




## Exploring AI Progress in GNSS Remote Sensing: A Deep Learning Based Framework for Real-Time Detection of Earthquake and Tsunami Induced Ionospheric Perturbations

Michela Ravanelli<sup>1</sup> , Valentino Constantinou<sup>2</sup> , Hamlin Liu<sup>3</sup>, and Jacob Bortnik<sup>3</sup> 

<sup>1</sup>DICEA, Sapienza University of Rome, Rome, Italy, <sup>2</sup>Terran Orbital Corporation, Irvine, CA, USA, <sup>3</sup>Atmospheric and Oceanic Sciences Department, University of California, Los Angeles, CA, USA

**Key Points:**

- The proposed deep learning-based framework can detect earthquake and tsunami induced ionospheric perturbations in Total Electron Content observations
- We used a multi-event approach, focusing on the Pacific area, to train and to test the framework, achieving 91% of F1 score and 84% of recall
- The framework is well-suited for real-time applications, making it readily deployable to enhance tsunami early warning systems

**Correspondence to:**

M. Ravanelli,  
michela.ravanelli@uniroma1.it

**Citation:**

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, e2024RS008016. <https://doi.org/10.1029/2024RS008016>

Received 26 APR 2024  
Accepted 20 AUG 2024

**Abstract** Global Navigation Satellite System Ionospheric Seismology investigates the ionospheric response to earthquakes and tsunamis. These events are known to generate Traveling Ionospheric Disturbances (TIDs) that can be detected through GNSS-derived Total Electron Content (TEC) observations. Real-time TID identification provides a method for tsunami detection, improving tsunami early warning systems (TEWS) by extending coverage to open-ocean regions where buoy-based warning systems are impractical. Scalable and automated TID detection is, hence, essential for TEWS augmentation. In this work, we present an innovative approach to perform automatic real-time TID monitoring and detection, using deep learning insights. We utilize Gramian Angular Difference Fields (GADFs), a technique that transforms time-series into images, in combination with Convolutional Neural Networks (CNNs), starting from VARION (Variometric Approach for Real-time Ionosphere Observation) real-time TEC estimates. We select four tsunamigenic earthquakes that occurred in the Pacific Ocean: the 2010 Maule earthquake, the 2011 Tohoku earthquake, the 2012 Haida-Gwaii, the 2015 Illapel earthquake. The first three events are used for model training, whereas the out-of-sample validation is performed on the last one. The presented framework, being perfectly suitable for real-time applications, achieves 91.7% of F1 score and 84.6% of recall, highlighting its potential. Our approach to improve false positive detection, based on the likelihood of a TID at each time step, ensures robust and high performance as the system scales up, integrating more data for model training. This research lays the foundation for incorporating deep learning into real-time GNSS-TEC analysis, offering a joint and substantial contribution to TEWS progression.

**Plain Language Summary** Global Navigation Satellite System Ionospheric Seismology investigates how the ionosphere responds to earthquakes and tsunamis, detecting TIDs through GNSS-derived TEC observations. Real-time TID identification aids tsunami detection, enhancing early warning systems by extending coverage to open-ocean regions. Automated TID detection is crucial for early warning system improvement. In this study, we propose an innovative approach using deep learning insights to perform automatic real-time TID monitoring and detection. We leverage GADFs and CNNs with VARION real-time TEC estimates. We train the model on four tsunamigenic earthquakes in the Pacific Ocean and validate it on an out-of-sample event. The framework achieves promising performance metrics, highlighting its potential for real-time applications. Our approach improves false positive detection, ensuring robustness and scalability as the system integrates more data for training. This research paves the way for integrating deep learning into real-time GNSS-TEC analysis, contributing significantly to the advancement of early warning systems.

### 1. Introduction

Nowadays, it is commonly renowned that the Total Electron Content (TEC) retrieved from Global Navigation Satellite System (GNSS) can be used to estimate fundamental information to assess the genesis of a tsunami (Astafeyeva, 2019; Manta et al., 2020; Meng et al., 2019; Ravanelli et al., 2021). This particular branch of GNSS Remote Sensing is referred as GNSS Ionospheric Seismology since it studies the ionospheric response to natural/man-made hazards through GNSS-TEC observations.

Indeed, earthquakes and tsunamis, as well other natural hazards, can produce acoustic and gravity waves (AGWs) which, due to the decrease in atmospheric density, can reach ionospheric heights, causing disturbances in the

electron content (Daniels, 1952; Hines, 1972, 1974; Peltier & Hines, 1976). Some of these perturbations can be classified as acoustic gravity wave epicenter ( $AGW_{epi}$ ) and internal gravity waves tsunami ( $IGW_{tsuna}$ ), following the nomenclature firstly introduced by (Occhipinti, 2015). Particularly,  $AGW_{epi}$  can provide timely information (after 8 min from the earthquake) of the vertical displacement at the source, a fundamental parameter for the estimation of tsunami genesis (Astafyeva, 2019; Ravanelli et al., 2021). Furthermore,  $IGW_{tsuna}$  can share insights about the propagation of the tsunami offshore, that is where data are lacking the most (Hohensinn et al., 2024).  $AGW_{epi}$  and  $IGW_{tsuna}$  can be generally referred as Traveling Ionospheric Disturbances (TIDs) (Azeem et al., 2017; Vergados et al., 2020).

This is the reason why GNSS-TEC information needs to be used in real-time since it provides continuing updates on tsunami potential and arrival times, reducing the risk of false alarms (LaBrecque et al., 2019) and, hence, it can be used to enhance tsunami early warning systems (TEWS) (Kamogawa et al., 2016; Manta et al., 2020; Ravanelli et al., 2021; Savastano et al., 2017). This goal was expressed in Resolution #4 of the International Union of Geodesy and Geophysics General Assembly 2015 which encourages the real-time GNSS Augmentation of the Tsunami Early Warning System (International union of Geodesy and Geophysics, 2015). Indeed, existing systems such as the Deep-ocean Assessment and Reporting of Tsunamis (DART) can be effective, but they are often limited to specific geographic locations, due to the use of specialized hardware (the DART system utilizes ocean-based buoys placed in strategic locations around the Pacific Ocean).

Operative tools like the GUARDIAN system, developed by NASA Jet Propulsion Laboratory and based on near real-time GNSS-TEC estimates for natural hazard monitoring, do not provide a real-time detection module yet (Martire et al., 2023). Indeed, what is still lacking is a scalable, real-time and automated detection of TIDs: this feature stands as the pivotal element to augment TEWS. Furthermore, the volume, variety and velocity of available GNSS data cause to also explore machine learning approaches which can effectively capture complex non-linear relationships. In this direction, it is worth mentioning the work from (Brissaud & Astafyeva, 2022) that explores Random Forest models for detecting TIDs, but to date these models require significant feature engineering, leveraging not only TEC and but also ionospheric spectrograms. Furthermore (Fuso et al., 2024), also explored classification architectures for automatically detecting earthquake-induced TIDs, albeit with a smaller data set and testing on a single event. Recently, deep learning has been employed thanks to its ability to automatically discover complex features without the need for domain knowledge (Munir et al., 2019). This trait can be leveraged and used toward TID detection. TID detection may be characterized as a distinct anomaly detection problem involving multivariate time-series data. The advancements of deep learning and the development of new network architectures have led to improvements in the performance of time series anomaly detection solutions (Bontemps et al., 2017; Chauhan & Vig, 2015; Hundman et al., 2018; Malhotra et al., 2015, 2016; Nanduri & Sherry, 2016; Taylor et al., 2016) which—together with new data transformation approaches (Wang & Oates, 2015)—may present an alternative approach for TID detection over Random Forests. Automatic TID detection approaches have the potential to greatly improve the detection of tsunami waves and the characterization of tsunami events by providing open-ocean detection capability unconstrained by a fixed geographic locations such as the DART system.

In this work, for the first time, we decided to use the insights of deep learning for TIDs detection on a multiple tsunamigenic events. This approach is far from self-evident, since GNSS-TEC observations are time series and not images. Yet encoding the data as images can provide significant advantages, as it allows the use of modeling approaches not reliant on continuous data streams. We applied Gramian Angular Difference Fields (GADFs) (Wang & Oates, 2015)—a method which converts time-series data to images—and Convolutional Neural Networks (CNNs)—a type of deep neural network (Ranzato et al., 2007; Szegedy et al., 2015; Wickramaratne & Mahmud, 2021)—for tsunami-induced TID detection starting from TEC observations. In detail, we employed VAR-ION (Variometric Approach for Real-time Ionosphere Observation), that is, a well-known algorithm to estimate TEC variations in real-time (Ravanelli et al., 2021; Savastano et al., 2017).

We selected four tsunamigenic earthquakes that occurred in the Pacific Ocean, as the most devastating tsunamigenic events often occur in that region: the 2010 Maule earthquake, the 2011 Tōhoku earthquake (coastal Japan tsunami), the 2012 Haida-Gwaii earthquake (coastal Hawaii tsunami) and the 2015 Illapel earthquake (coastal Chile tsunami). The first three events were used for model training, whereas the out-of-sample validation was performed on the last one.

In detail, extending the research conducted in the previous study by (Constantinou et al., 2023a, 2023b), this work provides an in-depth description of the proposed deep learning framework, offering comprehensive insights into its structure and methodology.

The paper is structured as follows. In Section 3, we describe the framework for using deep learning to detect TIDs, and assess the generalizability of the framework by utilizing multiple, separate events for model training and an unseen, future event for validation. Section 4 describes the proposed framework and the resulting model performance in this paper. In Section 5, we present an in-depth discussion of the experimental results, valuable lessons learned and observations from our efforts which we believe may be important in further developing deep learning-based TID detection systems. A description of the practical trade-offs of employing simple false positive mitigation strategies and the importance of employing data management best practices is also provided. Finally, Section 6 illustrates conclusions and future works.

## 2. Earthquake Information and Data Set

As stated before, we focused on four major Pacific tsunamigenic earthquake events. In detail, we used the 2010 Maule earthquake, the 2011 Tōhoku earthquake, the 2012 Haida-Gwaii earthquake and the 2015 Illapel to assess the viability of data transformation and modeling strategies that we will describe in Section 3.

To this aim, different GNSS data sets were employed for the respective events.

On 27 February 2010 an earthquake of magnitude 8.8 struck the Chilean coastal region of Maule (36.1221°S 72.8981°W) (USGS, 2021b). The shock triggered a tsunami that propagated all over the Pacific region, reaching over 700 km of coastline (Istituto Nazionale di Geofisica e Vulcanologia, 2021). The maximum run-up (i.e., the maximum topographic height reached by the tsunami) peak (29 m) was recorded at Constitución, Chile (Yue et al., 2014). The Pacific Tsunami Warning Center (PTWC) issued the warning 12 min after the earthquake (EQ) (Soulé, 2014). The tsunami arrived within 30 min at many locations in Chile, therefore, official evacuations and warnings by local authorities were not available at many places prior to the arrival of the tsunami (Center, 2021b). Finally, the tsunami accounts for 124 victims concentrated in the coastal regions of Maule and Biobío, Juan Fernández Archipelago's Robinson Crusoe Island and Mocha Island (Fritz et al., 2011). For this event, we used a data set of 30 GPS stations located in Chile from the UNAVCO network (UNAVCO, 2021). For that event, we don't have the data from the regional Chilean network of Centro Sismológico Nacional, Universidad de Chile (CSN) network that started at collecting data in 2010, but after the EQ. Nominally, we investigated a data set of 34 days, 20 of which before the EQ and 13 after the EQ at a 30 s rate.

The infamous Tōhoku-Oki earthquake of magnitude Mw 9.1 occurred near the northeast coast of Honshu (38.297°N 142.373°E) in the Tōhoku region on 11 March 2011 (USGS, 2021c). The shock triggered powerful tsunami waves that struck the Pacific coast of Honshu within about 20 min and that was observed all over the Pacific region. 15270 and 8499 people were reported to be killed and missed respectively because of the earthquake and tsunami. In Sendai, maximum tsunami run-up heights (15–20 m range) were registered. The Japan Meteorological Agency's national tsunami warning center issued a tsunami warning 3 min after the earthquake triggering the alerting process that immediately broadcasted by mass media and locally activated sirens and other mitigation countermeasures such as flood gate closures. Nevertheless, many casualties resulted: waves overtopped tsunami walls and destroyed many structures, especially wooden homes (Center, 2021a). In this case, we had available an extraordinary data set composed of 1237 GPS receivers belonging to the very dense Japanese GEONET network (Geospatial Information Authority of Japan, 2021). Precisely, stations sampling data at 30 s rate were employed. The data set includes 30 days, nine of which before the EQ and 19 after the EQ.

On 28 October 2012 an earthquake of moment magnitude 7.8 hit offshore of the island of Haida Gwaii, Canada (52.788°N 132.101°W) (USGS, 2020). The quake engendered a non-destructive tsunami that was registered throughout the Pacific, hitting the coast of Alaska, of British Columbia, of California and of Hawaii. The PTWC issued a tsunami warning (19:09 HST 27 October) that was then downgraded (01:01 HST 28 October). Observations of tsunami waves up to 1.5 m were registered in Maui and the Hawaii Island, while no great damages were reported statewide (Center, 2021c). For this case study, we collected a data set composed of 56 GPS stations placed on Hawaii islands belonging to UNAVCO network (UNAVCO, 2021). Thoroughly, 15 days were analyzed: 12 days before the EQ and 2 after the EQ at 30 s rate.

**Table 1**  
Events Are Sorted According to the Event Year

Earthquake characteristics		
	Year	Magnitude
<i>Maule</i>	2010	8.8
<i>Tohoku</i>	2011	9.1
<i>Haida-Gwaii</i>	2012	7.8
<i>Illapel</i>	2015	8.3

Finally, we used the 2015 Illapel earthquake and tsunami to test our model. The event took place 46 km offshore of the Coquimbo region on 16 September 2015 with a moment magnitude of 8.3 (31.573°S 71.674°W) (USGS, 2021a). The tremor generated a tsunami that spread across the Pacific Ocean. Tsunami waves heights up to 9 m on the coast were measured between 29°S and 32°S and smaller further south and north. Along the Chilean coast, the PTWC and National Hydrographic and Oceanic Service (SHOA) issued tsunami threat messages 7 and 8 min following the earthquake, respectively. Tsunami linked casualties were minimized by these prompt messages and evacuation (Satake & Heidarzadeh, 2017). Specifically, we collected data from 80 GPS stations located in Chile from Centro Sismológico Nacional,

Universidad de Chile (CSN) for the 2015 Illapel earthquake and tsunami. These GPS receivers gather data at 15 s rate (Centro Sismológico Nacional, 2021). Nominally we analyzed 26 days starting from 19 days before the EQ to 6 days after.

To have a visual overview of the selected events, we reported the epicenters and magnitudes of the selected earthquakes in Table 1 and in Figure 1, respectively.

Finally, it is important to recall that all these events were selected since they were recognized to generate clear ionospheric perturbations as previously studied in the literature: the 2010 Maule earthquake (Galvan et al., 2011; Ravanelli et al., 2020), the 2011 Tohoku earthquake (Astafyeva et al., 2011; Rolland et al., 2011), the 2012 Haida-Gwaii earthquake (Grawe & Makela, 2015; Savastano et al., 2017) and the 2015 Illapel earthquake (Meng et al., 2022; Ravanelli et al., 2021; Shrivastava et al., 2021).

### 3. Methods

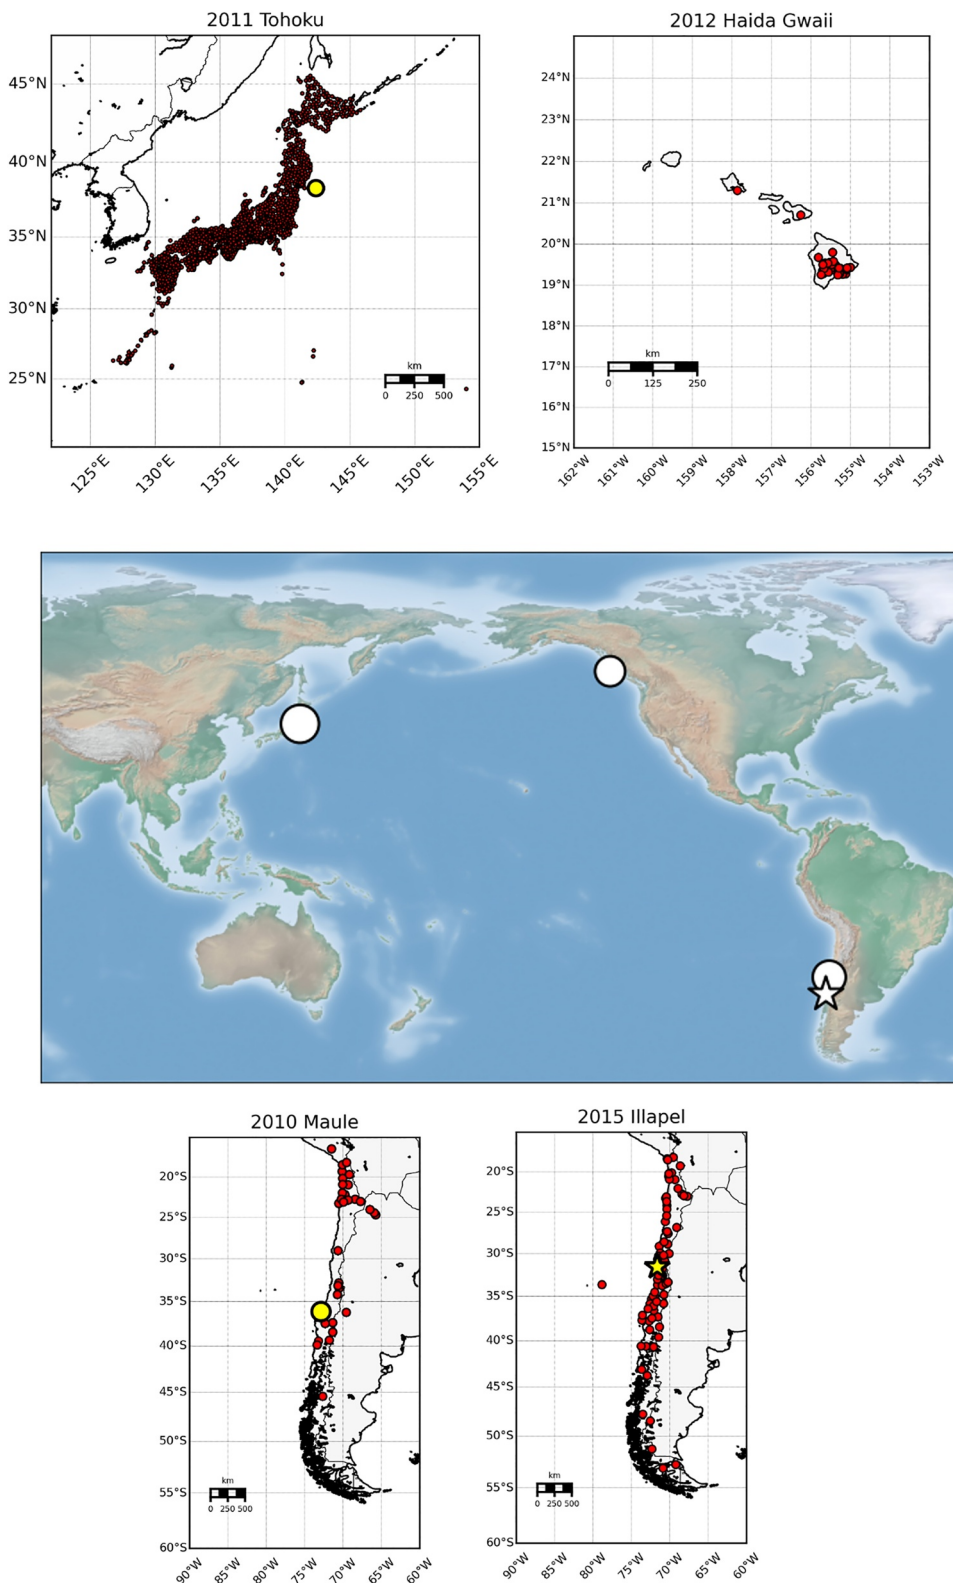
#### 3.1. Real-Time TEC Data

The retrieval of TEC data in real-time represents the starting point of our work. In this work, we employed the VARION (Variometric Approach for Real-Time Ionosphere Observation) algorithm to retrieve TEC estimates in real-time (Ravanelli et al., 2021; Savastano et al., 2017). In detail, VARION is based on the single time differences of geometry-free combination of GNSS carrier-phase observations, being suitable for real-time applications (Fortunato et al., 2019; Fratarcangeli et al., 2018; Ravanelli et al., 2020, 2023; Savastano et al., 2017, 2019). Indeed, VARION only relies on a standalone GNSS receiver and on standard GNSS broadcast products: these information are available in real-time. Here, we refer to slant TEC (sTEC, TEC on the line of sight satellite-receiver) variations related to every station-satellite pairs (Ravanelli et al., 2021). More specifically, we employed sTEC variations over time ( $\delta sTEC/\delta t$  [TECU/s]), that represent the first VARION output and that is obtainable in real-time.

The Ionospheric Pierce Point (IPP) is the point to which every TEC measurement is referred in a 2D representation of the ionosphere and is given by the intersection of the line of sight satellite-receiver with the single thin layer used to approximate ionosphere. The sub-Ionospheric Pierce Point (SIP) represent the projections onto the ellipsoid of IPP. Finally, it is important to highlight that we can apply a machine or deep learning algorithms since with the high volume and breadth of data produced. Indeed, several TEC measurements can be retrieved for every GNSS station: nominally, one for each satellite in view (Ravanelli et al., 2021). Nowadays, thanks to the actual multiplication of GNSS systems -i.e., GPS (US), GLONASS (Russia), Galileo (EU), BeiDou (China), to mention only the ones that reached the full operational capability—the ionospheric coverage is increasing day by day, opening ambitious perspectives for GNSS Ionospheric Seismology especially in a machine learning outlook.

#### 3.2. Deep Learning Framework

Deep learning approaches have been developed that effectively model temporal information. Recurrent Neural Networks (RNNs), and more specifically Long Short-Term Memory networks, have shown the ability to capture long-term dependencies and their ability to handle high-complexity, sequential and temporal data has resulted in a wide variety of applications using time-series data (Bontemps et al., 2017; Hundman et al., 2018; Malhotra et al., 2015, 2016). Similarly, CNNs—neural networks well suited for image recognition and classification problems (Ranzato et al., 2007; Szegedy et al., 2015; Wickramaratne & Mahmud, 2021)—have demonstrated



**Figure 1.** Central map showing the epicenters of the tsunamigenic earthquakes used in model training. Those used in the training set are represented as circles, with any events used in out-of- sample validation shown as stars. Side maps representing the data set related to each event used in this study.

impressive performance several benchmark data sets such as the Modified Institute of Standards and Technology (MNIST) database of hand-written digits, where CNNs were demonstrated to achieve around 99.77% accuracy (Ciresan et al., 2012). Recently, the introduction of the Gramian Angular Summation Field (GASF), Gramian Angular Difference Field (GADF) and Markov Transition Field (MTFs) (Wang & Oates, 2015) has led to applications of CNNs for time series anomaly detection (Chen & Tsai, 2019; Xu et al., 2020), a relatively new area of research. The use of CNNs—a computer vision (CV) technique—with methods like GADFs provide significant advantages.

GADF—more generally Gramian Angular Fields (GAF)—are methods for encoding time series data into images by converting scaled time-series data observations into a polar coordinate system, and then using the angular components and a specified inner product to construct a Gram matrix. Specifically, GADF computes the inner product with sine of the differences between the angular components of the converted polar coordinates. In effect, the GAF and related methods convert a one-dimensional time series into a two-dimensional image. By converting windows of time series into images, characteristics of the time series (and the relationship to the classification labels) can be captured visually, negating the need to utilize sequence-to-sequence deep learning methods such as RNNs which have difficulty being applied in the presence of significant missing data without feature or model engineering.


Identifying TIDs in time series data is achievable with a variety of signal processing and applications of machine learning (Brissaud & Astafyeva, 2022) due to similarities in the scale and format of data and the task (i.e., classification of a window in time) with other applications of anomaly detection, such as anomaly detection for spacecraft (Hundman et al., 2018). However, the introduction of Gramian Angular Summation Fields (GASF), GADFs (GADF) and Markov Transition Fields (MTF) (MTFs) (Wang & Oates, 2015) and advancements in time series anomaly detection approaches allow the conceptualization and development of a deep-learning-based framework for TID detection which leverages these advancements.

The source code for the proposed deep learning framework is available at <https://doi.org/10.5281/zenodo.12571500> (Constantinou & Liu, 2024).

### 3.2.1. Step 1 Data Sources, Format and Requirements

The input data utilized in this framework are a collection univariate time series  $X = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}\}$  of  $n$  real-valued observations (TEC estimates), with each univariate time series representing a connection between a satellite and ground station (there is a one to many relationship between satellite and ground station (receiver)). Data refresh rates and sampling rates depend on the GNSS and supporting software used to observe, downlink and process data—Glonass, Global Positioning System (GPS), Beidou etc. Different approaches to TEC estimation may be utilized depending on how the system will be applied in practice, the source of the data and other elements described earlier. Machine and deep learning approaches utilized in the context of TID detection—and more broadly anomaly detection—are agnostic to *how* data is produced. Rather, the data—more specifically data from the minority target variable (the anomaly or TID)—must contain distinguishable differences when compared to the majority class (normal behavior) to enable modeling success. In many cases, these differences are visually interpretable as those shown in Figure 2.

Data may be resampled to a different rate from the source data (e.g., from every second to 1 min averages of the original data). Several trade-offs are made when electing to resample the originating data stream. While resampling the data to larger time scales, the computational resources needed to train machine and deep learning models is significantly reduced. However, care must be taken when resampling as to not dampen behaviors in the time series important for detecting TIDs or other disturbances in time series. An additional trade-off is made between computational resources required for modeling and the alert latency of the overall system. Illustrated through a brief example, if data is processed in near-real-time at the minute label, in practice several positive TID detections (e.g., a sustained period of at least several minutes) may first need to be before being considered a true signal of an earthquake or resulting tsunami. Sampling at larger levels (e.g., 1 min) could add several seconds of latency not present when processing data each second. Finer sampling rates could mitigate system latency, but require additional training time, with larger data sets and predictive models. Researchers and developers of deep learning based anomaly detection systems assess these trade offs and select ideal parameters for their system. The “sweet spot” is unique combination of data, decisions made, software and supporting ground systems and capabilities.



**Figure 2.** The diagram illustrates how chronological windows of time series data are encoded into images using Gramian Angular Difference Fields (GADF) for both the anomalous TID (left) and normal (right) classes in a time-series. The pattern shown in the left image is distinct to the anomalous class, an important characteristic that must be produced regardless of encoding method used (GASF, GADF or MTF).

### 3.2.2. Step 2 Data Labeling and Transformation

Described in more detail in the following subsections, we leverage supervised models in this framework, necessitating the need for Subject Matter Expert (SME) labels for model training and validation. As with many anomaly detection applications of machine and deep learning, the target for detection—the earthquake and tsunami induced TID—is a rare event (powerful earthquakes and tsunamis are rare) and short-lived (earthquakes on average last only 10–30 s), adding difficulty to the detection process. Additionally, TIDs generated from tsunami waves may share similar patterns to ionospheric disturbances generated from other phenomena, such as meteorite explosions in the atmosphere (Luo et al., 2020; Yang et al., 2014), volcanic eruptions (Igarashi et al., 1994; Rozhnoi et al., 2014) or large human-driven explosions (Kundu et al., 2021). Careful examination, labeling and SME validation of the data is needed prior to effectively leverage in machine and deep learning.

Thorough understanding of the data and subject matter is of utmost importance to the success of any framework which leverages machine and deep learning. The labels used in model training and validation are of primary importance. Similar to other time series anomaly detection problems, each univariate time series (a TEC estimate for a particular satellite and ground station, e.g., *ahup-G10*) is labeled such that any known anomalies are represented by a start time and finish time (the data contained between those times is considered to be anomalous). These labels are used together with data transformation approaches to alter the data into images, which are then used in both model training and validation.

In early stages of data and model development, training data (historical or simulated) must be labeled—an often arduous and time-consuming process. However, while adjustments can be made to improve model performance, updates to data and labels often have the most profound effect on all elements of the modeling process (by existing further forward in the modeling pipeline). Incorporating labels from multiple SMEs is also an important component of generalizing a model. To ensure models remain relevant after initial training and validation, a human-in-the-loop (HIL) process could be leveraged that presents possible TIDs identified to the system for SMEs for review, in turn providing labels that can be used to update and improve the performance of the model. In the case of supervised learning, these science-driven HIL labels can update previously existing labels in the training data (Waszczak et al., 2017).

Described in more detail in the following subsections, computer vision approaches are utilized to detect tsunami-induced TIDs in this framework. As such, image data is needed for model training, validation and (later) inference when using the trained model in a real-world scenario.

TEC estimates for each satellite and ground station and labels are converted into images. More specifically, for each univariate time series  $X = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}\}$  of  $n$  TEC estimates, a consistent window size  $w$  is used to split

the data into a set of windows. Each window of data (TEC estimates) are converted into images using one of a set of available transformation approaches. Images are sorted according to the ground truth provided. If *any* of the image was generated from a window of data overlapping ground truth ranges, this image is categorized as representing a TID—if not, normal ionospheric TEC. The labels provided earlier in the framework are utilized directly in the creation of training and validation data.

The transformation from float data to image data is made using one of a variety of available approaches. Gramian Angular Summation Fields (GASF), GADFs (GADF) and MTF are described in detail in (Wang & Oates, 2015). Each approach utilizes a separate method for converting time-series data into a 2D matrix which may be represented as an image suitable for CNNs.

As GADF is utilized in our experimental work, we further detail this transformation approach. Specifically, GADF computes the inner product with sine of the differences between the angular components of the converted polar coordinates. Given a time series  $X = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}\}$  of  $n$  real-valued observations, the method first entails rescaling the data to an interval of values between  $[-1, 1]$  which we will denote with  $\tilde{x}_i$  where  $i$  represents the time step in the series and  $\tilde{X}$  as the rescaled time series.

$$\tilde{x}_i = \frac{(x_i - \max(X)) + (x_i - \min(X))}{\max(X) - \min(X)} \quad (1)$$

Now given these rescaled values, the implementation computes the polar coordinates where at the  $i$ th time step,  $\phi_i$  is the angular component and  $r_i$  is the radial component that is computed from the time step and a regularizing factor  $R$ .

$$\begin{aligned} \phi_i &= \arccos(\tilde{x}_i), \quad -1 \leq \tilde{x}_i \leq 1, \quad \tilde{x}_i \in \tilde{X} \\ r_i &= \frac{i}{R}, \quad 1 \leq i \leq n \end{aligned} \quad (2)$$

With the angular components of each time step, a vector  $\phi = \{\phi_1, \phi_2, \dots, \phi_n\}$  can be used to create the GADF by using the sine of the angular difference. This resulting matrix can then be read in as an image (Wang & Oates, 2015).

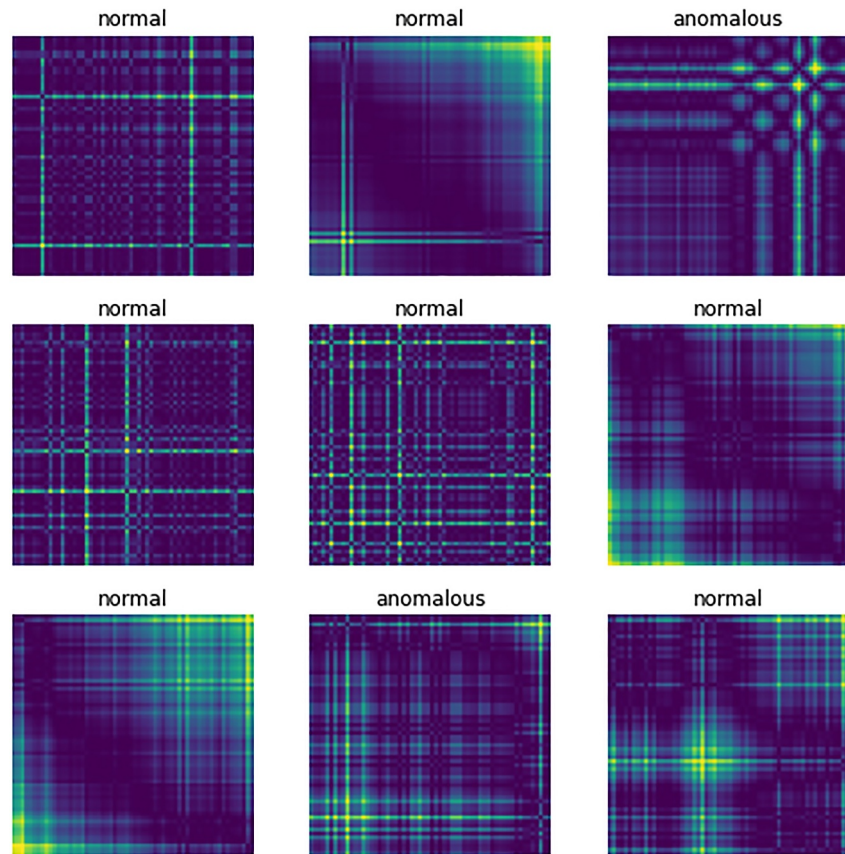
$$GADF = \sin(\phi_i - \phi_j), 1 \leq i, j \leq n \quad (3)$$

This method provides several advantages with regards to time-series data. First, it preserves temporal dependence since the time step increases as position moves from the top-left to the bottom-right of the Gramian Angular Field. The matrix is bijective, that is, the inverse function yields an absolute reconstruction the original time series data given the image (Hong et al., 2020). Second, temporal correlations between each different time step of our windowed data images are easily detectable due to how we construct the Gram Matrix and how it shows every possible angular difference between each time step in the window. Since temporal information is captured in the encoded resulting image and not during model training, GADF—as well as GASF and MTF—are robust to missing data. The ability to capture time-series behavior in images allows artificial intelligence practitioners the ability to use image-based deep learning model architectures for time-series anomaly detection in the face of large periods of missing data, specifically CNN (CNN) architectures that are well-understood, high-performing and readily available as open source software. Lastly, these image transformation approaches produce visually-interpretable differences in images across classes as shown in Figure 3.

### 3.2.3. Step 3 Selecting a Deep Learning Architecture

In this framework, data from multiple univariate time series are utilized to train a deep learning model. A single model is trained for TID detection across data from different satellites and ground stations. Additionally, a single model is trained over multiple distinct tsunami-inducing earthquake events. Training a single model across examples from various geolocations, events of interest, satellites and ground stations provides the model with the exposure to many different scenarios. Since information and characteristics about the time series are transformed into images in this framework, artificial intelligence approaches which learn information from images are needed. This single model is then tested over another set of distinct tsunami-inducing events mostly to measure the





**Figure 3.** Normal and anomalous classes for GADF-generated images. Note the pattern in the top-right image, which is distinct to the anomalous class (a result of using GADF and patterns in the data).

classification performance and test if the model can generalize toward unseen data. Using this validation performance, we can then select the architecture and/or the parameters of the architecture.

A variety of deep learning architectures have been developed specifically for image data, including CNNs. The type of deep learning architecture used in this framework is dependent only on the modeling performance (i.e., whether or not it is satisfactory) on time series data from unseen tsunami-inducing earthquake events and the type of image-encoding approach used. Unsatisfactory performance with one type of architecture does not preclude the failure of another—certain types of architectures, depending on the data and image encoding, may be highly successful and used in production. Certain image encoding (data transformation) approaches like GASF, GADF, and MTF might provide different modeling and classification performance when paired with different types of deep learning architectures, specifically CNNs like AlexNet, ResNet, or VGGNet. This training framework over multiple TID events can be used with any type of deep learning architecture, provided the architecture can be used effectively and reliably with images created using the selected encoding methodology.

### 3.2.4. Step 4 False Positive Mitigation Strategies

Time ranges predicted to represent TIDs may not always span large periods of time, depending on the data used in model training and validation, as well as other factors in the modeling process such as hyperparameter selection. Depending on the modeling results, short time periods may disproportionately represent false positives, or periods classified as TIDs but are in fact representative of noise or normal behavior. In these cases, false positive mitigation strategies may be employed to reduce false positive rates, improving precision.

The simplest strategy for mitigating false positives is to delay TID detection until they are consecutively classified by the machine learning model. Thus, only periods of sustained detection are considered seriously. This simple strategy can have a dramatic improvement in cases when false positive periods are far shorter in overall length

**Table 2**  
The Number of Samples Available in Each Class in the Original Data and the Balanced Data Set, the Latter Which Was Utilized in This Work

Number of samples	Number of samples	
	Original	Balanced
<i>anomalous</i>	90201	90201
<i>normal</i>	18756848	900329

compared to true positives. However, this strategy can also have a significant adverse impact on real-world alerting speed, depending on the sampling rate of the time series used in modeling. For example, a minute-level sampling rate could be delayed by 5 min if it is determined that five consecutive periods of classification as a TID are needed prior to alerting without producing significant amounts of false positives.

An alternative strategy instead considers all ground stations downlinking data from a satellite. More specifically, the set of all univariate time series for a satellite  $\mathbf{X}_s$ —which contains many univariate time series  $X_g = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}\}$  of  $n$  TEC estimates, one time series  $X$  for each ground station—is considered in unison in near-real-time. A boolean vector  $\mathbf{X}_{s,t}$  of  $1 \times g$  dimension, where  $g$  is the number of ground stations, is generated by representing TID detections (anomalies detected) as 1s and normal behavior as 0s. Vector values are summed and divided by  $g$ , generating a float value  $F_{s,t}$  representing the overall share of time series representing a possible TID. A threshold  $T_{s,t}$  is selected such that any time periods  $t$  where  $F_{s,t} > T_{s,t}$  are considered as strong candidates of representing TIDs. These time periods (indices) are compared to individual TID detections within each satellite univariate time series  $X_s$ , with periods of overlap considered as true TID detections by the system. The threshold parameter  $T_{s,t}$  is adjustable, with higher values reducing recall but improving precision. This approach provides a means of ensuring some level of agreement is reached across ground stations before a TID is detected.

#### 4. Results

The framework described is utilized in the context of training a model for TID detection. Precisely,  $sTEC$  variations over time ( $\delta sTEC/\delta t$ ) were used in model training and out-of-sample validation, with labels provided by a SME indicating the start and finish times of TIDs in the training and validation data. Start times and finish times were provided for distinct sets of satellites and ground stations for each earthquake event. These labels were used to generate modeling data using GADF, with the 2010 Maule, 2011 Tōhoku and 2012 Haida-Gwaii tsunami events used for model training and testing. The 2015 Illapel event was used for out-of-sample validation (in out-of-sample validation, the model is exposed to previously unseen data). During validation, a false positive mitigation strategy described earlier in methods is applied. The ResNet (He et al., 2015) architecture (50-layers) is used in this work. Training and validation metrics from using this framework configuration are reported, along with summary statistics about the data set.

We placed our research emphasis on evaluating the effectiveness of the described framework when using data from multiple tsunami-inducing earthquake events. Utilizing multiple historical events for training data is needed to train a generalized model that can be used across geographic regions, satellites and ground stations. While we utilize a CNN architecture in our experiments (ResNet with no pre-trained weights), no comparisons are made between various CNN architectures such as DenseNet (Huang et al., 2016) or VGG (Simonyan & Zisserman, 2015). Similarly, no comparisons are made between various types of image encoding methodologies described earlier in the paper. A “balanced” data set was created by undersampling the *normal* (majority) class such that the minority class represented 10% of the number in the normal class. Table 2 describes the number of samples in each class before and after balancing the data.

##### 4.1. Setup

Each time series data stream corresponds to a particular satellite and ground station. For each unique stream of  $sTEC$  data, 60-min windows of values were converted to images using the described GADF approach and assigned to the appropriate class using the SME derived ground-truth start and finish times for each TID. A single model was trained on the set of training events and subsequently evaluated on multiple unique streams of  $sTEC$  data from ground stations and satellites in a validation set.

Two sets of tsunami-inducing earthquake events are utilized in our experiments. The first set of data is used in model training and generating testing scores, leveraging data available from three tsunami-inducing earthquake events—the 2010 Maule, 2011 Tōhoku and 2012 Haida-Gwaii earthquake—available in the training set. Both anomalous data (which contains TIDs) and normal data from each event are utilized in model training. Training data is randomly sampled prior to model training, with precision, recall, accuracy, and F1-score metrics provided on the test set (see Table 4).

**Table 3**  
*The Parameters Employed in the Experimental Model*

Model parameters	
Architecture	resnet50
Max training iterations	50
Batch size	512
Optimizer	Adam
Beginning learning rate	0.00025
Loss function	Cross-Entropy Loss
Image size (in pixels)	224 × 224

The second set of data contains tsunami-inducing earthquake events occurring after (chronologically) the events in the training set, ensuring validation is performed on data unseen by the model during model training. To simulate real-world use and evaluate the effectiveness of the trained model, the model is utilized on out-of-sample validation data occurring several years after the events in the training data, more specifically the 2015 Illapel earthquake off the coast of Chile. Each minute, every unique stream of data is processed chronologically with a 60-min window of sTEC data, with a GADF-generated image created and assessed by the trained model to contain a TID (anomalous) or not (normal). These labels are concatenated to produce an anomalous sequence (i.e., a range of time representing anomalous behavior). Each labeled ground truth anomalous sequence  $x_a \in \mathbf{x}_a$  of values across the satellites and ground stations in the validation

data are then evaluated against the final set of predicted anomalous sequences identified according to the below rules (Hundman et al., 2018). Regions that overlap or touch after expansion are combined into a single region to account for situations where multiple anomalous regions represent a single event (Hundman et al., 2018).

1. A **true positive** is recorded if:

$$\left| e_a^{(t)} \in e_{seq} \in \mathbf{e}_{seq} : x_i^{(t)} \in x_a \right| > 0$$

for any  $x_a \in \mathbf{x}_a$ . In other words, a true positive results if any portion of a predicted sequence of anomalies falls within any true labeled sequence. Only one true positive is recorded even if portions of multiple predicted sequences fall within a labeled sequence.

2. If no predicted sequences overlap with a positively labeled sequence, a **false negative** is recorded for the labeled sequence.
3. For all predicted sequences that do not overlap a labeled anomalous region, a **false positive** is recorded.

We do not make any scoring adjustments based on how early an anomaly was detected or the distance between false positives and labeled regions (Lavin & Ahmad, 2015) for simplicity. The recall, precision, accuracy, and  $F_1$  score performance metrics on the validation set are calculated and provided according using the above (bulleted) definitions. The validation metrics reported provide a representative viewpoint of model's performance and generalizability.

The false positive mitigation strategy described in the last paragraph of Section 3.2 utilizing data across all ground stations for each satellite is used as part of our experiments. The results of utilizing this strategy are reported in precision, recall and  $F_1$  score metrics. Labels time ranges indicating TIDs produced by the false positive mitigation strategy are considered against available ground truth as described in the preceding paragraphs, and provide a reliable point of comparison between performance metrics with and without using a false positive mitigation strategy.

**Table 4**  
*Performance Metrics From Various Stages of the Experiment*

Metrics	
testing–recall	96.1%
testing–precision	93.5%
testing– $F_1$ score	94.8%
validation–recall	96.2%
validation–precision	34.7%
validation– $F_1$ score	51.0%
validation (false positive mitigation)–recall	84.6%
validation (false positive mitigation)–precision	100.0%
validation (false positive mitigation)– $F_1$ score	<b>91.7%</b>

## 4.2. Model Parameters and Evaluation

The CNN architecture and model parameters used in this experiment are described in Table 3. A thorough description of the ResNet architecture is provided in (He et al., 2015). We utilize a sequence length of  $l_s = 60$  min, but this can be adjusted and in addition to the model architecture and input image sizes. Steps are taken to minimize over-fitting, such as reducing the learning rate on loss plateauing and using early stopping. As earlier described, performance metrics were generated for both test and validation sets, with accompanying visualizations. The threshold  $T_{s,t}$  used in false positive mitigation was kept constant at 0.75 (the threshold has a range of  $[0, 1]$ ). Predictions and metrics were generated on the test set and validation set. Visualizations

were also generated illustrating the ground truth labels and data, the predictions and the model's confidence in that prediction (see Figure 4).

## 5. Discussion

As illustrated in the metrics shown in Table 4, the false positive mitigation strategy has a profound impact on reducing false positives, thus improving the precision from 34.7% to 100.0% and  $F_1$  score from 51.0% to 91.7%. While this reduces the overall recall from 96.2% to 84.6%, it reduces the false positive rate considerably which is an important consideration for a real-world system. The threshold  $T_{s,t}$  was kept constant in these experiments, but may be varied such that a higher precision or a higher recall can be achieved in future work. The 91.7%  $F_1$  score is achieved with a significantly smaller amount of events used for training compared to recent work leveraging Random Forest models (Brissaud & Astafyeva, 2022).

Model training, testing and validation took approximately 59 hr and 40 min, consuming approximately 16 gigabytes of memory across three NVIDIA Tesla P100 Graphical Processing Units. The Python (van Rossum, 1995) programming language was used, together with SciPy (Virtanen et al., 2020) and pyts (Faouzi & Janati, 2020) libraries for image encoding, Numpy (van der Walt et al., 2011) and FastAI (Howard & Gugger, 2020) for model training and inference, and the Matplotlib (Hunter, 2007) and Seaborn (Waskom et al., 2020) libraries for visualization.

The parameters utilized in the experiments were chosen based on previous work and trial and error, however a grid-based or evolutionary-algorithm-based parameter search may be an effective means of improving performance in future work. As with previous work, short-length sequences (within a few minutes) are a common source of false positives in our experiment. However, contrary to previous work many true positives also share short sequence lengths in these results despite a much longer average length (13.81 min for true positives, 2.38 min for false positives). Figure 5 shows the distributions of sequence lengths for true positives and false positives.

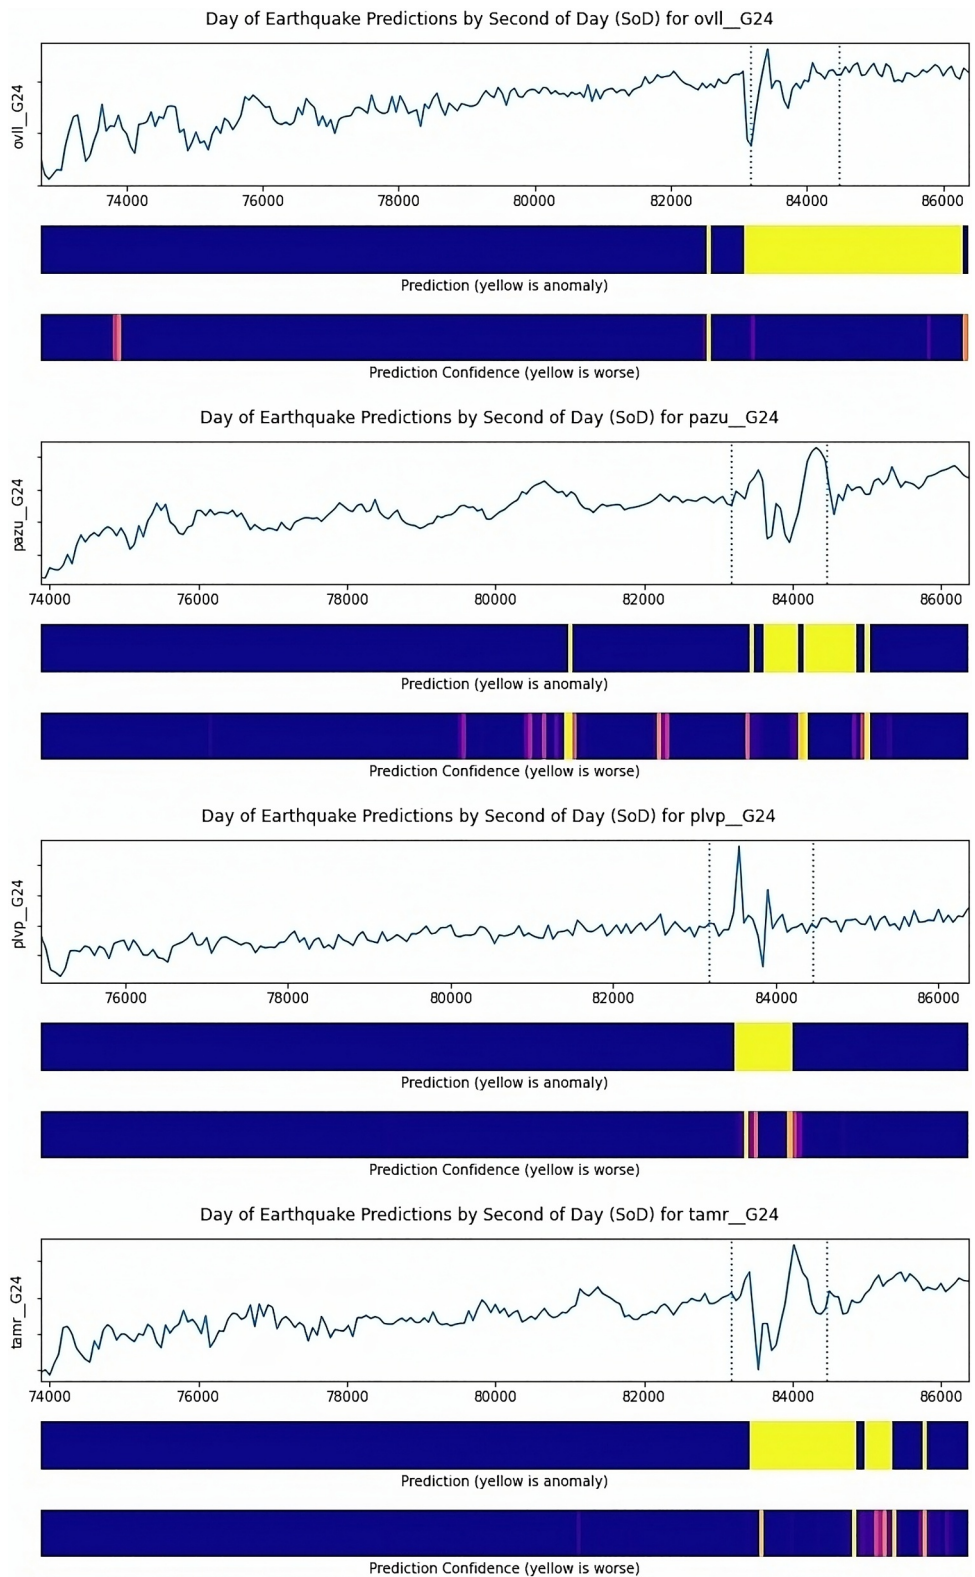
### 5.1. Impacts of Employing a False Positive Mitigation Strategy

It is clear from the metrics provided in Table 4 that employing the satellite-based false positive mitigation strategy has a profound impact on the overall performance of the approach. In Figure 6, the differences between utilizing this strategy and not mitigating false positives are shown. Areas shown in darker shades on the plot indicate periods classified as a TID in isolation, but are re-classified to normal behavior when considered in aggregate (lighter shades in the plot indicate time periods considered representative of a TID). The final TID classifications for satellites G12 and G25 are illustrated geographically in Figure 7.


The ideal threshold  $T_{s,t}$  can be selected based on the requirements of the end-users or systems accessing predictions made using this framework. For example, a cautious system may provide false positives but a significant ability to capture all relevant instances of a TID. However, a system tuned to capture all relevant instances without regard for false positives would not be practical as an alerting system. In general, a higher threshold reduces the false positives (improving precision) but decreases recall.

### 5.2. Good Data Delivers

Data from multiple tsunami-inducing earthquake events were used in this work to support the goal of training a generalized model as described in Table 1. A single set of labels is produced for each earthquake event, by a single SME. This expert utilized his or her scientific knowledge to best represent TID start and finish times, which were used to generate training data as well as perform out-of-sample validation. While this approach has shown respectable results in this work, utilizing the labels of multiple SMEs would provide additional training data and a broader perspective on what constitutes a TID—a perspective translated to the model during training. Broad consensus is better than a single opinion, and this is best captured in data by utilizing multiple SMEs for labeling. Additionally, introducing a human-in-the-loop (HIL) process would provide a means of continuously reviewing and assessing time periods predicted as TIDs by a near-real-time system by multiple SMEs. A HIL process could be considered as part of a software package employing this framework in near-real-time. Even with the performance level of 91.7%  $F_1$  score achieved using this framework for TID detection, additional performance gains are achievable by exposing the model to a larger variety of events and geographic regions

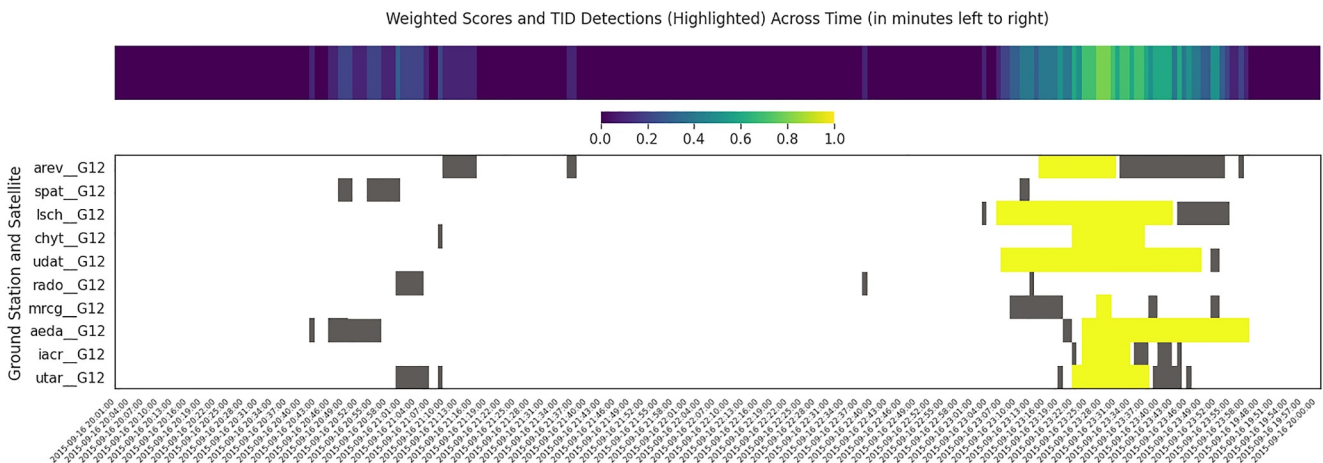


**Figure 4.** The VARION-produced  $\delta sTEC/\delta t$  estimates are visualized together with ground-truth labels, alongside classifications (anomalous or not) and classification confidence (on bottom) for the G24 satellite and *ovll*, *pazu*, *plvp*, *tamr* ground stations during the 2015 Illapel earthquake.




**Figure 5.** The distribution of predicted sequence lengths for true positives and false positives. Note that while shorter lengths of sequences are observed in the false positives, a significant amount of true positives are short length sequences.

during model training (Brissaud & Astafyeva, 2022). Data from additional tsunami-producing earthquake events such as the 2010 Sumatra (Banyaks), 2009 Fiordland or 2013 Solomon Islands earthquakes could be incorporated into model training, among others. While components of this approach such as the model hyperparameters, batch size or false positive mitigation strategy remain adjustable, large gains in model performance are likely best achieved by training the deep learning model with data from a larger number of historical events than this work. Future work should seriously consider continuing to focus on data preparation and management, from increasing the number of events to improving data labeling processes as this area will also be critical in a real-world system.



**Figure 6.** The TID classifications across all univariate time series during the 2015 Illapel earthquake and tsunami for satellite G12 are shown, one stream for each ground station. The weighted score  $F_{s,t}$  is shown on top. Shown on bottom, the highlighted time periods are those considered as TIDs following mitigation of false positives, with those in the darker shade no longer considered truly representative of a TID.



**Figure 7.** Satellite tracks (gray) with TID detections (yellow) for the G12 and G25 satellites, 2015 Illapel earthquake (epicenter shown on bottom left). The TID detections shown correspond to those remaining following pruning of false likely false positives.

### 5.3. Detecting Ionospheric Perturbation Caused by Other Phenomena

It is also important to underline the generability of the presented framework for other kind of TID detection (e.g., from volcanoes, earthquakes, large explosions and other events), creating a database of trained events. As described throughout this work, earthquakes and tsunami waves can cause acoustic and gravity wave which generate perturbations in the TEC. As such, TIDs can be an important indicator to monitoring the progression of a tsunami wave, especially in open-ocean areas not covered by buoy systems. However, ionospheric perturbation can also be caused by other natural events such as volcanic eruptions (Igarashi et al., 1994; Ravanelli et al., 2023; Rozhnoi et al., 2014) and meteorite explosions in the atmosphere (Luo et al., 2020; Yang et al., 2014). Additionally, ionospheric perturbation can be caused by man-made hazards such as large explosions, rocket and missile launches, mine blast (Astafyeva, 2019; Calais et al., 1998; Savastano et al., 2019). While this framework was developed with the intent of detecting earthquake and tsunami-induced TIDs, other data sets could be utilized with this approach to detect TIDs more broadly or detect ionospheric perturbation from different sources (creating a multi-classification problem). Certainly, it is important to underline that detecting various kind of TIDs involves discrimination, and discrimination necessitates understanding features. To this point, the creation of a database with several TIDs scenario could be effective. Finally, while challenges will be present with data curation and labeling, a system that can identify and differentiate ionospheric perturbation in near-real-time would represent a significant new capability in continuously monitoring the ionosphere for various signs of Earth activity.

## 6. Conclusions

GNSS-TEC observations are routinely used to detect earthquake and tsunami induced TIDs. In this work, we presented an innovative approach to perform real-time TID detection, on the basis of computer vision and deep learning. Being aware of the ionospheric condition variability, we used a multi-event approach, focusing on the Pacific area, to train and to test our framework. Our model is trained using data from multiple events and subsequently tested on a completely independent event. This aspect underscores the high generalizability of our approach.

The presented framework is based on the joint use of GADFs and of CNNs, achieving 91% of F1 score, 84% of recall and 100% of precision. The False Positive Mitigation Strategy—an approach for identifying TIDs, based on the likelihood of a TID at each time step—proves highly effective as it enhances precision, offering a means to ensure the robust performance of the system as it scales up, incorporating more data for model training. Further improvements in performance can be attained by incorporating data from additional events and using data at higher rate (1 Hz). Leveraging labels from multiple

SMEs in conjunction with a human-in-the-loop (HIL) process would offer a more nuanced perspective on defining and identifying TIDs.

Furthermore, the presented framework should be also able to detect ionospheric perturbations induced by moderate  $M_w$  earthquakes as long they can produce a signature in the ionosphere. Indeed, it is known that the characteristics of the ionospheric TEC perturbations depend of seismic and non-seismic factors (such as the geometry of GNSS sounding or the background ionospheric electron density (Astafyeva, 2019; Meng et al., 2019)). Nonetheless, as outlook of this work, it could be interesting to test the proposed framework on ionospheric perturbations induced by moderate earthquakes (ranging between  $M_w$  6.5 and  $M_w$  7.5).

Although this framework was initially designed for identifying earthquake and tsunami induced TIDs, it has broader applicability. This approach can be extended to detect ionospheric perturbations stemming from various sources, thereby addressing a multi-classification challenge, or identifying time series anomalies in other geodetic applications such as glacier and ice sheet monitoring.

In conclusion, this work represents a first feasibility demonstration of applying deep learning to real-time TID detection with the final aim to augment tsunami early warning system. Finally, this work is perfectly aligned with the objectives and research interests of the AI for GNSS Remote Sensing study group of the GGOS Focus Area on Artificial Intelligence for Geodesy (AI4G) (Global Geodetic Observing System (GGOS), 2023), representing a relevant and valuable contribution to developments of AI-driven techniques for the TEWS enhancement through real-time GNSS Ionospheric Seismology.

### Data Availability Statement

Global Navigation Satellite System data for the 2012 Haida Gwaii earthquake and the 2010 Maule earthquake are available from UNAVCO Data Archive Interface (UNAVCO, 2021) at <https://www.unavco.org/data-gps-gnss/data-access-methods/gnss-data-access-notebooks/gnss-permanent-station-data-access-notebook-embed.html>. Thirty second GNSS data for the 2011 Tōhoku earthquake are available under request from GEONET network (of Japan (GSI), 2021) from Geospatial Information Authority of Japan website [https://www.gsi.go.jp/ENGLISH/geonet\\_english.html](https://www.gsi.go.jp/ENGLISH/geonet_english.html). Global Navigation Satellite System data for the 2015 Illapel earthquake are available from Centro Sismológico Nacional, Universidad de Chile (CSN) (Centro Sismológico Nacional, 2021) <http://gps.csn.uchile.cl/data/2015/>, accessed on 2021-10-18. Source code is available at: <https://doi.org/10.5281/zenodo.12571500> (Constantinou & Liu, 2024).

### References

- Astafyeva, E. (2019). Ionospheric detection of natural hazards. *Reviews of Geophysics*, 57(4), 1265–1288. <https://doi.org/10.1029/2019rg000668>
- Astafyeva, E., Lognonné, P., & Rolland, L. (2011). First ionospheric images of the seismic fault slip on the example of the Tohoku-Oki earthquake. *Geophysical Research Letters*, 38(22). <https://doi.org/10.1029/2011gl049623>
- Azeem, I., Vadas, S. L., Crowley, G., & Makela, J. J. (2017). Traveling ionospheric disturbances over the United States induced by gravity waves from the 2011 Tohoku tsunami and comparison with gravity wave dissipative theory. *Journal of Geophysical Research: Space Physics*, 122(3), 3430–3447. <https://doi.org/10.1002/2016ja023659>
- Bontemps, L., Cao, V. L., McDermott, J., & Le-Khac, N.-A. (2017). Collective anomaly detection based on long short term memory recurrent neural network. In *Future data and security engineering: Third International Conference, FDSE 2016, can Tho City, Vietnam, November 23-25, 2016, Proceedings 3* (pp. 141–152). Springer International Publishing.
- 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>
- Calais, E., Bernard Minster, J., Hofton, M., & Hedlin, M. (1998). Ionospheric signature of surface mine blasts from global positioning system measurements. *Geophysical Journal International*, 132(1), 191–202. <https://doi.org/10.1046/j.1365-246x.1998.00438.x>
- Center, I. T. I. (2021a). 11 March 2011, MW 9.1, near the east coast of Honshu Japan Tsunami. Retrieved from [http://itic.ioc-unesco.org/index.php?option=com\\_content&view=category&id=2362&Itemid=2365](http://itic.ioc-unesco.org/index.php?option=com_content&view=category&id=2362&Itemid=2365)
- Center, I. T. I. (2021b). 2010 February 27, MW 8.8 south-central Chile. Retrieved from [http://itic.ioc-unesco.org/images/stories/list\\_of\\_tsunamis/2010/27feb2010\\_Chile/995\\_027%20DOC%20PTWS%20book\\_2010Chile.pdf](http://itic.ioc-unesco.org/images/stories/list_of_tsunamis/2010/27feb2010_Chile/995_027%20DOC%20PTWS%20book_2010Chile.pdf)
- Center, I. T. I. (2021c). PTWC issues warning for m7.7 Haida Gwaii, British Columbia earthquake. Retrieved from [http://itic.ioc-unesco.org/index.php?option=com\\_content&view=article&id=1839&Itemid=2986](http://itic.ioc-unesco.org/index.php?option=com_content&view=article&id=1839&Itemid=2986)
- Centro Sismológico Nacional, U. d. C. (2021). Data-2015 [Dataset]. <http://gps.csn.uchile.cl/data/2015/>
- Chauhan, S., & Vig, L. (2015). Anomaly detection in ECG time signals via deep long short-term memory networks. In *2015 IEEE international conference on data science and advanced analytics (DSAA)*. IEEE. <https://doi.org/10.1109/dsaa.2015.7344872>
- Chen, J.-H., & Tsai, Y.-C. (2019). Encoding candlesticks as images for patterns classification using convolutional neural networks. *Financial Innovation*, 6(1), 26. <https://doi.org/10.1186/s40854-020-00187-0>
- Ciresan, D., Meier, U., & Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. In *2012 IEEE conference on computer vision and pattern recognition*. IEEE. <https://doi.org/10.1109/cvpr.2012.6248110>
- Constantinou, V., & Liu, H. (2024). vc1492a/tidd: 0.1.2 (0.1.2) [Software]. Zenodo. <https://doi.org/10.5281/zenodo.12571500>
- Constantinou, V., Ravanelli, M., Liu, H., & Bortnik, J. (2023a). A deep learning approach for detection of internal gravity waves in earth's ionosphere. In *IGARSS 2023-2023 IEEE international geoscience and remote sensing symposium* (pp. 1178–1181).
- Constantinou, V., Ravanelli, M., Liu, H., & Bortnik, J. (2023b). Deep learning driven detection of tsunami related internal gravitywaves: A path towards open-ocean natural hazards detection. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV) workshops* (pp. 3748–3753).
- Daniels, F. B. (1952). Acoustical energy generated by the ocean waves. *Journal of the Acoustical Society of America*, 24(1), 83. <https://doi.org/10.1121/1.1906855>
- Faouzi, J., & Janati, H. (2020). pyts: A python package for time series classification. *Journal of Machine Learning Research*, 21(46), 1–6. Retrieved from <http://jmlr.org/papers/v21/19-763.html>
- Fortunato, M., Ravanelli, M., & Mazzoni, A. (2019). Real-time geophysical applications with android GNSS raw measurements. *Remote Sensing*, 11(18), 2113. <https://doi.org/10.3390/rs11182113>

### Acknowledgments

M.R. was supported by the fellowship granted by AXA Research Fund and the Intergovernmental Oceanographic Commission of UNESCO (IOC-UNESCO) on Coastal Livelihood Grant (B83C23002490007). Open access publishing facilitated by Università degli Studi di Roma La Sapienza, as part of the Wiley - CRUI-CARE agreement.



- Fratarcangeli, F., Ravanelli, M., Mazzoni, A., Colosimo, G., Benedetti, E., Branzanti, M., et al. (2018). The variometric approach to real-time high-frequency geodesy. *RENDICONTI LINCEI. SCIENZE FISICHE E NATURALI*, 29(S1), 95–108. <https://doi.org/10.1007/s12210-018-0708-5>
- Fritz, H. M., Petroff, C. M., Catalán, P. A., Cienfuegos, R., Winckler, P., Kalligeris, N., et al. (2011). Field survey of the 27 February 2010 Chile tsunami. *Pure and Applied Geophysics*, 168(11), 1989–2010. <https://doi.org/10.1007/s00024-011-0283-5>
- 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>
- Galvan, D. A., Komjathy, A., Hickey, M. P., & Mannucci, A. J. (2011). The 2009 Samoa and 2010 Chile tsunamis as observed in the ionosphere using GPS total electron content. *Journal of Geophysical Research*, 116(A6). <https://doi.org/10.1029/2010ja016204>
- Geospatial Information Authority of Japan. (2021). Geonet (GNSS earth observation network system) [Dataset]. [https://www.gsi.go.jp/ENGLISH/geonet\\_english.html](https://www.gsi.go.jp/ENGLISH/geonet_english.html)
- Global Geodetic Observing System (GGOS). (2023). AI for GNSS Remote sensing. <https://ggos.org/about/org/fa/ai-for-geodesy/gnss-remote-sensing/>
- Grawe, M. A., & Makela, J. J. (2015). The ionospheric responses to the 2011 Tohoku, 2012 Haida Gwaii, and 2010 Chile tsunamis: Effects of tsunami orientation and observation geometry. *Earth and Space Science*, 2(11), 472–483. <https://doi.org/10.1002/2015ea000132>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep residual learning for image recognition. *CoRR*. Retrieved from <http://arxiv.org/abs/1512.03385>
- Hines, C. O. (1972). Gravity waves in the atmosphere. *Nature*, 239(5367), 73–78. <https://doi.org/10.1038/239073a0>
- Hines, C. O. (1974). Internal atmospheric gravity waves at ionospheric heights. In *The upper atmosphere in motion* (pp. 248–328). American Geophysical Union. <https://doi.org/10.1029/gm018p0248>
- Hohensinn, R., Aichinger-Rosenberger, M., Wareyka-Glaner, M. F., & Ravanelli, M. (2024). Natural-hazard monitoring with global navigation satellite systems (GNSS). In *Advances in Geophysics, Space Geodesy for Environmental Monitoring*. (Vol. 65). Elsevier.
- Hong, Y., Martinez, J. J. F., & Fajardo, A. C. (2020). Day-ahead solar irradiation forecasting utilizing Gramian angular field and convolutional long short-term memory. *IEEE Access*, 8, 18741–18753. <https://doi.org/10.1109/ACCESS.2020.2967900>
- Howard, J., & Gugger, S. (2020). Fastai: A layered API for deep learning. *Information*, 11(2), 108. <https://doi.org/10.3390/info11020108>
- Huang, G., Liu, Z., & Weinberger, K. Q. (2016). Densely connected convolutional networks. *CoRR*. Retrieved from <http://arxiv.org/abs/1608.06993>
- Hundman, K., Constantinou, V., Laporte, C., Colwell, I., & Soderstrom, T. (2018). Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. ACM. <https://doi.org/10.1145/3219819.3219845>
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>
- 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)
- International union of Geodesy and Geophysics. (2015). International union of geodesy and geophysics resolutions. Retrieved from [https://iag.dgfi.tum.de/fileadmin/LAG-docs/IUGG\\_Resolutions\\_2015.pdf](https://iag.dgfi.tum.de/fileadmin/LAG-docs/IUGG_Resolutions_2015.pdf)
- Istituto Nazionale di Geofisica e Vulcanologia. (2021). Tsumemo: 27.02.2010, the Maule tsunami, Chile. Retrieved from <https://www.ingv.it/cat/en/media-and-documents/news-en/241-tsu-memo-27-02-2010-the-maule-tsunami-chile>
- Kamogawa, M., Orihara, Y., Tsurudome, C., Tomida, Y., Kanaya, T., Ikeda, D., et al. (2016). A possible space-based tsunami early warning system using observations of the tsunami ionospheric hole. *Scientific Reports*, 6(1), 1–7. <https://doi.org/10.1038/srep37989>
- Kundu, B., Senapati, B., Matsushita, A., & Heki, K. (2021). Atmospheric wave energy of the 2020 August 4 explosion in Beirut, Lebanon, from ionospheric disturbances. *Scientific Reports*, 11(1), 2793. <https://doi.org/10.1038/s41598-021-82355-5>
- LaBrecque, J., Rundle, J., & Bawden, G. (2019). Global navigation satellite system enhancement for tsunami early warning systems. Global Assessment Report on Disaster Risk Reduction.
- Lavin, A., & Ahmad, S. (2015). Evaluating real-time anomaly detection algorithms—The Numenta anomaly benchmark. In *Machine learning and applications (ICMLA), 2015 IEEE 14th international conference on* (pp. 38–44).
- Luo, Y., Yao, Y., & Shan, L. (2020). Analysis of ionospheric disturbances caused by the 2018 Bering Sea meteor explosion based on GPS observations. *Sensors*, 20(11), 3201. <https://doi.org/10.3390/s20113201>
- Malhotra, P., Lovekesh, V., Shroff, G., & Argarwal, P. (2015). Long short term memory networks for anomaly detection in time series. In *Proceedings of the European symposium on artificial neural networks (ESANN), computational intelligence and machine learning*.
- Malhotra, P., Ramakrishnan, A., Anand, G., Vig, L., Agarwal, P., & Shroff, G. (2016). LSTM-based encoder-decoder for multi-sensor anomaly detection. *CoRR*. [abs/1607.00148](https://arxiv.org/abs/1607.00148).
- Manta, F., Occhipinti, G., Feng, L., & Hill, E. M. (2020). Rapid identification of tsunami-genic earthquakes using GNSS ionospheric sounding. *Scientific Reports*, 10(1), 1–10. <https://doi.org/10.1038/s41598-020-68097-w>
- Martire, L., Krishnamoorthy, S., Vergados, P., Romans, L. J., Szilágyi, B., Meng, X., et al. (2023). 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>
- Meng, X., Ravanelli, M., Komjathy, A., & Verkhoglyadova, O. P. (2022). On the north-south asymmetry of co-seismic ionospheric disturbances during the 16 September 2015 Illapel m8. 3 earthquake. *Geophysical Research Letters*, 49(8), e2022GL098090. <https://doi.org/10.1029/2022gl098090>
- 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>
- Munir, M., Siddiqui, S. A., Dengel, A., & Ahmed, S. (2019). DeepAnT: A deep learning approach for unsupervised anomaly detection in time series. *IEEE Access*, 7, 1991–2005. <https://doi.org/10.1109/access.2018.2886457>
- Nanduri, A., & Sherry, L. (2016). Anomaly detection in aircraft data using recurrent neural networks (RNN). 2016 Integrated Communications Navigation and Surveillance (ICNS), 5C2-1-5C2-8.
- Occhipinti, G. (2015). The seismology of the planet mongo. In *Geophysical monograph series* (pp. 169–182). John Wiley & Sons, Inc. <https://doi.org/10.1002/9781118888865.ch9>
- Peltier, W. R., & Hines, C. O. (1976). On the possible detection of tsunamis by a monitoring of the ionosphere. *Journal of Geophysical Research*, 81(12), 1995–2000. <https://doi.org/10.1029/jc081i012p01995>
- Ranzato, M. A., Boureau, Y. L., & Cun, Y. (2007). Sparse feature learning for deep belief networks. *Advances in Neural Information Processing Systems*, 20.

- Ravanelli, M., Astafyeva, E., Munaibari, E., Rolland, L., & Mikesell, T. (2023). Ocean-ionosphere disturbances due to the 15 January 2022 Hunga-Tonga Hunga-Ha'apai eruption. *Geophysical Research Letters*, *50*(10), e2022GL101465. <https://doi.org/10.1029/2022gl101465>
- Ravanelli, M., Crespi, M., & Foster, J. (2020). Tides detection from ship-based GNSS receiver: First test on 2010 Maule tsunami. In *IGARSS 2020-2020 IEEE international geoscience and remote sensing symposium* (pp. 6846–6849).
- 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), 1–12. <https://doi.org/10.1038/s41598-021-82532-6>
- Rolland, L. M., Lognonné, P., & Munkane, H. (2011). Detection and modeling of Rayleigh wave induced patterns in the ionosphere. *Journal of Geophysical Research*, *116*(A5). <https://doi.org/10.1029/2010JA016060>
- Rozhnoi, A., Hayakawa, M., Solovieva, M., Hobar, Y., & Fedun, V. (2014). Ionospheric effects of the mt. Kirishima volcanic eruption as seen from subionospheric VLF observations. *Journal of Atmospheric and Solar-Terrestrial Physics*, *107*, 54–59. <https://doi.org/10.1016/j.jastp.2013.10.014>
- Satake, K., & Heidarzadeh, M. (2017). A review of source models of the 2015 Illapel, Chile earthquake and insights from tsunami data. *Pure and Applied Geophysics*, *174*(1), 1–9. [https://doi.org/10.1007/978-3-319-57822-4\\_1](https://doi.org/10.1007/978-3-319-57822-4_1)
- Savastano, G., Komjathy, A., Shume, E., Vergados, P., Ravanelli, M., Verkhoglyadova, O., et al. (2019). Advantages of geostationary satellites for ionospheric anomaly studies: Ionospheric plasma depletion following a rocket launch. *Remote Sensing*, *11*(14), 1734. <https://doi.org/10.3390/rs11141734>
- 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*(1), 46607. <https://doi.org/10.1038/srep46607>
- Shrivastava, M. N., Maurya, A. K., Gonzalez, G., Sunil, P. S., Gonzalez, J., Salazar, P., & Aranguiz, R. (2021). Tsunami detection by GPS-derived ionospheric total electron content. *Scientific Reports*, *11*(1), 12978. <https://doi.org/10.1038/s41598-021-92479-3>
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International conference on learning representations*.
- Soulé, B. (2014). Post-crisis analysis of an ineffective tsunami alert: The 2010 earthquake in maule, Chile. *Disasters*, *38*(2), 375–397. <https://doi.org/10.1111/disa.12045>
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., et al. (2015). Going deeper with convolutions. In *2015 IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 1–9). <https://doi.org/10.1109/CVPR.2015.7298594>
- Taylor, A., Leblanc, S., & Japkowicz, N. (2016). Anomaly detection in automobile control network data with long short-term memory networks. In *2016 IEEE international conference on data science and advanced analytics (DSAA)*. IEEE. <https://doi.org/10.1109/dsaa.2016.20>
- UNAVCO. (2021). Data archive interface - UNAVCO [Dataset]. <https://www.unavco.org/data/gps-gnss/data-access-methods/gnss-data-access-notebooks/gnss-permanent-station-data-access-notebook-embed.html>
- USGS. (2020). M 7.8 - 206 km SW of Prince Rupert, Canada. <https://earthquake.usgs.gov/earthquakes/eventpage/usp000juhz/executive>
- USGS. (2021a). M 8.3 - 48 km w of Illapel, Chile. Retrieved from <https://earthquake.usgs.gov/earthquakes/eventpage/us20003k7a/executive>
- USGS. (2021b). M 8.8 - 36 km WNW of Quirihue, Chile. Retrieved from [https://earthquake.usgs.gov/earthquakes/eventpage/official20100227063411530\\_30/executive](https://earthquake.usgs.gov/earthquakes/eventpage/official20100227063411530_30/executive)
- USGS. (2021c). M 9.1 - 2011 great Tohoku earthquake, Japan. Retrieved from [https://earthquake.usgs.gov/earthquakes/eventpage/official20110311054624120\\_30/executive](https://earthquake.usgs.gov/earthquakes/eventpage/official20110311054624120_30/executive)
- van der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The Numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, *13*(2), 22–30. <https://doi.org/10.1109/mcse.2011.37>
- van Rossum, G. (1995). Python tutorial, technical report cs-r9526. (Tech. Rep.).
- Vergados, P., Komjathy, A., & Meng, X. (2020). GNSS observation for detection, monitoring, and forecasting natural and man-made hazardous events. Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications, *1*, 939–969.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Wang, Z., & Oates, T. (2015). Imaging time-series to improve classification and imputation. *CoRR*. Retrieved from <http://arxiv.org/abs/1506.00327>
- Waskom, M., Botvinnik, O., Ostblom, J., Lukauskas, S., Hobson, P., MaozGelbart, et al. (2020). *mwaskom/seaborn: v0.10.0 (January 2020)*. Zenodo. <https://doi.org/10.5281/zenodo.3629446>
- Waszczak, A., Prince, T. A., Laher, R., Masci, F., Bue, B., Rebbapragada, U., et al. (2017). Small near-earth asteroids in the palomar transient factory survey: A real-time streak-detection system. *Publications of the Astronomical Society of the Pacific*, *129*(973), 034402. <https://doi.org/10.1088/1538-3873/129/973/034402>
- Wickramaratne, S. D., & Mahmud, M. S. (2021). A deep learning based ternary task classification system using Gramian angular summation field in FNIRS neuroimaging data.
- Xu, H., Li, J., Yuan, H., Liu, Q., Fan, S., Li, T., & Sun, X. (2020). Human activity recognition based on Gramian angular field and deep convolutional neural network. *IEEE Access*, *8*, 199393–199405. <https://doi.org/10.1109/ACCESS.2020.3032699>
- Yang, Y.-M., Komjathy, A., Langley, R. B., Vergados, P., Butala, M. D., & Mannucci, A. J. (2014). The 2013 chelyabinsk meteor ionospheric impact studied using GPS measurements. *Radio Science*, *49*(5), 341–350. <https://doi.org/10.1002/2013rs005344>
- Yue, H., Lay, T., Rivera, L., An, C., Vigny, C., Tong, X., & Baez Soto, J. C. (2014). Localized fault slip to the trench in the 2010 Maule, Chile  $m_w = 8.8$  earthquake from joint inversion of high-rate GPS, teleseismic body waves, INSAR, campaign GPS, and tsunami observations. *Journal of Geophysical Research: Solid Earth*, *119*(10), 7786–7804. <https://doi.org/10.1002/2014jb011340>