopIS, DOA, and ODBASE* (Vol. 2888, pp. 986–996). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-39964-3_62
- Guryanov, A. (2019). Histogram-Based Algorithm for Building Gradient Boosting Ensembles of Piecewise Linear Decision Trees (pp. 39–50). https://doi.org/10.1007/978-3-030-37334-4_4
- Haaser, R. A., Lay, E. H., & Junor, W. (2017). Analysis framework for systematically studying ionospheric response to impulsive events from below. *Radio Science*, 52(9), 1149–1169. <https://doi.org/10.1002/2016RS006196>
- Habarulema, J. B., Katamzi, Z. T., & Yizengaw, E. (2015). First observations of poleward large-scale traveling ionospheric disturbances over the African sector during geomagnetic storm conditions. *Journal of Geophysical Research: Space Physics*, 120(8), 6914–6929. <https://doi.org/10.1002/2015JA021066>
- Habarulema, J. B., Katamzi, Z. T., Yizengaw, E., Yamazaki, Y., & Seemala, G. (2016). Simultaneous storm time equatorward and poleward large-scale TIDs on a global scale. *Geophysical Research Letters*, 43(13), 6678–6686. <https://doi.org/10.1002/2016GL069740>
- Hall, C., Johnsen, M. G., & Smyth, A. R. (2020). On the relation between ionospheric parameters and sunspot number. <https://doi.org/10.1017/EXP.2020.19>
- Haralambous, H., Guerra, M., Chum, J., Verhulst, T. G. W., Barta, V., Altadill, D., Cesaroni, C., Galkin, I., Márta, K., Mielich, J., Kouba, D., Buresova, D., Segarra, A., Spogli, L., Rusz, J., & Zedník, J. (2023). Multi-Instrument Observations of Various Ionospheric Disturbances Caused by the 6 February 2023 Turkey Earthquake. *Journal of Geophysical Research: Space Physics*, 128(12), e2023JA031691. <https://doi.org/10.1029/2023JA031691>
- Hastie, T., Rosset, S., Zhu, J., & Zou, H. (2009). Multi-class AdaBoost. *Statistics and Its Interface*, 2(3), 349–360. <https://doi.org/10.4310/SII.2009.v2.n3.a8>
- He, M., Liu, L., Wan, W., & Zhao, B. (2011). A study on the nighttime midlatitude ionospheric trough. *Journal of Geophysical Research*. <https://doi.org/10.1029/2010JA016252>
- Heki, K. (2006). Explosion energy of the 2004 eruption of the Asama Volcano, central Japan, inferred from ionospheric disturbances. *Geophysical Research Letters*, 33(14). <https://doi.org/10.1029/2006GL026249>
- Heki, K. (2021). Ionospheric Disturbances Related to Earthquakes. In *Ionosphere Dynamics and Applications* (pp. 511–526). American Geophysical Union (AGU). <https://doi.org/10.1002/9781119815617.ch21>
- Heki, K., & Ping, J. (2005). Directivity and apparent velocity of the coseismic ionospheric disturbances observed with a dense GPS array. *Earth and Planetary Science Letters*, 236(3), 845–855. <https://doi.org/10.1016/j.epsl.2005.06.010>

- Hernández-Pajares, M., Juan, J. M., & Sanz, J. (2006). Medium-scale traveling ionospheric disturbances affecting GPS measurements: Spatial and temporal analysis. *Journal of Geophysical Research: Space Physics*, 111(7). <https://doi.org/10.1029/2005JA011474>
- Hernández-Pajares, M., Juan, J. M., Sanz, J., & Aragón-Ángel, A. (2012). Propagation of medium scale traveling ionospheric disturbances at different latitudes and solar cycle conditions. *Radio Science*, 47(4). <https://doi.org/10.1029/2011RS004951>
- Hickey, M. P. (2011). Atmospheric Gravity Waves and Effects in the Upper Atmosphere Associated with Tsunamis. <https://doi.org/10.5772/13577>
- Hocke, K., & Schlegel, K. (1996). A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. *Annales Geophysicae*, 14(9), 917–940. <https://doi.org/10.1007/s00585-996-0917-6>
- Igarashi, K., Kainuma, S., Nishimuta, I., Okamoto, S., Kuroiwa, H., Tanaka, T., & Ogawa, T. (1994). Ionospheric and atmospheric disturbances around Japan caused by the eruption of Mount Pinatubo on 15 June 1991. *Journal of Atmospheric and Terrestrial Physics*, 56(9), 1227–1234. [https://doi.org/10.1016/0021-9169\(94\)90060-4](https://doi.org/10.1016/0021-9169(94)90060-4)
- Ikotun, A. M., Ezugwu, A. E., Abualigah, L., Abuhaija, B., & Heming, J. (2023). K-means clustering algorithms: A comprehensive review, variants analysis, and advances in the era of big data. *Information Sciences*, 622, 178–210. <https://doi.org/10.1016/j.ins.2022.11.139>
- Iyemori, T. (1990). Storm-Time Magnetospheric Currents Inferred from Mid-Latitude Geomagnetic Field Variations. *Journal of Geomagnetism and Geoelectricity*, 42(11), 1249–1265. <https://doi.org/10.5636/jgg.42.1249>
- Jakowski, N., Wilken, V., Tsybulya, K., & Heise, S. (2006). Earthquake Signatures in the Ionosphere Deduced from Ground and Space Based GPS Measurements. https://doi.org/10.1007/3-540-29522-4_4
- Janssen, V. (2012). Likely Impact of the Approaching Solar Maximum on GNSS Surveys: Be Alert but Not Alarmed.
- Jewtoukoff, V., Plougonven, R., & Hertzog, A. (2013). Gravity waves generated by deep tropical convection: Estimates from balloon observations and mesoscale simulations. *Journal of Geophysical Research*. <https://doi.org/10.1002/JGRD.50781>
- Jin, S., Jin, S., Jin, R., & Liu, X. (2019). Tsunami Ionospheric Disturbances. https://doi.org/10.1007/978-981-10-3178-6_12
- Jin, Y., Kotova, D., Xiong, C., Brask, S. M., Clausen, L. B. N., Kervalishvili, G., Stolle, C., & Miloch, W. J. (2022). Ionospheric Plasma Irregularities—IPIR - Data Product Based on Data From the Swarm Satellites. *Journal of Geophysical Research: Space Physics*, 127(4), e2021JA030183. <https://doi.org/10.1029/2021JA030183>
- Johansson, U., & Gabrielsson, P. (2019). Are Traditional Neural Networks Well-Calibrated? 2019 International Joint Conference on Neural Networks (IJCNN), 1–8. <https://doi.org/10.1109/IJCNN.2019.8851962>
- Katamzi, Z. T., Smith, N. D., Mitchell, C. N., Spalla, P., & Materassi, M. (2012). Statistical analysis of travelling ionospheric disturbances using TEC observations from geostationary satellites. *Journal of Atmospheric and Solar-Terrestrial Physics*, 74, 64–80. <https://doi.org/10.1016/j.jastp.2011.10.006>
- Kil, H., & Paxton, L. J. (2017). Global Distribution of Nighttime Medium-Scale Traveling Ionospheric Disturbances Seen by Swarm Satellites. *Geophysical Research Letters*, 44(18), 9176–9182. <https://doi.org/10.1002/2017GL074750>
- King, G. A. M. (1956). Electron distribution in the ionosphere. *Journal of Atmospheric and Solar-Terrestrial Physics*. [https://doi.org/10.1016/0021-9169\(56\)90188-X](https://doi.org/10.1016/0021-9169(56)90188-X)
- Komjathy, A., Yang, Y. M., Meng, X., Verkhoglyadova, O., Mannucci, A. J., & Langley, R. B. (2016). Review and perspectives: Understanding natural-hazards-generated ionospheric perturbations using GPS measurements and coupled modeling. *Radio Science*, 51(7), 951–961. <https://doi.org/10.1002/2015RS005910>
- Kouris, S. S., & Fotiadis, D. N. (2002). Ionospheric variability: a comparative statistical study. *Advances in Space Research*. [https://doi.org/10.1016/S0273-1177\(02\)00045-5](https://doi.org/10.1016/S0273-1177(02)00045-5)
- Krzysztofowicz, R. 2024. Probabilistic forecasts and optimal decisions. John Wiley & Sons Ltd, UK.
- Kubota, T., Saito, T., & Nishida, K. (2022). Global fast-traveling tsunamis driven by atmospheric Lamb waves on the 2022 Tonga eruption. *Science*, 377(6601), 91–94. <https://doi.org/10.1126/science.abo4364>
- Kulichkov, S., Chunchuzov, I. P., Popov, O. E., Gorchakov, G. I., Mishenin, A. A., Perepelkin, V. G., Bush, G. A., Skorokhod, A., Vinogradov, Y. A., Semutnikova, E., Šepić, J., Medvedev, I. P., Gushchin, R. A., Kopeikin, V. M., Belikov, I. B., Gubanova, D. P., Karpov, A. V., & Tikhonov, A. V. (2022). Acoustic-

- Gravity Lamb Waves from the Eruption of the Hunga-Tonga-Hunga-Hapai Volcano, Its Energy Release and Impact on Aerosol Concentrations and Tsunami. *Pure and Applied Geophysics*. <https://doi.org/10.1007/s00024-022-03046-4>
- Lee, W. K., Kil, H., Kwak, Y.-S., Wu, Q., Cho, S., & Park, J.-U. (2011). The winter anomaly in the middle-latitude F region during the solar minimum period observed by the Constellation Observing System for Meteorology, Ionosphere, and Climate. *Journal of Geophysical Research*. <https://doi.org/10.1029/2010JA015815>
- Leonard, R. S., & Barnes Jr., R. A. (1965). Observation of ionospheric disturbances following the Alaska earthquake. *Journal of Geophysical Research* (1896-1977), 70(5), 1250–1253. <https://doi.org/10.1029/JZ070i005p01250>
- Lloyd, S. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2), 129–137. <https://doi.org/10.1109/TIT.1982.1056489>
- Lognonné, P., Artru, J., Garcia, R., Crespon, F., Ducic, V., Jeansou, E., et al. (2006a). Ground-based GPS imaging of ionospheric post-seismic signal. *Planetary and Space Science*, 54(5), 528–540. <https://doi.org/10.1016/j.pss.2005.10.021>
- Lognonné, P., Garcia, R., Crespon, F., Occhipinti, G., Kherani, A., & Artru-Lambin, J. (2006b). Seismic waves in the ionosphere. *Europhysics News*, 37(4), 11–14. <https://doi.org/10.1051/eprn:2006401>
- Li, M., Yang, Z., Song, J., Zhang, Y., Jian, X., & Shan, X. (2023). Statistical seismo-ionospheric influence with the focal mechanism under consideration. *Atmosphere*, 14(3), 455. <https://doi.org/10.3390/atmos14030455>
- Li, S., Wang, X., Tao, T., Zhu, Y., Qu, X., Li, Z., et al. (2023). Source model of the 2023 Turkey earthquake sequence imaged by Sentinel-1 and GPS measurements: Implications for heterogeneous fault behavior along the East Anatolian Fault zone. *Remote Sensing*, 15(10), 2618. <https://doi.org/10.3390/rs15102618>
- Liu, S., & Li, L. (2002). Study on Relationship between Southward IMF Events and Geomagnetic Storms. *Chinese Journal of Geophysics*. <https://doi.org/10.1002/CJG2.243>
- Liu, L., Zhao, B., Wan, W., Ning, B., Zhang, M.-L., & He, M. (2009). Seasonal variations of the ionospheric electron densities retrieved from Constellation Observing System for Meteorology, Ionosphere, and Climate mission radio occultation measurements. *Journal of Geophysical Research*. <https://doi.org/10.1029/2008JA013819>
- Lu, G., Goncharenko, L. P., Nicolls, M. J., Maute, A., Coster, A. J., & Paxton, L. J. (2012). Ionospheric and thermospheric variations associated with prompt penetration electric fields. *Journal of Geophysical Research*. <https://doi.org/10.1029/2012JA017769>
- Luhrmann, F., Park, J., Wong, W.-K., Corcoran, F., & Lewis, C. (2022). Detecting Traveling Ionospheric Disturbances with LSTM Based Anomaly Detection. 3002–3011. <https://doi.org/10.33012/2022.18343>
- Lundberg, S., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions (No. arXiv:1705.07874). arXiv. <https://doi.org/10.48550/arXiv.1705.07874>
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics: Vol. 5.1* (pp. 281–298). University of California Press. <https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fifth-Berkeley-Symposium-on-Mathematical-Statistics-and/chapter/Some-methods-for-classification-and-analysis-of-multivariate-observations/bsmsp/1200512992>
- Maletckii, B., & Astafyeva, E. (2021). Determining spatio-temporal characteristics of coseismic travelling ionospheric disturbances (CTID) in near real-time. *Scientific Reports*, 11(1). <https://doi.org/10.1038/s41598-021-99906-5>
- Maletckii, B., & Astafyeva, E. (2022). Near-Real-Time Analysis of the Ionospheric Response to the 15 January 2022 Hunga Tonga-Hunga Ha’apai Volcanic Eruption. *Journal of Geophysical Research: Space Physics*, 127(10), e2022JA030735. <https://doi.org/10.1029/2022JA030735>
- Maletckii, B., Astafyeva, E., Sanchez, S. A., Kherani, E. A., & de Paula, E. R. (2023). The 6 February 2023 Türkiye Earthquake Sequence as Detected in the Ionosphere. *Journal of Geophysical Research: Space Physics*, 128(9), e2023JA031663. <https://doi.org/10.1029/2023JA031663>
- Maletckii, B. (2023). Near-real-time analysis of traveling ionospheric disturbances for the fundamental science, natural hazards and space weather applications. *Astrophysics*. Université Paris Cité, 2023. English. NNT : 2023UNIP7194. tel-04749877

- Mannucci, A. J., Wilson, B. D., Yuan, D. N., Ho, C. H., Lindqwister, U. J., & Runge, T. F. (1998). A global mapping technique for GPS-derived ionospheric total electron content measurements. *Radio Science*, 33(3), 565–582. <https://doi.org/10.1029/97RS02707>
- Mansilla, G. A., Mansilla, G. A., Mosert, M., Ezquer, R. G., & Ezquer, R. G. (2005). Seasonal variation of the total electron content, maximum electron density and equivalent slab thickness at a South-American station. *Journal of Atmospheric and Solar-Terrestrial Physics*. <https://doi.org/10.1016/J.JASTP.2005.02.024>
- Martire, L., Krishnamoorthy, S., Vergados, P., Romans, L. J., Szilágyi, B., Meng, X., Anderson, J. L., Komjáthy, A., & Bar-Sever, Y. E. (2022). The GUARDIAN system-a GNSS upper atmospheric real-time disaster information and alert network. *GPS Solutions*, 27(1), 32. <https://doi.org/10.1007/s10291-022-01365-6>
- Maruyama, T., Yusupov, K., & Akchurin, A. (2016a). Interpretation of deformed ionograms induced by vertical ground motion of seismic Rayleigh waves and infrasound in the thermosphere. *Annales Geophysicae*, 34(2), 271–278. <https://doi.org/10.5194/angeo-34-271-2016>
- Maruyama, T., Yusupov, K., & Akchurin, A. (2016b). Ionosonde tracking of infrasound wavefronts in the thermosphere launched by seismic waves after the 2010 M8.8 Chile earthquake. *Journal of Geophysical Research: Space Physics*, 121(3), 2683–2692. <https://doi.org/10.1002/2015JA022260>
- Mathews, J. D., Sulzer, M. P., & Perillat, P. (1997). Aspects of layer electrodynamics inferred from high-resolution ISR observations of the 80–270 km ionosphere. *Geophysical Research Letters*. <https://doi.org/10.1029/97GL01320>
- Matoza, R. S., Fee, D., Assink, J. D., Iezzi, A. M., Green, D. N., Kim, K., Toney, L., Lecocq, T., Krishnamoorthy, S., Lalande, J.-M., Nishida, K., Gee, K. L., Haney, M. M., Ortiz, H. D., Brissaud, Q., Martire, L., Rolland, L., Vergados, P., Nippess, A., ... Wilson, D. C. (2022). Atmospheric waves and global seismoacoustic observations of the January 2022 Hunga eruption, Tonga. *Science*, 377(6601), 95–100. <https://doi.org/10.1126/science.abo7063>
- Meng, X., Vergados, P., Komjathy, A., & Verkhoglyadova, O. (2019). Upper Atmospheric Responses to Surface Disturbances: An Observational Perspective. *Radio Science*, 54(11), 1076–1098. <https://doi.org/10.1029/2019RS006858>
- Miyoshi, Y., & Fujiwara, H. (2008). Gravity Waves in the Thermosphere Simulated by a General Circulation Model. *Journal of Geophysical Research*. <https://doi.org/10.1029/2007JD008874>
- Motoba, T., Kikuchi, T., Okuzawa, T., & Yumoto, K. (2003). Dynamical response of the magnetosphere-ionosphere system to a solar wind dynamic pressure oscillation. *Journal of Geophysical Research*. <https://doi.org/10.1029/2002JA009696>
- Munro, G. H. (1962). Diurnal Variations in the Ionosphere Deduced from Satellite Radio Signals. *Journal of Geophysical Research*. <https://doi.org/10.1029/JZ067I001P00147>
- National Geodetic Survey. (2025). NOAA continuously operating reference stations (CORS) network. Retrieved from <https://www.fisheries.noaa.gov/inport/item/39662>
- 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>
- Newell, P. T., Sotirelis, T., Liou, K., Meng, C. I., & Rich, F. J. (2007). A nearly universal solar wind-magnetosphere coupling function inferred from 10 magnetospheric state variables. *Journal of Geophysical Research: Space Physics*, 112(A1). <https://doi.org/10.1029/2006JA012015>
- Nicholls, R. W. (1993). The Upper Atmosphere and the Ionosphere. https://doi.org/10.1007/978-94-011-2048-7_5
- Oyama, S., & Watkins, B. J. (2012). Generation of Atmospheric Gravity Waves in the Polar Thermosphere in Response to Auroral Activity. *Space Science Reviews*, 168(1), 463–473. <https://doi.org/10.1007/s11214-011-9847-z>
- Paul, A. K. (1995). Short term variability of the ionosphere. *Advances in Space Research*. [https://doi.org/10.1016/S0273-1177\(99\)80020-9](https://doi.org/10.1016/S0273-1177(99)80020-9)
- Paul, A., & DasGupta, A. (2010). Characteristics of the equatorial ionization anomaly in relation to the day-to-day variability of ionospheric irregularities around the postsunset period. *Radio Science*. <https://doi.org/10.1029/2009RS004329>
- Pavlov, A., & Pavlova, N. M. (2017). Formation mechanisms of the midlatitudinal NmF2 semiannual anomaly under daytime quiet geomagnetic conditions at low solar activity. *Geomagnetism and Aeronomy*. <https://doi.org/10.1134/S0016793217040132>

- Pignalberi, A., Pietrella, M., Pezzopane, M., Nava, B., & Cesaroni, C. (2022). The Ionospheric Equivalent Slab Thickness: A Review Supported by a Global Climatological Study Over Two Solar Cycles. *Space Science Reviews*, 218(4), 37. <https://doi.org/10.1007/s11214-022-00909-z>
- Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2019). CatBoost: Unbiased boosting with categorical features (No. arXiv:1706.09516). arXiv. <https://doi.org/10.48550/arXiv.1706.09516>
- Ravanelli, M., Constantinou, V., Liu, H., & Bortnik, J. (2024). Exploring AI Progress in GNSS Remote Sensing: A Deep Learning Based Framework for Real-Time Detection of Earthquake and Tsunami Induced Ionospheric Perturbations. *Radio Science*, 59(9), e2024RS008016. <https://doi.org/10.1029/2024RS008016>
- Ravanelli, M., Occhipinti, G., Savastano, G., Komjathy, A., Shume, E. B., & Crespi, M. (2021). GNSS total variometric approach: First demonstration of a tool for real-time tsunami genesis estimation. *Scientific Reports*, 11(1). <https://doi.org/10.1038/s41598-021-82532-6>
- RING WORKING GROUP. (2016). RETE INTEGRATA NAZIONALE GNSS. 10.13127/RING.
- Rishbeth, H., Müller-Wodarg, I. C. F., Müller-Wodarg, I. C. F., Zou, L., Fuller-Rowell, T., Millward, G. H., Millward, G. H., Millward, G. H., Moffett, R. J., Idenden, D. W., & Aylward, A. D. (2000). Annual and semiannual variations in the ionospheric F2-layer: II. Physical discussion. *Annales Geophysicae*. <https://doi.org/10.1007/S00585-000-0945-6>
- Roberts, D. H., Klobuchar, J. A., Fougere, P. F., & Hendrickson, D. H. (1982). A large-amplitude traveling ionospheric disturbance produced by the May 18, 1980, explosion of Mount St. Helens. *Journal of Geophysical Research: Space Physics*, 87(A8), 6291–6301. <https://doi.org/10.1029/JA087iA08p06291>
- Roberts, S. J., Husmeier, D., Rezek, I., & Penny, W. (1998). Bayesian approaches to Gaussian mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1133–1142. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. <https://doi.org/10.1109/34.730550>
- Rodger, A. S., Wells, G. D., Wells, G. D., Moffett, R. J., & Bailey, G. J. (2001). The variability of Joule heating, and its effects on the ionosphere and thermosphere. *Annales Geophysicae*. <https://doi.org/10.5194/ANGE0-19-773-2001>
- Rolland, L., Lognonné, P., & Mune Kane, H. M. (2011). Detection and modeling of Rayleigh wave induced patterns in the ionosphere. *Journal of Geophysical Research Space Physics*, 116(A5), 5320. <https://doi.org/10.1029/2010JA016060i>
- Rolland, L. M., Lognonné, P., Astafyeva, E., Kherani, E. A., Kobayashi, N., Mann, M., & Mune Kane, H. (2011). The resonant response of the ionosphere imaged after the 2011 off the Pacific coast of Tohoku Earthquake. *Earth, Planets and Space*, 63(7), 853–857. <https://doi.org/10.5047/eps.2011.06.020>
- Rolland, L. M., Vergnolle, M., Nocquet, J. M., Sladen, A., Dessa, J. X., Tavakoli, F., Nankali, H. R., & Cappa, F. (2013). Discriminating the tectonic and non-tectonic contributions in the ionospheric signature of the 2011, Mw7.1, dip-slip Van earthquake, Eastern Turkey. *Geophysical Research Letters*, 40(11), 2518–2522. <https://doi.org/10.1002/grl.50544>
- Sagiya, T. (2004). A decade of GEONET: 1994–2003 —The continuous GPS observation in Japan and its impact on earthquake studies—. *Earth, Planets and Space*, 56(8), xxix–xli. <https://doi.org/10.1186/BF03353077>
- Saito, S. (2022). Ionospheric disturbances observed over Japan following the eruption of Hunga Tonga-Hunga Ha’apai on 15 January 2022. *Earth, Planets and Space*, 74(1), 57. <https://doi.org/10.1186/s40623-022-01619-0>
- Savastano, G., Komjathy, A., Verkhoglyadova, O., Mazzoni, A., Crespi, M., Wei, Y., & Mannucci, A. J. (2017). Real-time detection of tsunami ionospheric disturbances with a stand-alone GNSS receiver: A preliminary feasibility demonstration. *Scientific Reports*, 7. <https://doi.org/10.1038/srep46607>
- Scali, J. L. (1992). The Mid-Latitude Trough. A Review.
- Schafer, R. (2011). What Is a Savitzky-Golay Filter? [Lecture Notes]. *IEEE Signal Processing Magazine - IEEE SIGNAL PROCESS MAG*, 28, 111–117. <https://doi.org/10.1109/MSP.2011.941097>
- Segarra, A., Altadill, D., Paula, V. de, & Navas-Portella, V. (2025). Climatology of large-scale Traveling Ionospheric Disturbances above Europe during the 2014–2023 period. *Journal of Space Weather and Space Climate*, 15, 28. <https://doi.org/10.1051/swsc/2025024>
- Shalimov, S. L., Rozhnoi, A., Solov’eva, M. S., & Ol’shanskaya, E. V. (2019). Impact of Earthquakes and Tsunamis on the Ionosphere. *Izvestiya-Physics of The Solid Earth*. <https://doi.org/10.1134/S1069351319010087>

- Shiokawa, K. (2023). Introduction of Space Weather Research on Magnetosphere and Ionosphere of the Earth. In: Kusano, K. (eds) *Solar-Terrestrial Environmental Prediction*. Springer, Singapore. https://doi.org/10.1007/978-981-19-7765-7_4
- Silva Filho, T., Song, H., Perello-Nieto, M., Santos-Rodriguez, R., Kull, M., & Flach, P. (2023). Classifier calibration: A survey on how to assess and improve predicted class probabilities. *Machine Learning*, 112(9), 3211–3260. <https://doi.org/10.1007/s10994-023-06336-7>
- Smith, P. A., & King, J. W. (1981). Long-term relationships between sunspots, solar faculae and the ionosphere. *Journal of Atmospheric and Solar-Terrestrial Physics*. [https://doi.org/10.1016/0021-9169\(81\)90020-9](https://doi.org/10.1016/0021-9169(81)90020-9)
- Soicher, H., & Gorman, F. J. (1985). Seasonal and day-to-day variability of total electron content at mid-latitudes near solar maximum. *Radio Science*. <https://doi.org/10.1029/RS020i003P00383>
- Spogli, L., Alberti, T., Bagiacchi, P., Cafarella, L., Cesaroni, C., Cianchini, G., Coco, I., Mauro, D. D., Ghidoni, R., Giannattasio, F., Ippolito, A., Marocci, C., Pezzopane, M., Pica, E., Pignalberi, A., Perrone, L., Romano, V., Sabbagh, D., Scotto, C., ... Viola, M. (2024). The effects of the May 2024 Mother's Day superstorm over the Mediterranean sector: From data to public communication. *Annals of Geophysics*, 67(2), PA218–PA218. <https://doi.org/10.4401/ag-9117>
- Spogli, L., Ghobadi, H., Cicone, A., Alfonsi, L., Cesaroni, C., Linty, N., Romano, V., & Cafaro, M. (2021). Adaptive Phase Detrending for GNSS Scintillation Detection: A Case Study Over Antarctica. *IEEE Geoscience and Remote Sensing Letters*, PP, 1–5. <https://doi.org/10.1109/LGRS.2021.3067727>
- Spogli, L., Piersanti, M., Cesaroni, C., Materassi, M., Cicone, A., Alfonsi, L., Romano, V., & Ezquer, R. G. (2019). Role of the external drivers in the occurrence of low-latitude ionospheric scintillation revealed by multi-scale analysis. *Journal of Space Weather and Space Climate*, 9, A35. <https://doi.org/10.1051/swsc/2019032>
- Stallone, A., Cicone, A., & Materassi, M. (2020). New insights and best practices for the successful use of Empirical Mode Decomposition, Iterative Filtering and derived algorithms. *Scientific Reports*, 10(1), Article 1. <https://doi.org/10.1038/s41598-020-72193-2>
- Stolle, C., Manoj, C., Manoj, C., Lühr, H., Maus, S., & Alken, P. (2008). Estimating the daytime Equatorial Ionization Anomaly strength from electric field proxies. *Journal of Geophysical Research*. <https://doi.org/10.1029/2007JA012781>
- Tanskanen, E. I. (2009). A comprehensive high-throughput analysis of substorms observed by IMAGE magnetometer network: Years 1993–2003 examined. *Journal of Geophysical Research: Space Physics*, 114(A5). <https://doi.org/10.1029/2008JA013682>
- Thaganyana, G. P., Habarulema, J. B., Ngwira, C., & Azeem, I. (2022). Equatorward Medium to Large-Scale Traveling Ionospheric Disturbances of High Latitude Origin During Quiet Conditions. *Journal of Geophysical Research: Space Physics*, 127(3), e2021JA029558. <https://doi.org/10.1029/2021JA029558>
- Themens, D. R., Watson, C., Žagar, N., Vasylykevych, S., Elvidge, S., McCaffrey, A., Prikryl, P., Reid, B., Wood, A., & Jayachandran, P. T. (2022). Global Propagation of Ionospheric Disturbances Associated With the 2022 Tonga Volcanic Eruption. *Geophysical Research Letters*, 49(7), e2022GL098158. <https://doi.org/10.1029/2022GL098158>
- Tiwari, R., Strangeways, H. J., Tiwari, S., & Ahmed, A. (2013). Investigation of ionospheric irregularities and scintillation using TEC at high latitude. *Advances in Space Research*. <https://doi.org/10.1016/J.ASR.2013.06.010>
- Tornatore, V., Cesaroni, C., Pezzopane, M., Alizadeh, M. M., & Schuh, H. (2021). Performance evaluation of VTEC GIMs for regional applications during different solar activity periods, using RING TEC values. *Remote Sensing*, 13(8). <https://doi.org/10.3390/rs13081470>
- Torr, D. G., Torr, M. R., & Richards, P. G. (1980). Causes of the F region winter anomaly. *Geophysical Research Letters*. <https://doi.org/10.1029/GL007i005P00301>
- Torres, J. A., Altamimi, Z., Boucher, C., Brockmann, E., Bruyninx, C., Caporali, A., Gurtner, W., Habrich, H., Hornik, H., Ihde, J., Kenyeres, A., Mäkinen, J., Marel, H. v. d., Seeger, H., Simek, J., Stangl, G., & Weber, G. (2009). Status of the European Reference Frame (EUREF). In M. G. Sideris (Ed.), *Observing our Changing Earth* (pp. 47–56). Springer. https://doi.org/10.1007/978-3-540-85426-5_6
- Tsugawa, T., Saito, A., & Otsuka, Y. (2004). A statistical study of large-scale traveling ionospheric disturbances using the GPS network in Japan. *Journal of Geophysical Research: Space Physics*, 109(A6). <https://doi.org/10.1029/2003JA010302>
- Tsurutani, B. T., Tsurutani, B. T., Verkhoglyadova, O. P., Verkhoglyadova, O. P., Mannucci, A. J., Saito, A., Araki, T., Yumoto, K., Tsuda, T., Abdu, M. A., Sobral, J. H. A., Gonzalez, W. D., McCreadie, H., Lakhina,

- G. S., & Vasyliūnas, V. M. (2008). Prompt penetration electric fields (PPEFs) and their ionospheric effects during the great magnetic storm of 30-31 October 2003. *Journal of Geophysical Research*. <https://doi.org/10.1029/2007JA012879>
- Urbar, J., Spogli, L., Cicone, A., Clausen, L. B. N., Jin, Y., Wood, A. G., Alfonsi, L., Cesaroni, C., Kotova, D., Høeg, P., & Miloch, W. J. (2022). Multi-scale response of the high-latitude topside ionosphere to geospace forcing. *Advances in Space Research*. <https://doi.org/10.1016/j.asr.2022.06.045>
- US Geological Survey. (2023). Event page of the M7.8 - Pazarcik earthquake, Kahramanmaraş earthquake sequence. Retrieved from <https://earthquake.usgs.gov/earthquakes/eventpage/us6000jllz/executive>
- Vaishnav, R., Jacobi, C., & Berdermann, J. (2019). Long-term trends in the ionospheric response to solar extreme-ultraviolet variations. *Annales Geophysicae*. <https://doi.org/10.5194/ANGE0-37-1141-2019>
- Ventriglia, V., Guerra, M., Cesaroni, C., Spogli, L., Altadill, D., Segarra, A., Galkin, I., Barta, V., Verhulst, T. G. W., Paula, V. de, Navas-Portella, V., Berényi, K. A., & Belehaki, A. (2025). An explainable Machine Learning model for Large-Scale Travelling Ionospheric Disturbances forecasting. *Journal of Space Weather and Space Climate*, 15, 25. <https://doi.org/10.1051/swsc/2025020>
- Verhulst, T. G. W., Altadill, D., Barta, V., Belehaki, A., Burešová, D., Cesaroni, C., Galkin, I., Guerra, M., Ippolito, A., Herekakis, T., Kouba, D., Mielich, J., Segarra, A., Spogli, L., & Tsaouri, I. (2022). Multi-instrument detection in Europe of ionospheric disturbances caused by the 15 January 2022 eruption of the Hunga volcano. *Journal of Space Weather and Space Climate*, 12, 35. <https://doi.org/10.1051/swsc/2022032>
- Volland, H. (1974). Coupling Between the Neutral and the Ionized Atmosphere. https://doi.org/10.1007/978-94-010-2265-1_4
- Vovk, V., & Petej, I. (2014). Venn-Abers predictors (No. arXiv:1211.0025). arXiv. <https://doi.org/10.48550/arXiv.1211.0025>
- Wang, W., Wang, W., Lei, J., Burns, A. G., Solomon, S. C., Wiltberger, M., Xu, J., Zhang, Y., Paxton, L. J., & Coster, A. J. (2010). Ionospheric response to the initial phase of geomagnetic storms: Common features. *Journal of Geophysical Research*. <https://doi.org/10.1029/2009JA014461>
- Watanabe, S. (2025). Tree-Structured Parzen Estimator: Understanding Its Algorithm Components and Their Roles for Better Empirical Performance (No. arXiv:2304.11127). arXiv. <https://doi.org/10.48550/arXiv.2304.11127>
- Wei, C., Bühler, O., & Tabak, E. G. (2015). Evolution of Tsunami-Induced Internal Acoustic-Gravity Waves. *Journal of the Atmospheric Sciences*. <https://doi.org/10.1175/JAS-D-14-0179.1>
- Wilkinson, P., Wu, J., Du, J., & Wang, Y. J. (2001). Real-time total electron content estimates using the international reference ionosphere. *Advances in Space Research*. [https://doi.org/10.1016/S0273-1177\(00\)00147-2](https://doi.org/10.1016/S0273-1177(00)00147-2)
- Willis, D. M. (1964). The sudden commencement and first phase of a geomagnetic storm. *Journal of Atmospheric and Solar-Terrestrial Physics*. [https://doi.org/10.1016/0021-9169\(64\)90189-8](https://doi.org/10.1016/0021-9169(64)90189-8)
- Wright, C., Hindley, N., Alexander, M., Barlow, M., Hoffmann, L., Mitchell, C., Prata, F., Bouillon, M., Carstens, J., Clerbaux, C., Osprey, S., Powell, N., Randall, C., & Yue, J. (2022). Surface-to-space atmospheric waves from Hunga Tonga-Hunga Ha'apai eruption. *Nature*, 609. <https://doi.org/10.1038/s41586-022-05012-5>
- Yamazaki, Y., Matzka, J., Stolle, C., Kervalishvili, G., Rauberg, J., Bronkalla, O., Morschhauser, A., Bruinsma, S., Shprits, Y. Y., & Jackson, D. R. (2022). Geomagnetic Activity Index Hpo. *Geophysical Research Letters*, 49(10), e2022GL098860. <https://doi.org/10.1029/2022GL098860>
- Yang, H., Moreno, E. M., & Hernández-Pajares, M. (2019). ADDTID: An alternative tool for studying earthquake/tsunami signatures in the ionosphere. Case of the 2011 Tohoku earthquake. *Remote Sensing*, 11(16). <https://doi.org/10.3390/rs11161894>
- Yeh, K. C., Liu, C. H., & Hearn, A. L. (1979). Propagation of gravity wave spectra in the thermosphere. *Journal of Geophysical Research*. <https://doi.org/10.1029/JA084IA03P00834>
- Yermolaev, Y. (2014). Whether duration of the recovery phase of magnetic storm depends on the development rate of storm at its main phase. arXiv: Space Physics. <https://doi.org/10.1134/S0016793215040039>
- Yiğit, E., & Medvedev, A. S. (2015). Internal wave coupling processes in Earth's atmosphere. *Advances in Space Research*. <https://doi.org/10.1016/J.ASR.2014.11.020>
- Yildirim, O., Salgin, O., & Bakici, S. (2011). The Turkish CORS network (TUSAGA-Aktif). In FIG working week (pp. 18–22). Retrieved from http://fig.net/resources/proceedings/fig_proceedings/fig2011/papers/ts03g/ts03g_yildirim_algin_et_al_5244.pdf

- Zakharenkova, I., Astafyeva, E., & Cherniak, I. (2016). GPS and GLONASS observations of large-scale traveling ionospheric disturbances during the 2015 St. Patrick's Day storm. *Journal of Geophysical Research: Space Physics*, 121(12), 12,138-12,156. <https://doi.org/10.1002/2016JA023332>
- Zettergren, M. D., & Snively, J. B. (2019). Latitude and Longitude Dependence of Ionospheric TEC and Magnetic Perturbations From Infrasonic-Acoustic Waves Generated by Strong Seismic Events. *Geophysical Research Letters*, 46(3), 1132–1140. <https://doi.org/10.1029/2018GL081569>
- Zhang, S. R., Vierinen, J., Aa, E., Goncharenko, L. P., Erickson, P. J., Rideout, W., Coster, A. J., & Spicher, A. (2022). 2022 Tonga Volcanic Eruption Induced Global Propagation of Ionospheric Disturbances via Lamb Waves. *Frontiers in Astronomy and Space Sciences*, 9. <https://doi.org/10.3389/fspas.2022.871275>
- Zolesi, B., & Cander, L. R. (2014). Ionospheric Spatial and Temporal Variations. https://doi.org/10.1007/978-3-642-38430-1_4

Appendix A

A.1 Ionosonde Measurements for the Türkiye Earthquake




Figure A.1: Successive ionograms at European stations (from top to bottom: (a, b) Juliusruh, (c, d) Dourbes, (e, f) Roquetes, (g, h) Fairford, and (i, j) El Arenosillo) registering a clear multiple-cusp signature. The left column shows the last quiet ionogram, while the right shows the first disturbed ionogram.




Figure A.2: Successive ionograms at European stations (from top to bottom: (a, b) Juliusruh, (c, d) Dourbes, (e, f) Roquetes, (g, h) Fairford, and (i, j) El Arenosillo) registering a clear multiple-cusp signature. The left column shows the last quiet ionogram, while the right shows the first disturbed ionogram.




Figure A.3: Travel-time diagram of Rayleigh waves detected in seismograms and co-seismic ionospheric disturbances causing multiple-cusp signatures (MCS) at each ionosonde. Red and blue dashed lines show the linear best fit to seismic wave signatures in blue and to the appearance of earthquake-related disturbances (MCS) in red.

A.2 Additional results for the August 2023 tropospheric storm

A.2.1 GNSS receivers




Figure A.4: The top panel shows the GNSS single station TEC, while the middle panel shows the dTEC and its moving standard deviation calculated on a 45-minute window. Bottom panel shows the wavelet spectra at the stations co-located with San Vito ionosonde from 25 to 29 August 2023.




Figure A.5: The top panel shows the GNSS single station TEC, while the middle panel shows the dTEC and its moving standard deviation calculated on a 45-minute window. Bottom panel shows the wavelet spectra at the stations co-located with the Ebro ionosonde from 25 to 29 August 2023.

A.2.2 Doppler sounder




Figure A.6: The top panel shows Doppler shift (signals from individual transmitters are artificially offset by about 4 Hz) for the analysed period at 4.65 MHz (left side, a–c) and at 7.04 MHz (right side, d–f), while the panels in the middle and at the bottom give information on RMS values of Doppler shift in different frequency (period) bands and azimuths of the wave propagation, respectively.

A.2.3 Ionosondes

Table A.1: Summary of the TID activity observed in the traces of the ionograms at European stations during the front passage between 25 and 29 August 2023 (the dates can be seen in the first column). The columns indicate the hours in UT. The investigated stations are Dourbes (DB049), Ebro (EB040), El Arenosillo (EA036) from North to South in the Western Europe longitudinal sector, and Juliusruh (JR055), Pruhonice (PQ052), Sopron (SO148), and San Vito (VT139) from North to South in the Central Europe longitudinal Sector. The colours indicate the intensity of the TID activity as seen on the ionograms (if at least two ionograms during that hour were affected). Light orange is used when the affected frequency range is less than 0.5 MHz, while the darkest orange denotes the case where most of the F trace was distorted. The presence of a Sporadic E layer is indicated with the label “Es”.

Date; UT->	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	<- UT; Date
Irregularity < 0,5 MHz																									
Irregularity > 0,5 MHz																									
Irregularity > 1,5 MHz																									
Most of the F trace is distorted																									
Western Europe longitudinal sector																									
DB049																									
2023.08.25	Es	Es	Es	Es	Es		Es	Es				Es			Es	Es			Es	Es	Es	Es	Es	Es	2023.08.25
2023.08.26	Es	Es	Es				Es			Es		Es	Es	Es	Es	Es			Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27					Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es		Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28				Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es		Es	Es	Es	Es	Es	Es	2023.08.28
2023.08.29			Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es		Es	Es	Es	Es	Es	Es	2023.08.29
EB040																									
2023.08.25	Es				Es	Es									Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.25
2023.08.26				Es					Es							Es	Es		Es			Es	Es	Es	2023.08.26
2023.08.27									Es		Es	Es							Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28									Es										Es	Es	Es	Es	Es	Es	2023.08.28
2023.08.29					Es				Es										Es	Es	Es	Es	Es	Es	2023.08.29
EA036																									
2023.08.25				Es	Es	Es	Es	Es			Es	Es					Es	Es		Es	Es	Es	Es	Es	2023.08.25
2023.08.26										Es	Es					Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27	Es							Es	Es	Es							Es	Es	Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28																	Es			Es	Es	Es	Es	Es	2023.08.28
2023.08.29								Es				Es	Es	Es	Es				Es	Es	Es	Es	Es	Es	2023.08.29
Central Europe longitudinal sector																									
JR055																									
2023.08.25	Es	Es	Es		Es	Es	Es	Es	Es	Es									Es	Es	Es	Es	Es	Es	2023.08.25
2023.08.26				Es				Es						Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27	Es	Es	Es	Es	Es			Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28	Es	Es	Es	Es															Es	Es	Es	Es	Es	Es	2023.08.28
2023.08.29	Es	Es	Es	Es	Es	Es					Es	Es	Es	Es	Es				Es	Es	Es	Es	Es	Es	2023.08.29
PQ052																									
2023.08.25							Es	Es	Es							Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.25
2023.08.26	Es	Es					Es		Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27	Es	Es	Es				Es				Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28			Es	Es			Es	Es		Es	Es									Es					2023.08.28
2023.08.29	Es	Es	Es	Es															Es			Es			2023.08.29
SO148																									
2023.08.25										Es						Es									2023.08.25
2023.08.26									Es		Es	Es			Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27									Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28					Es	Es	Es															Es	Es	Es	2023.08.28
2023.08.29			Es				Es											Es				Es			2023.08.29
VT139																									
2023.08.25		Es		Es			Es				Es			Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.25
2023.08.26							Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.26
2023.08.27							Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	Es	2023.08.27
2023.08.28									Es	Es	Es	Es													2023.08.28
2023.08.29					Es		Es	Es	Es										Es	Es					2023.08.29

A.2.4 Atmospheric weather system




Figure A.7: Surface pressure maps provided by the Deutscher Wetterdienst available at <https://www.wetterzentrale.de/de/reanalysis.php?map=1&model=dwd&var=45>, accessed on 19 January 2025, geographical coordinates and a layer with a map of the Czech Republic was added. Surface pressure is represented by solid lines with 10 hPa steps. Atmospheric fronts (red line with semicircles indicating the direction of the warm front, blue line with blue triangles indicating the direction of the cold front, and purple lines with alternating triangles and semicircles indicating the direction of the occluded front) and the locations of the centres of high (H) and low (T) pressure systems are also shown. The orange branch line indicates a squall line.

Appendix B

B.1 Savitzky-Golay and polynomial detrending parameters

An additional study was performed to test the sensitivity of detrending techniques to parameter choices, with the aim of ensuring a fair comparison among methods and avoiding bias toward specific settings. This analysis focused primarily on the Polynomial and Savitzky–Golay (SGOLAY) filters, while analogous reasoning was applied to the remaining techniques.

B.1.1 Polynomial Detrending

Figures B.1 and B.2 illustrate the performance of Polynomial detrending using 4th, 5th, 8th, 8th and 10th orders. The results were evaluated in terms of amplitude and correlation coefficient errors for both LSTID and MSTID scenarios. The cumulative distribution functions (CDFs) of amplitude error show that lower-order polynomials (4th and 5th) can reduce errors in the MSTID case, while higher orders (9th and 10th) provide more consistent results in the LSTID case. The violin plots of correlation coefficient error further confirm these trends: 5th-order performs best for LSTIDs, while 10th-order ensures optimal results for MSTIDs. Based on these outcomes, the polynomial degree was set to 5th for LSTID analysis and 10th for MSTID analysis, representing a balanced choice between accuracy and robustness.




Figure B.1: CDF of amplitude error and amplitude error distribution for LSTID and MSTID scenarios for different polynomial detrending settings.




Figure B.2: Violin plots of correlation coefficient error for different polynomial orders in LSTID and MSTID scenarios for different polynomial detrending settings.

B.1.2 Savitzky–Golay Filtering

Figures B.3 and B.4 show the results of the SGOLAY detrending method for different polynomial orders (2nd and 3rd) and varying sliding windows expressed as multiples of the modeled wave period (1.5T, 2T, 3T). The analysis highlights that the 2nd-order polynomial achieves the lowest amplitude and correlation errors across LSTID and MSTID scenarios. Regarding the time window, optimal values were found at 3T for LSTIDs and 2T for MSTIDs. Nevertheless, slightly larger windows were adopted to avoid unfair optimization tied to prior knowledge of the wave period: twice the moving average window length for LSTIDs and 2.5 times the MA window length for MSTIDs. This choice ensures applicability to real scenarios, where the true TID period is unknown in advance.




Figure B.3: CDF of amplitude error and amplitude error distribution for different polynomial orders and sliding windows in LSTID and MSTID scenarios for different Savitzky-Golay filter settings.




Figure B.4: violin plots of correlation coefficient error for different polynomial orders and sliding windows in LSTID and MSTID scenarios for different Savitzky-Golay filter settings.

Appendix C

C.1 Impact of GNSS network size on keogram IPP per pixel

No specific study was performed to identify the optimal number of clusters; instead, the focus was on finding a number of stations that ensured that the keogram image pixels were mostly covered. Below is shown a keogram that, instead of the dTEC, shows the number of IPP that fall inside each pixel. The three keograms show, respectively, the full EUREF network for the area of interest (Figure C.2), the 100-station configuration (Figure C.3), and the results for 50 stations (Figure C.4). Figure C.1 shows the three different station networks in different colors. Comparing Figure C.2, Figure C.3, and Figure C.4, it is possible to note how the amount of empty pixels (colored in white) for the 100 stations and the full network is quite similar. At the same time, the number of pixels that show more than 5 IPP per pixel reduces in the 100-station network. On the other hand, the 50-station network shows many empty pixels, especially at high and low latitudes. While small gaps can be filled with interpolation, huge gaps would make the interpolation less reliable and more prone to introducing outliers and noise in the image. At the same time, a high amount of highly covered pixels (shown in red, more than 10 IPP per pixel) is detrimental from a computational point of view. Thus, the choice was to use the 100-station solution, as it should balance IPP coverage and computational time well.




Figure C.1: Different network sizes considered for the generation of the IPP number keograms. In green is the full EUREF network for the area of interest. The 100 and 50 station subnetworks obtained with K-means clustering are shown in red and blue.




Figure C.2: IPP density distribution over time and latitude using data from 175 GNSS stations on February 23, 2024.




Figure C.3: IPP density distribution over time and latitude using data from 100 GNSS stations on February 23, 2024.




Figure C.4: IPP density distribution over time and latitude using data from 50 GNSS stations on February 23, 2024.

C.2 Investigating the amplitude threshold impact

The first parameter investigated is the amplitude of the threshold used to obtain the positive and negative masks. The parameter changed is the amplitude threshold value, which varies across 0.01, 0.015, 0.025, and 0.05 TECu.

The detection rate is higher than 100% for a SymH value of 60, visible in Figure C.5, due to the reduction of the time resolution of the LSTID ROI timing to 30 minutes. This means that two LSTIDs might have the same reference time, thus increasing the detection rate above 100%.

As expected, increasing the amplitude threshold decreases both the number of detections and the detection rate. Additionally, the hourly distribution changes depending on the chosen threshold. Specifically, when the threshold is raised to 0.025, detections after sunset decrease consistently. This behaviour is justified because post-sunset LSTIDs typically have lower amplitude due to reduced background ionization. The Hp30 detection rate distribution illustrates how the distribution shifts to higher Hp30 values with increasing amplitude thresholds. The same pattern is visible in the IE and SymH distributions. Thus, a lower amplitude threshold would increase the overall number of detections, but it would also raise the number of possible false detections (such as LSTIDs occurring without geomagnetic activity, which are considered extremely rare).

Looking at the LSTID parameters, one can observe that the average period decreases as the threshold increases, while the velocity decreases slightly along with its standard deviation. This period change is explained by larger amplitude LSTIDs tending to have shorter periods, as shown in the period and amplitude versus IE plots. In contrast, the velocity is prone to overestimation because it is retrieved based on slope. Lowering the amplitude threshold can lead to false detections with high velocity, which could affect the statistics, especially for LSTIDs detected during periods of low geomagnetic activity.

After carefully evaluating the impact of the amplitude threshold, it was fixed at 0.015 TECu. It is important to note that these parameters can still be fine-tuned to meet the user's specific requirements.




Figure C.5: LSTID statistics obtained with 0.35 minimum solidity, 45 minutes of maximum distance between ROI. The amplitude threshold is set to 0.01 TECu.




Figure C.6: LSTID statistics obtained with 0.35 minimum solidity, 45 minutes of maximum distance between ROI. The amplitude threshold was set to 0.015 TECu.




Figure C.7: LSTID statistics obtained with 0.35 minimum solidity, 45 minutes of maximum distance between ROI. The amplitude threshold was set to 0.025 TECu.




Figure C.8: LSTID statistics obtained with 0.35 minimum solidity, 45 minutes of maximum distance between ROI. The amplitude threshold was set to 0.05 TECu.

C.3 Investigating the minimum solidity impact

Here, the impact of the minimum solidity parameter is studied. The amplitude threshold was fixed at 0.01 TECu, and the maximum time separation was set to 45 minutes. The solidity parameter was tested at 0.2, 0.35, and 0.5. The solidity of a region is defined as the ratio between the area (number of pixels) of the region and the area of its convex hull. The convex hull is the smallest convex polygon that contains all the pixels of the region.

As expected, increasing the minimum solidity threshold impacts the number of detections similarly to the amplitude threshold. Specifically, being stricter on the solidity threshold removes more irregular regions, which are typically associated with less confident detections. Moreover, the difference between a minimum solidity of 0.2 and 0.35 is subtle. The detection rate remains essentially the same for high geomagnetic activity, indicating that the higher number of detections seen in the hourly and monthly distributions is mainly due to low

geomagnetic activity. The literature shows that LSTIDs unrelated to geomagnetic activity are rare, so this increase in detections could be linked to unreliable detections. Raising the solidity threshold to 0.5 has a more significant impact. Specifically, the number of detections and the detection rate decrease, with a slightly larger drop for low geomagnetic activity.

Unlike the amplitude threshold, the solidity threshold has a lower impact on the LSTID parameters. Taking everything into account and keeping in mind that these parameters are not set in stone but are adjustable to meet the required performance of the algorithm, the paper presents results based on a solidity threshold of 0.35.




Figure C.9: LSTID statistics obtained with 0.01 TECu as amplitude threshold and 45 minutes of maximum distance between ROI. The minimum solidity of regions was set to 0.2.




Figure C.10: LSTID statistics obtained with 0.01 TECu as amplitude threshold and 45 minutes of maximum distance between ROI. The minimum solidity of regions was set to 0.35.




Figure C.11: LSTID statistics obtained with 0.01 TECu as amplitude threshold and 45 minutes of maximum distance between ROI. The minimum solidity of regions was set to 0.5.

C.4 Investigating the maximum time separation impact

Here, the impact of the maximum time separation between paired ROIs is investigated. The amplitude threshold was fixed at 0.015 TECu, and the minimum solidity at 0.35. The algorithm results are evaluated for two maximum time separation parameter values, 45 minutes and 60 minutes. This parameter directly determines the maximum period of the LSTIDs that the algorithm can detect: 90 minutes and 120 minutes, respectively.

Figure C.12 compares LSTID ROI occurrences across different geomagnetic disturbance levels, represented by Hp30 and IE. Examining Figure C.12, and in line with the expected behavior of LSTIDs, one can see that the main difference occurs at low levels of geomagnetic disturbance, while during disturbed periods (Hp30 > 4 and IE > 350), the two distributions are nearly identical.

At first glance, one might assume that the increased number of occurrences for the higher time separation (3600 seconds) accounts for higher-period LSTIDs expected during moderately disturbed periods. However, the largest relative increase occurs when there is no geomagnetic activity ($Hp30 < 3$). The solar terminator could also be responsible for this increase in occurrences.

Figure C.13 shows the ratio between detections for 3600 and 2700 seconds, indicating that the highest ratios are observed when there is no geomagnetic disturbance. Since LSTIDs are primarily generated by energy injection at high latitudes, the maximum time separation was fixed at 2700 seconds for the results presented in the manuscript.

As with the other parameters, if a user is specifically interested in longer-period LSTIDs, this parameter can be fine-tuned to meet those needs.




Figure C.12: Comparison between the LSTID ROI number for different levels of geomagnetic indexes for the maximum time separation of 2700 seconds (orange line) and 3600 seconds (blue line).




Figure C.13: Ratio between the number of detections for 60 minutes and the number for 45 minutes. The ratio is plotted against the Hp30 value on the left, and on the right, it is plotted against the IE value.

C.5 Investigating the max area of the largest concavity impact

Here, the impact of a secondary parameter is investigated. The results are obtained using an amplitude threshold of 0.015 TECu, a minimum solidity of 0.35, and a time separation of 45 minutes. The parameter of interest is the maximum area of concavity. This parameter serves a specific purpose: after connected regions have been identified and before they are filtered, there is an intermediate step. Connected regions with a latitudinal elongation greater than 10 degrees and a temporal extent of more than 3000 seconds undergo a sub-filtering process. In this step, the alphashape algorithm is applied to find the concave boundary of the connected region. After the concave boundary is determined, the concave and convex hulls are compared to identify all the different concavities within the connected region. If the largest connected concavity area exceeds the defined threshold, clustering with the Bayesian Gaussian Mixture Model (BGMM) is performed to split the original connected region to remove the large concavity within the data. A simple sketch is presented in Figure C.14 to better illustrate the scenario this step addresses.

Figures C.15 and C.16 help compare the impact of a larger maximum concave area. The overall distribution remains similar for both values, even though the maximum concave areas considered (1400 and 700 pixels) differ by a factor of two. The main difference is the slightly higher number of detections for the threshold of 700 pixels. Additionally, the 1400-pixel threshold shows fewer detections, particularly between 19 and 24. Given this small difference, a maximum largest concave area of 700 pixels was used to obtain the results presented in this thesis.




Figure C.14: Sketch of the max concavity filter and BGMM step. The starting point is the red connected region, with its convex hull drawn in green. The yellow curve represents the concave hull obtained with alphashape. The blue regions show the difference between the convex and concave hulls. As the largest concave region is large, BGMM is applied to obtain the three new areas represented with numbers 1, 2, and 3, which will be filtered as different regions.




Figure C.15: LSTID statistics obtained with 0.015 TECu as amplitude threshold, minimum solidity 0.35, and a time separation of 45 minutes. The maximum area of the largest concavity here is set to 700 pixels.




Figure C.16: LSTID statistics obtained with 0.015 TECu as amplitude threshold, minimum solidity 0.35, and a time separation of 45 minutes. The maximum area of the largest concavity here is set to 1400 pixels.