Denoising Daily Displacement GNSS-Time series using Deep Neural Networks In a Near Real-Time Framing: a Single-Station Method

Mastella G.[1], Bedford J.[2], Corbi F.[3], Funiciello F.[4]

[1] Sapienza University of Rome, Earth Sciences, Italy.

[2] Institut für Geologie, Mineralogie und Geophysik, Ruhr-Universität Bochum, Bochum, 44801, Germany.

[3] Istituto di Geologia Ambientale e Geoingegneria – CNR c/o Dipartimento di Scienze della Terra, Sapienza Università di Roma, Rome, Italy.

[4] Università "Roma TRE,", Rome, Italy, Dip Scienze, Laboratory of Experimental Tectonics.

***Corresponding author: Giacomo Mastella (giacomo.mastella@uniroma1.it)**

This is a non-peer-reviewed manuscript being submitted to Geophysical International Journal (GJI)

Abstract

Recent ground observations from Global Navigation Satellite Systems (GNSS) displacement time series have provided compelling evidence that the motion of tectonic plates is ubiquitously non-steady-state. In some cases, these anomalous transient motions have been identified as potential precursors occurring months, days, or hours before large-magnitude earthquakes. However, effectively detecting these signals in daily geodetic time-series at the earliest opportunity remains challenging due to the levels of high-frequency noise. Currently, there is a lack of established methodologies to reduce this noise in near-real-time thereby hindering our ability to timely monitor tectonic transient motions. Precursors are typically modeled retrospectively, and the use of geodetic data for seismic hazard surveillance remains limited. To address this limitation, this study demonstrates an approach to model high-frequency noise in daily GNSS displacement time-series, with the removal of this modeled noise allowing for tectonic transients to be potentially more clearly identified. Using Deep Neural Networks (DNNs), we develop a denoising approach that removes noise from GNSS displacement time-series on a station-by-station basis. To more effectively train our DNN models, we generate a comprehensive and diverse dataset by combining synthetic trajectories with synthetic noise time-series created using Generative Adversarial Networks (GAN). To train the GAN, we use noise time series extracted from ~5000 GNSS displacement time series distributed globally. Validating our approach with real data confirms its capability to significantly reduce the high-frequency noise that characterizes GNSS time-series. The flexibility of the method allows for near-real-time noise removal (with a latency of a few days), opening up the possibility of detecting and modeling small tectonic transients in a timely fashion. By introducing this novel approach, we present exciting opportunities to advance the geodetic surveillance of tectonic motions and usher in a new era of improved monitoring of seismic activity.

Plain Language Summary

Time-series of daily displacements obtained from global navigation satellite systems (GNSS) involve precise position measurements of the Earth's surface acquired by collecting position and time data from satellites. These data are crucial for monitoring the deformation of our planet, particularly for studying tectonic plate motion. However, the presence of unwanted noise poses challenges to the characterization of these data. This includes detecting slow earthquakes, which can act as potential precursors to large earthquakes but often go unnoticed by humans. Currently, there are no established methodologies for real-time noise reduction aimed at the geodetic monitoring of transient tectonic mo-

tions, limiting scientists' ability to fully harness the potential of these data in earthquake hazard surveillance. To tackle this challenge, we propose a data-driven method that removes noise from the GNSS data by leveraging Deep Neural Networks. Our method effectively removes high-frequency noise commonly found in GNSS data on a global scale. Notably, it operates with near real-time capabilities, enabling prompt detection of small-scale Earth surface movements as they occur in the daily time series. This advancement opens up exciting possibilities for enhancing our understanding of tectonic activity, and potentially unlocking new avenues for predicting future seismicity.

1. Introduction

There is a growing recognition within the earthquake research community that transient tectonic motions - accelerations from the trajectory expected from long-term plate motions and cyclic seasonal motion - are widespread. The occurrence of slow-slip events, which are the most commonly observed type of transients, has been documented in numerous active fault systems around the world [Jolivet and Frank, 2016]. These events are currently being intensely investigated for their potential link to earthquakes in both natural settings [e.g. Wallace et al., 2018; Rolandone et al., 2018; Kato and Ben-Zion, 2020; Cruz-Atienza et al., 2021], numerical models [e.g. Schaal and Lapusta, 2019; Meng and Duan, 2022], and laboratory studies [e.g. McLaskey 2019; Cebry et al., 2022; Marty et al., 2023]. Recent observations have shown that a slow aseismic phase of slip on the fault can anticipate large earthquakes over temporal scales ranging from minutes to years [Bouchon et al., 2011; Bletery and Nocquet, 2023; Kato et al., 2012; Bouchon et al., 2013; Ruiz et al., 2017; Caballero et al., 2021; Ruiz et al., 2014; Schurr et al., 2014; Bouchon et al., 2016; Radiguet et al., 2016; Socquet et al., 2017; Bedford et al., 2020; Mavrommatis et al., 2014]. This growing understanding suggests that transient motion may provide valuable precursor information on the likelihood of impending large earthquakes [Obara and Kato, 2016].

The detection and characterization of transient tectonic movements primarily rely on ground deformation data derived from daily displacement time-series obtained from Global Navigation Satellite Systems (GNSS). GNSS time-series exhibit complex signals that operate over various time-scales, and when properly modeled, can provide valuable insights into the dynamic behavior, deformation, and movement of Earth's crust. Geodesists typically employ trajectory fitting techniques to extract these multiscale signals, which involve mathematical approximations to interpret the observed data [see Bevis and Brown, 2014]. The residuals obtained after time series decomposition represent the remaining scatter, commonly referred to as noise, without specifying its specific nature or underlying causes. This unquantified scatter acts as a significant impediment, limiting the accurate detection and characterization of tectonic signals. Such noise in GNSS time-series can arise from several different noise sources, such as atmospheric loading or tropospheric delay [e.g., Tregoning and Watson, 2009], mis-modeling of orbits [Griffiths and Ray, 2009; Steigenberger et al., 2009], monument instability [King and Williams, 2009], data processing choices (e.g. jitter shown in Bevis and Brown [2014]) or near-field multipath [e.g., King and Watson, 2010]. Studies into frequency content suggest this non-tectonic, non-seasonal scatter as a combination of colored and white-noise [e.g., Dmitrieva et al., 2015; Langbein, 2008]. Depending on the processing strategy, the noise present in GNSS displacement time-series can exhibit various degrees of spatial correlation. This correlation is referred to as common-mode signal (CMS) and is thought to be the result of large-scale processes such as ocean and atmosphere dynamics, mass changes [e.g. Fu and Freymueller, 2012; Liu et al., 2015]. Some of the CMS is also thought to result from processing biases such as errors in the estimation of orbits for precise point positioning mode (PPP) processing. Such errors can be classed as Common Mode Errors (CME). Differentiating between CMS and CME is not always straightforward, often leading to the removal of spatially correlated noise without undergoing interpretation through common-mode filters (CMFs) [Wdowinski et al., 1997; Kreemer, and Blewitt, 2021]. CMFs utilize the correlation of noise across GNSS stations in close spatial proximity, enabling spatial filtering to unveil local tectonic transients [e.g. Smith et al., 2004]. However, CMFs face limitations when dealing with sparsely distributed GNSS stations and may inadvertently introduce additional noise if the original trajectory model fits are suboptimal. Numerous other approaches have been proposed to mitigate GNSS noise. Common methods to spatially filter GNSS noise are Karhunen–Loeve expansion (KLE) [Dong et al. 2006], Principal Component Analysis (PCA) [Jackson and Chen 2004; He et al. 2015], which is very effective in eliminating long-period noise [e.g., He et al., 2015] and variational Bayesian Independent Component Analysis (vbICA) [e.g. Gualandi et al., 2016]. Additionally, data-driven methods combining wavelet transform (WT) [Satirapod and Rizos, 2005; Azarbad and Mosavi, 2014] with PCA and vbICA techniques have shown promise in reducing high-frequency noise [Li et al., 2018], albeit with increased computational time due to empirical mode decomposition.

In some geological studies, even simpler frequency-dependent filters are used to illuminate tectonically-related processes [e.g. Frank, 2016; Rousset et al., 2019; Nuyen and Schmidt, 2021]. Examples of such filters include sidereal filtering, moving averages, and match-filtering. Such simpler filters act as

low-pass filters, which restricts their effectiveness to frequencies below the chosen cutoff frequency. Furthermore, the filtering process introduces inherent temporal delays, making them unsuitable for real-time monitoring of daily sampled GNSS displacement time-series. In general, there is a lack of well-established methodologies to reduce the noise in daily GNSS displacement data on a near-real-time basis, which inhibits the real-time monitoring of anomalous tectonic motions. This shortcoming in our suite of methods prevents us from tracking precursors that are thought to appear before some large earthquakes, and therefore hampers our ability to fully exploit the potential of geodetic data for seismic hazard surveillance. This is not the case for seismic data, for which a near real-time automated detection and characterization of tectonic plates' deformation is possible for almost all earthquakes with a body wave magnitude (m_b) greater than 4 [Gaebler et al, 2021].

Deep Learning (DL) has recently emerged as a very useful tool in earthquake monitoring [Mousavi and Beroza, 2022; Lin et al., 2021; Thomas et al., 2021; Dittmann et al., 2022]. Thomas et al., 2023 introduced a novel Deep Learning method for denoising High-rate-GNSS (HR-GNSS) records, inspired by the denoising strategy developed by Zhu et al. [2019] for seismic data. This method decomposes the input HR-GNSS time-series into signal and noise by predicting their respective masks in the frequency domain, significantly improving the Signal-to-Noise Ratio (SNR). To train their Deep Learning models, Thomas et al. [2023] leveraged a synthetic dataset, as is often done in machine learning (ML) problems in Earth Science. The dataset consisted of earthquake-induced ground displacements in HR-GNSS records generated synthetically with the FakeQuakes code [Melgar et al., 2016], along with noise recordings. The noise waveforms were selected from real HR-GNSS recordings unaffected by earthquakes above magnitude 4.3. Consequently, the network was trained to reconstruct the kinematic rupture of synthetic earthquakes registered by HR-GNSS waveforms. However, applying this approach to daily displacement GNSS time-series is not feasible, because differently from HR records, daily time-series cannot capture the kinematic ground displacements caused by dynamic seismic waves. Instead, daily time series are more suitable for capturing longer temporal processes such as variations in interseismic plate motion [e.g. Materna et al., 2019; Yuzariyadi, M. & Heki, 2021], subsurface hydrological dynamics [e.g. D'Agostino et al, 2018] and mostly slow tectonic transients [e.g. Riel et al., 2014].

In this study, we present a DL-based method specifically designed to reduce noise in GNSS daily displacement time-series and investigate how close to real-time this method can provide accurate information about tectonic motion. The method works by predicting the high-frequency scatter present in the data on a station by station basis (i.e. spatially correlated information is not considered). The model for predicting the high-frequency scatter relies on the output of a model to predict the presence of steps in the time series. To optimize the training process of our DNN models, we first build a comprehensive and diverse synthetic dataset of GNSS displacement time-series. As described in Section 2, this is achieved by integrating smooth synthetic trajectories of a parameterized trajectory model with artificial noise generated by a Generative Adversarial Network (GAN) that has been trained on more than 5000 examples of time-series of what we define to be "real" noise. An important benefit of the synthetic noise generator is our ability to generate a potentially unlimited number of training examples for training the step detection and the scatter predicting networks. We firstly provide a detailed description of the Denoising models employed, while we later present the results obtained when applying the denoising method to real time series. We quantitatively compare the performance of our method with commonly used filters in GNSS data such as the CMF in Section 5, highlighting the improvements achievable with our DL-based denoising procedure. Furthermore, we present specific real case studies where we apply our method. By showcasing its efficacy in these practical applications, we emphasize the potential of our DL-based method for enhancing the accuracy and reliability of GNSS displacement time-series analysis in near-real-time.

2. Data

To train our DL models, we create synthetic daily displacement time-series. These synthetic displacement time series combine smooth trajectories and realistic, GAN-generated artificial noise. We first describe the dataset of geodetic data from real stations used for training the GAN and then explain the process of synthetic time-series generation.

2.1. Real Data

Our dataset of real GNSS records is composed of 8364 three directional components, daily displacement time-series. This dataset comprises the PPP solutions of GNSS time-series obtained from the Nevada Geodetic Laboratory (NGL) in the IGS14 reference frame selected based on specific spatial and temporal constraints. Firstly, we consider stations located in the latitude range between 30°N and

150°N at an altitude of less than 1500 m to prevent any influence caused by snow and ice. Second, we select only time-series with at least 4 years of continuous data, a maximum number of missing days lower than 146 (the 10% of 4 years) and a maximum number of consecutive data gaps of 5 five days. The whole dataset is randomly split into training and testing sets (60-40% respectively) based on the number of stations (**Fig.1**). The training set is used to train the GAN models employed to generate synthetic noise, while the remaining data are used in the testing phase of the denoising models (described in **section 3**).

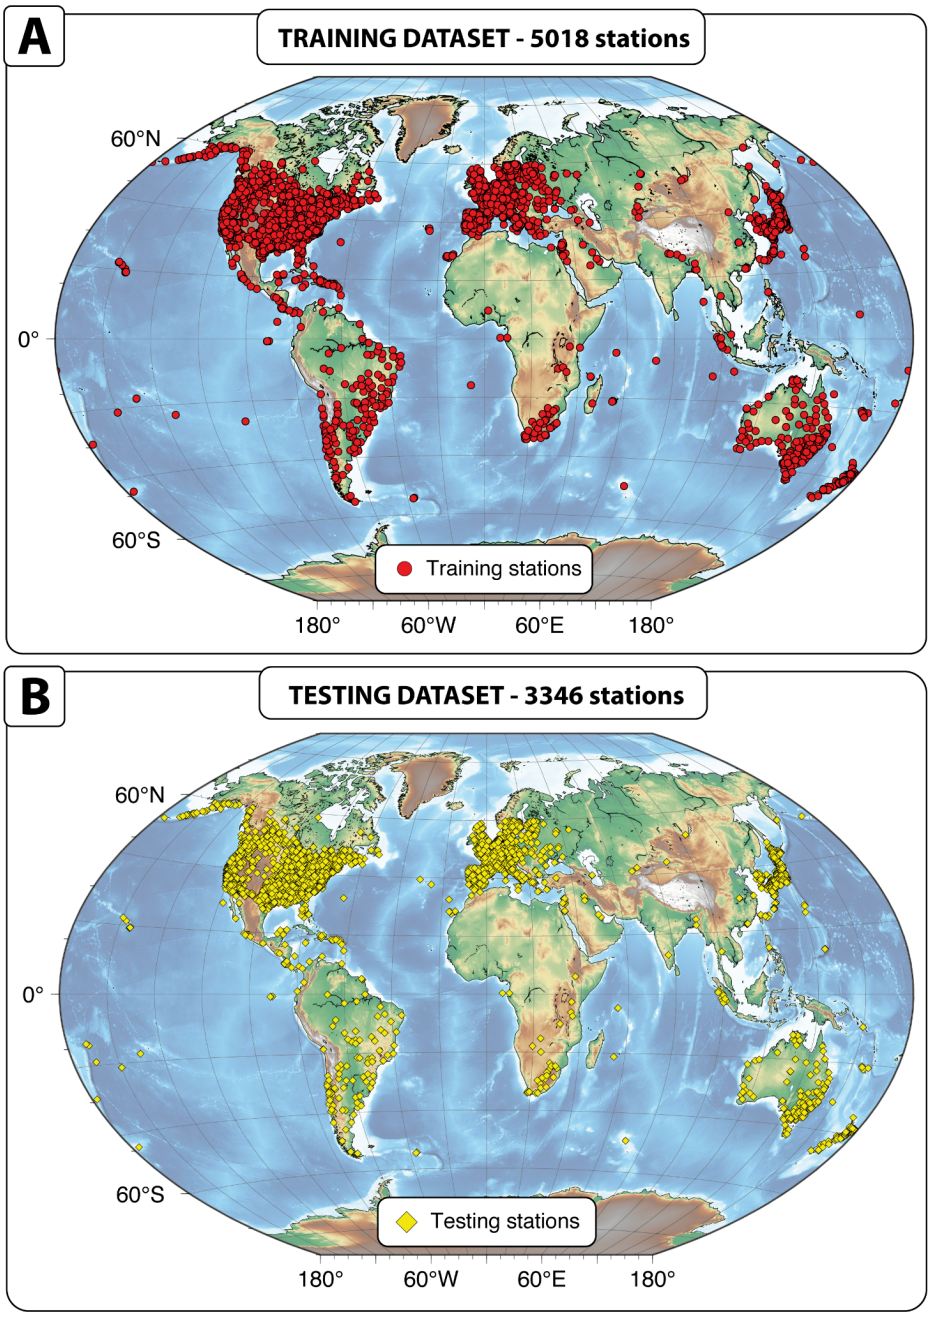

Fig. 1. Location of the GNSS stations used as the training dataset for realistic synthetic noise time series generation (**A**), and of stations utilized as the testing set to assess the performances of the Denoiser models (**B**).

2.2. Generating synthetic time series

Here, we use exclusively synthetic data for training our denoising DL models, while we evaluate the performance of the models using both synthetic and real data. This approach allows us to overcome

data scarcity and imbalance in real time-series by generating a large number of diverse samples, including patterns that are not frequently represented in the real data, such as transient slow earthquakes. This strategy allows training a more robust DL model. Synthetic data samples are generated by assuming that a GNSS displacement time series can be described as:

$$y(t) = S(t) + \xi(t)$$

(**Eq.1**)

where $y(t)$ is the displacement as a function of time t, $S(t)$ is a trajectory model (what we are able to model with a parametric fit) and $\xi(t)$ is the noise term, i.e. what is left after fitting the time series with a trajectory model (i.e., the residuals; **Figure 2A**). To generate highly realistic synthetic time-series, we divide the generation process into two stages. Firstly we create synthetic trajectories by randomly reassembling the parametric function of a prescribed trajectory model (**section 2.2.1**). Subsequently we create realistic synthetic residuals time-series by using a generative DL model (**section 2.2.2**).

2.2.1. Generating GNSS trajectories

To build surface displacement synthetic trajectory time-series we assume an extended version of Bevis and Brown [2014]'s trajectory model:

$$y(t) = mt + c + \sum_{k=1}^{n_k} A_k(t)\left[s_k sin(\omega_k t + \phi_k(t)) + c_k cos(\omega_k t + \phi_k(t))\right] + \sum_{j=1}^{n_j} b_j H(t - t_j) + \sum_{i=1}^{n_i} A_i H(t - t_i) log_{10}\left(1 + \left(\frac{t - t_i}{T_R}\right)\right) + \sum_{s=1}^{n_s} d_s tan^{-1}\left(\frac{t - t_s}{T_s}\right) + \sum_{g=1}^{n_g} F_g exp\left(\frac{-(t - t_g)^2}{2(E_g)^2}\right) + \xi(t)$$

(**Eq.2**)

According to this model (equation 2), displacement time-series are the result of the sum of a minimum number of transient basis functions and sometimes permanent functions. These permanent basis functions are the first-order polynomial term (linear secular trend, described by the coefficients m and q), the Fourier solutions for annual and semi-annual oscillations (sine and cosine, with n_k= 2, comprising coefficients $s_k$and c_k), and some possible steps (Heaviside function, $H(t - t_j)$, with b_j denoting the coefficients for events at time t_j), due to hardware changes, processing parameter discontinuities, or to the occurrence of large earthquakes (**Figure 2B**). These steps can be "smeared" across two samples in the daily time series. The smearing is generated by dividing the step widths over two days (two samples) following a random ratio. To reproduce the variability of seasonal motion that can be due to anomalous climatic and/or meteorological effects [e.g. Klos et al., 2018; Davis et al., 2012; Pintori et al., 2022], we incorporate a time-varying amplitude weighting $A_k(t)$ and a time-varying phase modulation $\phi_k(t)$ of the Fourier terms (**Figure 2C**), according to Davis et al. [2012] and Bedford and Mastella [2023]. Transient functions represent instead some time-dependent unexpected motions. Transients include postseismic decays following large earthquakes (modeled with a logarithmic decay, where T_r represents the characteristic time and A_i the amplitude), arctangent functions representing slow earthquakes (with d_s denoting the coefficients for events at time t_s with a given duration T_s), and Gaussian functions (at time T_g described by coefficients F_g and E_g), that can originate from volcanic inflations or that can arise from epistemic or aleatory nature of GNSS measurements. To generate a large spectrum of possible synthetic displacement time-series, we varied all the parameters characterizing each of the described signals reported in **EQ.2** sampling from a given normal distributed Gaussian range. The ranges employed for each parameter are set according to plausible values for real displacement time-series. By using a random generation process, we are able to create an augmented synthetic dataset of synthetic time-series that encompasses the entire range of potential variations observed in real GNSS time-series. Heaviside steps are generated by using a modified Gutenberg-Richter law, resulting in the addition of more smaller steps compared to larger ones. We generate trajectories with an arbitrary temporal length of 4 years.

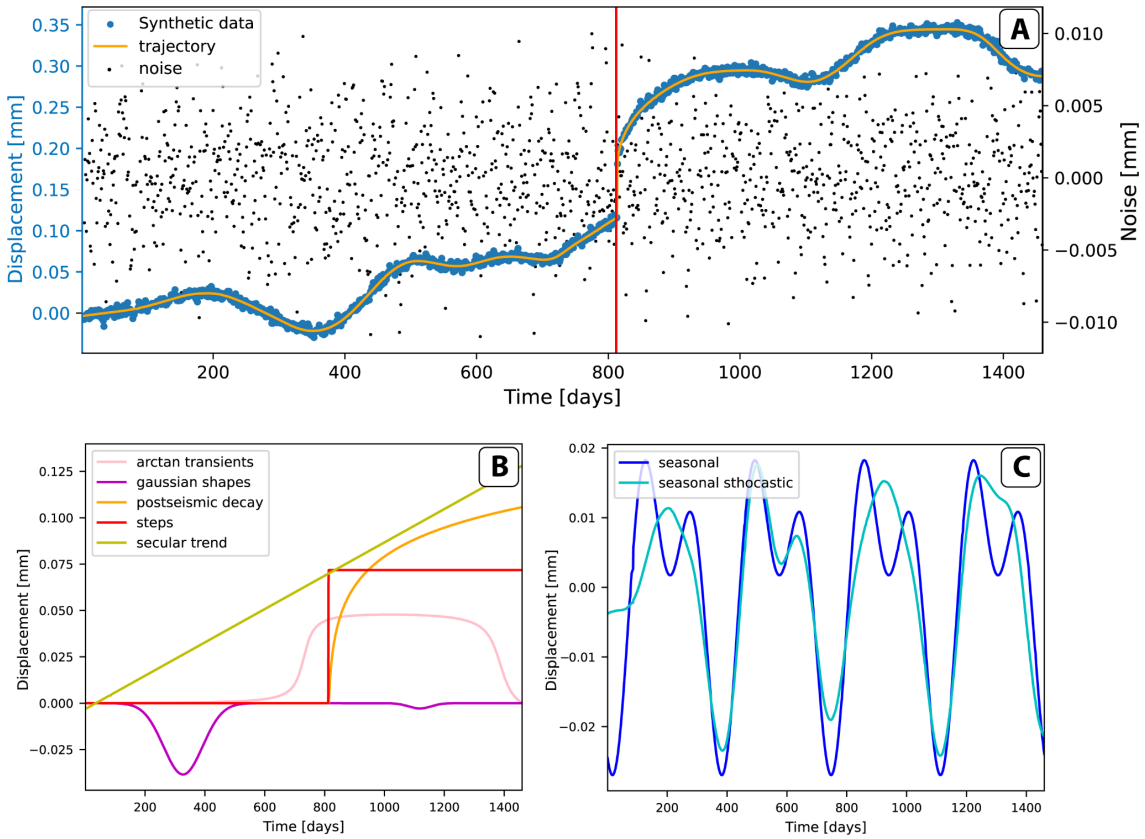

Figure 2. Example of a synthetic time-series described as the sum of a given trajectory model and noise (**A**). The trajectory model employed is defined as the sum of some permanent functions and time-dependent transients (**B-C**). The noise is generated with the GAN model.

2.2.2. Generating realistic synthetic noise with GANs

The nature of noise in GNSS displacement time-series is a matter of debate and each time series can be characterized by different typologies of noise [e.g. Amiri-Simkooei, 2009], like flicker noise (dominant at frequency below 18-20 days) and white noise (which prevails at frequencies higher than 8 days) [Santamaría-Gómez et al 2011; Rebischung et al. 2016; Niu et al., 2023]. White noise has no frequency dependence and seems to arise predominantly from the incorrect modeling of atmospheric delays [Williams et al. 2004; Gobron, 2021]. On the contrary, flicker noise is power law and a Gauss–Markov process [Santamaría-Gómez et al. 2011], but its physical generation mechanism is still unclear. As a consequence of flicker noise, GNSS displacement time series show time-correlated background noise [e.g. Williams et al. 2004]. These temporal correlations, well observable in the isolated residuals, can additionally arise from the underfitting of the original displacement time series as well as from some specific unknown signals that trajectory models are not able to fit. For example, it has been shown that the power spectrum of noise in geodetic time series can be biased by the estimation of trajectories' functional parameters [e.g. Gobron et al. 2022; Bedford & Bevis 2018]. Due to the variability of the power spectrum of noise in GNSS time-series, generating synthetic noise samples with a realistic distribution is thus a challenging task. Although one may think of applying a given noise spectrum to a synthetic time series [e.g. Xue and Freymueller, 2023], achieving a wide and balanced distribution of noise samples is actually non-trivial (**Figure S7**). To address this issue, we leverage Generative Adversarial Networks (GANs) to generate synthetic residuals that are realistic and representative of the variability of noise in daily GNSS time-series. GANs are a particular class of DL models that are able to generate synthetic samples by creating their own understanding of a given dataset [Goodfellow et al., 2014].

In this section, we first describe how we isolate residuals (the noise) from real GNSS time series to create a training dataset representing the input of GAN models. Subsequently we describe the basic structure of our generative models, while in the last section we report how we combine trajectories and residuals to assemble a realistic synthetic dataset of daily displacement time-series.

2.2.2.1. Isolating noise from real GNSS time series

To isolate the noise characterizing real GNSS displacement time series we use GrAtSiD, the Greedy Automatic Signal Decomposition algorithm [Bedford and Bevis, 2018]. GrAtSiD is an iterative greedy linear regression algorithm that decomposes the input time series into three parts: (1) the background seasonal signal, (2) an assortment of of secular and transient functions, and (3) the remaining residual (i.e., the raw time series minus parts 1 and 2, the noise). A detailed description of GrAtSiD is reported in the text **SI.1**. We applied GrAtSiD to all the time series presented in the training dataset (**Figure 3A**) after removing major outliers using a Hampel filter [Davies and Gather, 1993]. Without qualifying its nature, noise time-series are defined to be the residual signal relative to the GrAtSiD trajectory fitting (**Figure 3B**). These residual time series compose the dataset on which the GAN is trained to generate synthetic residual time series.

2.2.2.2. Converting residuals time-series to time-frequency spectrograms

To capture the temporal correlations of the noise, we represent residual displacement time-series (which have been obtained through the GrAtSiD trajectory decomposition) in the time-frequency domain by applying the Short-time Fourier transform (STFT). This operation computes a spectrogram of the input time signal by performing consecutive discrete Fourier transforms (DFT) over short overlapping windows. We apply the STFT by employing overlapping windows of 96 days. As a result, 3 years-long residual time series are transformed into spectrograms with a shape S = 48 x 48 (**Figure 3C**). The computed spectrogram highlights the change of the signal's frequency content over time, highlighting the temporal correlations of the residuals. We use the derived spectrograms as the input to train the generative models able to synthesize fake spectrograms of residuals.

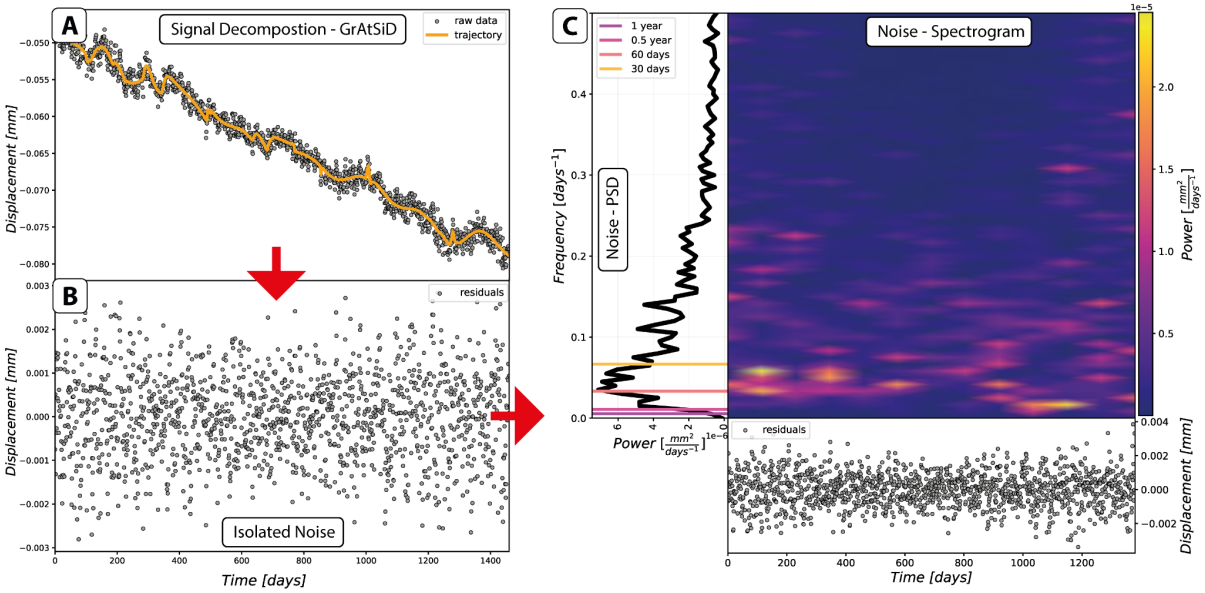

Figure 3. A) Example of a trajectory fit obtained using GrAtSiD. **B**) Residuals time-series, i.e. what is assumed to be noise left after time-series decomposition. **C**) Spectrogram of the residuals time-series in the time-frequency domain.

2.2.2.3. GAN model for noise generation

A GAN model consists of two deep networks called the generator and the discriminator, that are trained in a competitive manner in a zero-sum adversarial game framework [Goodfellow et al. 2014]. The generator is a regression model trained to synthesize realistic samples capable of fooling the discriminator model, a classification model trained to differentiate between real and synthetically generated data samples. The main structure of our GAN model is illustrated in **Figure 4**. The input vector z of the generator model is a Gaussian noise vector of length 100 with a uniform distribution between -1 and 1. This vector is subsequently manipulated through three transposed convolutional layers to synthesize a fake spectrogram. These layers upsample the input feature vector to generate an output feature image that has a greater spatial dimension. To capture more general patterns, each transposed convolutional layer has a larger kernel size that expands the network's receptive field. In between layers, batch normalization is applied between the layers to stabilize learning. The activation function after each layer is a ReLU, while the last layer is characterized by the sigmoid activation function. The discriminator network is a binary convolutional network that receives as input the 48 × 48 × 1 spec-

trogram. This network consists of three convolutional layers with ReLU activation that downscale the spectrogram. The kernel size employed in the successive layers is reduced to capture increasingly complex and abstract features while reducing the dimensionality of the data. After the first layer, a dropout layer with a rate=0.2 is used to prevent overfitting. At the end of the network, a dense layer gives the input spectrogram a score between 0 and 1 (fake or real sample) by applying a sigmoid activation function. We use Binary Cross Entropy for both generator and discriminator loss for gradient descent, with the Adam optimizer [Kingma and Ba, 2015]. During the training, we monitor the discriminator accuracy for fake and real examples. The optimal scenario for GAN is for the discriminator's accuracy to be about 0.5, which is when the discriminator's ability to distinguish between real and false data is no better than chance. However, this condition alone is insufficient to ensure correct training, as the discriminator accuracy can still be 0.5 even if the model undergoes mode collapse. Mode collapse in GANs occurs when the generator produces a limited set of examples that the discriminator finds difficult to distinguish from the real data, resulting in repetitive and less diverse generated samples that do not reassemble the full diversity of the target distribution. To address this, we carefully monitor the GAN's performance during training. When the accuracy stabilizes between 0.4 and 0.6, we stop the training and we evaluate model's performances by checking the distribution of the generated samples with respect to real ones. For this purpose, we apply different statistical tests, which are reported in Section **S3** in the **SI**. Once all tests are satisfied, it means the GAN model is able to generate realistic synthetic samples which are at the same time similar but indistinguishable from real ones, and the training is halted. Given that the residuals of vertical and horizontal components are characterized by strongly different amplitudes and statistics, we build two different GAN models: one using only the time series of the vertical component and one considering the time series of horizontal components (E and N).

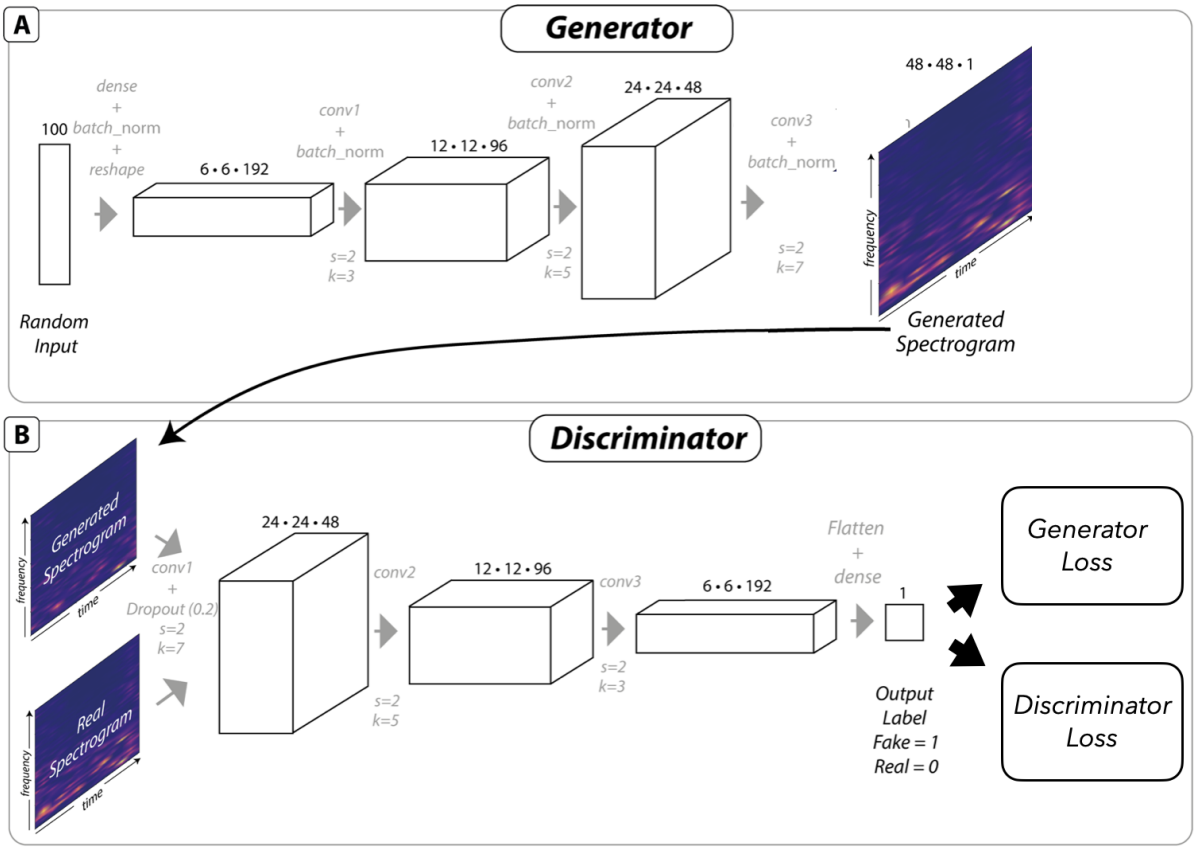

Figure 4. Schematic representation of the GAN Network architecture we employed to generate synthetic spectrograms of residuals time series.

2.2.3. Combining synthetic trajectories and synthetic residuals time series

Once the GANs models are successfully trained, the synthetically created spectrograms in the time-frequency domain are subsequently reverted in time-domain by employing the Griffin-Lim algorithm [Griffin and Lim, 1994] (**text S2**). In this way we obtain residuals time-series comparable to the real ones obtained through time-series decomposition with GrAtSiD. The last step to generate GNSS synthetic time series is represented by the simple random combination of the synthetic residual time se-

ries with the synthetic trajectories (**Figure 2A**). By applying this combination a dataset consisting of 20,000 synthetic time series of 4 years is built.

3. Deep Learning method

We take advantage of convolutional autoencoder networks and recurrent neural networks to jointly learn and model a compact representation of the features of trajectory signals and the features of variable noise that characterize GNSS displacement time-series. This choice was inspired by the outstanding ability of autoencoders in denoising seismic data [Zhu et al., 2019], and of recurrent neural networks in GPS time-series analysis [Yamaga and Mitsui; 2019; Lin et al., 2021; Costantino et al., 2023a, Xue and Freymueller, 2023]. Predicting the noise of a GNSS time-series is formally equal to predicting its trajectory model. However, predicting the noise assures the stationarity of the target, which is a fundamental concern in every kind of DL application for time-series analysis. Noise prediction is a challenging task in the presence of steps. Steps can be mistaken as offsets and detecting offsets in GNSS displacement time series can be a non-trivial and non-unique task [e.g. Gazeaux et al. 2013]. To facilitate the prediction of the noise in the presence of steps, in this work we decided to build two subsequent DL models. The first model, that we called the Step model, is trained to detect discontinuities. Subsequently, a second model, called the Noise model, is trained to predict the noise of a given time-series, by using as an a priori information the detection of discontinuities obtained with the previous model.

3.1. Step model

In the context of data analysis, the task of identifying a step in GNSS displacement time series can be considered as an anomaly (or outlier) detection problem. This type of issue pertains to the recognition of infrequent items, events, or observations that differ greatly from the majority of the data. A step in a GNSS time series is a rare event that can be both earthquake-induced as well as having anthropic origin. It has been already demonstrated by Crocetti et al., 2021 that ML can be used to automatically detect earthquake-related steps in GNSS time series. Following a similar approach, but using DL models and a more general input dataset, we here tackle the step identification problem as a supervised classification problem. The classification task is formally defined as follows. Consider $X[t_{i-w-1},t_i]$ as a one-dimensional (1-D) vector of daily GNSS displacement measurements with a temporal length of w days until the day t_i. This sequence of data is used to predict a probability binary vector $B[t_{i-w-1},t_i]$ of 0 or 1 depending if the point is a step or not. Specifically, if $X[t_s]$ represents the onset of a step, then $B[t_s]=1$, while if $X[t_s]$ is not a step, then $B[t_s]=0$ (**Figure S1**). Our classification network is a convolutional Autoencoder (**Figure 5A**). Autoencoders are a particular class of DL architectures which are capable of learning two functions: a streamlined encoding function that reduces the input data dimensionality, and a decoding function that generates a different representation of input data again. Our architecture draws inspiration from the Phase-Net architecture [Zhu and Beroza, 2019], devised for seismic waves picking. This architecture has been modified to deal with 1-D time-series data from U-Net [Ronneberger et al. 2015]. Our network consists of a stacked series of 1-D convolutional layers with descending and then ascending sizes. Especially in the encoding branch, three convolutional layers with ReLU activation downscale the input vector of w days. Subsequently, the shrunken vector is processed through three transposed convolutional layers, which perform upsampling. The process ultimately generates an output vector of the same dimension as the input vector, w. We incorporated skip connections to enhance training convergence and improve prediction accuracy. Our network processes single-component displacement time series individually. Note additionally that the supervised step model is trained only using synthetic data, for which we know exactly the beginning of a step. The binary cross entropy has been chosen as the loss function for gradient descent during the training.

3.2. Noise model

The information regarding the locations of steps, obtained from the probability output of the step model, is integrated with the raw time series to construct the input for the second model that we call the Noise model. The Noise model is a regression model trained to predict the noise of a given time-series. Formally, this regression task can be formulated as follows: the binary vector $B[t_{i-w-1},t_i]$ represents the output of the step model, while the one-dimensional vector $G[t_{i-w-1},t_i]$ corresponds to the daily GNSS position points with a temporal length of w days, previously used as an input by the Step model. These two vectors are combined to form a stacked vector $X[t_{i-w-1},t_i,2]$ with dimensions $[w,2]$. This data tensor is utilized to predict the relative noise $Y[t(P)]$ for a specific point located at position $P \in w$ within the window. Hence, this regression task can be classified as a many-to-one problem, involving

multiple inputs and a single output. The architecture of the noise model is depicted in **Figure 5B**. The network consists of 4 fully connected Bidirectional LSTMs (Long Short Term Memory) [Hochreiter and Schmidhuber, 1997] with descending size followed by a Dense layer at the end. LSTMs are a set of Recurrent Neural Networks (RNN) that can successfully model sequences of data learning temporal dependencies by recursively updating an internal status. With respect to standard LSTMs, Bidirectional LSTMs combine both forward and backward information flow, allowing the network to capture contextual information from both past and future states. LSTMs layers have a descending size, in order to decrease the latent feature vector towards the one-dimensional output. The activation function after each recurrent layer is a ReLU, while the last layer is characterized by the linear activation function. During the training process, the loss function for gradient descent has been selected as the mean squared error (MSE).

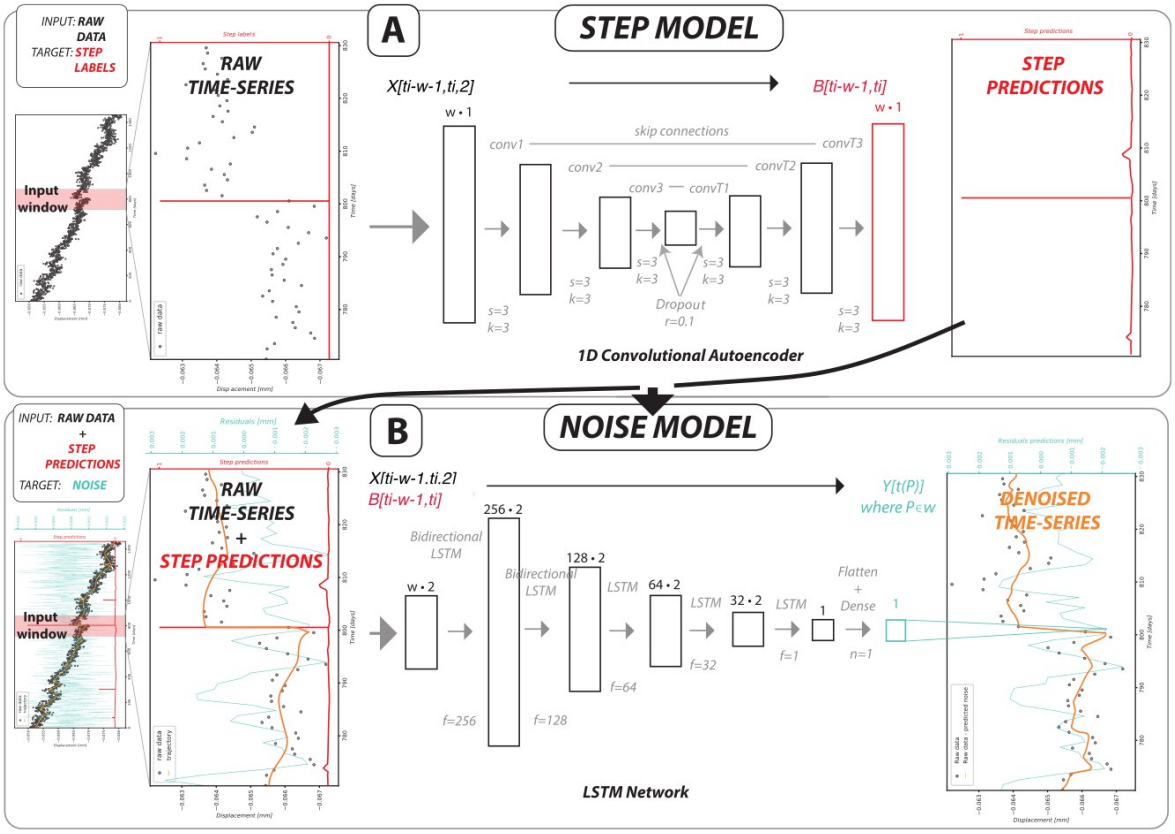

Figure 5. Schematic representation of the DL strategy and architectures we employed to denoise GNSS daily displacement time series.

3.3. Training strategy

Both the Step model and the Noise model are trained using 60% of the synthetic time-series (12000 time-series), while from the remaining time-series 20% are used as validation for model selection and hyperparameters fine-tuning and the last 20% for testing. As a result of the time-series windowing, the training set consists of a total of 20 million examples. Given that in the synthetically generated time series steps are relatively rare (we generate 3-years synthetic time-series with a variable number of steps between 0 to 5), the step identification task is a strongly imbalanced classification problem. We therefore added 20 million additional steps to augment the training set of the Step model. These additional jumps are characterized by random amplitudes and are located at random positions in the input temporal windows to try to balance the disproportion of the two classes (0 and 1). Note that the training set of the Noise model is not augmented by these additional steps. The Step and Noise models are trained independently with a learning rate of 5e-4 and 7e-4 respectively. For both models, we used a batch size of 256. Input sequences are randomly shuffled during training to promote generalization and prevent overfitting. To tune networks' hyperparameters we employed the Hyperband Ran-

dom search tuner implemented in the machine learning platform of Weight and Biases (https://docs.-wandb.ai) (**Text S4**). To speed up the tuning and prevent the overfitting of the training set, we implement thresholding early stopping with 4 epochs of patience based both on the validation and training loss. This early stopping mechanism does not impact training performance but prevents overfitting. Data are standardized to have zero mean and a standard deviation of one, therefore a scaling factor needs to be applied to the inputs. To additionally level the data, we remove its median to each input vector *X* of daily displacements.

4. Results

4.1. Steps Predictions

The performances of the Step detector model are quantitatively evaluated by using the testing dataset of synthetic examples, for which we know exactly the location of steps. Our model makes predictions of "positive" or "negative" outcomes based on whether a point is identified as a step or not. The step model is a binary classification model, and depending on the correctness of the predicted outcome, four cases are possible: true positive TP, true negative TN, false positive FP, and false negative FN. We calculate Precision and Recall to measure the model's performance. Recall represents the proportion of correctly identified positives out of all actual positives (R=TP/(TP+FN)), while Precision indicates the proportion of correct positive identifications (P=TP/(TP+FP)). These metrics are particularly useful for assessing the performance of the model in detecting the minority class (i.e., steps). Since our Step model provides a continuous output, we set an arbitrary threshold of 0.95 above which predictions are considered as steps, corresponding to class 1. We underline that this threshold is employed only to evaluate the Step model predictions but does not influence the Noise model (with the probability value between 0 and 1 flowing into the Noise model). **Figure S2A** shows the prediction performance as a function of this threshold and of the position of the predicted point inside the input window *w*. As the threshold decreases, the number of FP drastically increases, leading the Precision to decrease from 1, if the central point of the input window is predicted, to 0.2 if the last point is predicted. **Figure 6** depicts the step predictions of the E component of station J083, obtained with a model trained with an input of 31 days. This station is located and has been influenced by the 2011 Mw 9.1 Tohoku earthquake, which caused a large jump in the displacement time series. This discontinuity is perfectly detected by our DL model when the central point of the input window is the target (**Figure 6A**). When instead the model tries to predict the last day of the input window, predictions become more scattered and the model struggles to differentiate between a step and an outlier, producing a large number of false detections (**Figure 6B**). This tendency is also dependent on the temporal length chosen to build input examples (**Figure 7**): as this length increases from 16 to 61 days, the predictions become more reliable and the number of false detections strongly decreases (**Figure S2B**). Considering the central point framing, Precision varies from 1 (w=61 days) to 0.6 (w=16 days). Another two factors that have an impact on the predictions are the standard deviation of input examples, which is an indicator of the level of noise, and the amplitude of steps. This intuitively makes sense since small amplitude steps within noisier time series tend to make the detections appear more trivial (**Figure S2C and S3D**). Interestingly, steps with an amplitude higher than 0.09 m are always correctly detected, even when the standard deviation of the residuals is very high (0.004 m). This suggests that steps can be detected if their amplitude is around 2 standard deviations of the residuals.

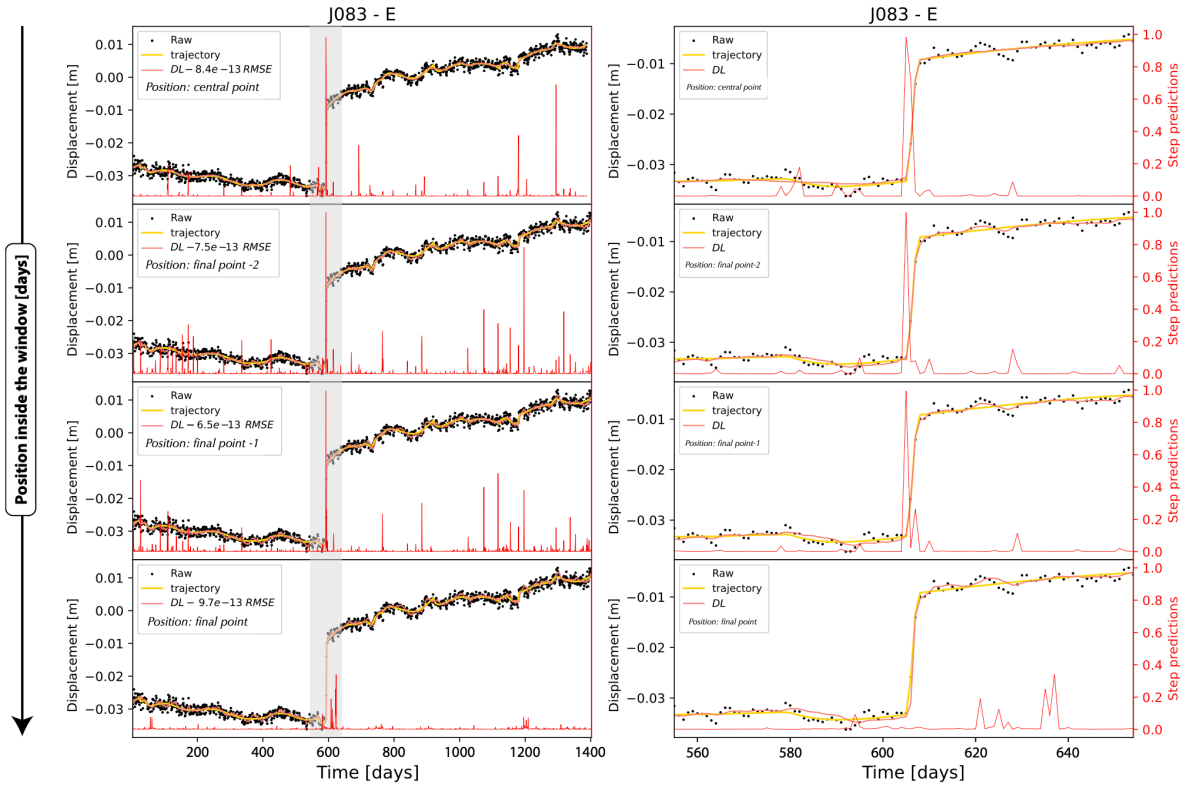

Figure 6. Denoised time-series and step predictions as a function of which target point (i.e. position) inside the input window (*w*=31 days*)* is predicted for the E component time-series of the J083 GNSS station. Shaded areas in the first column to the left represent the zoom reported in the second column to the right.

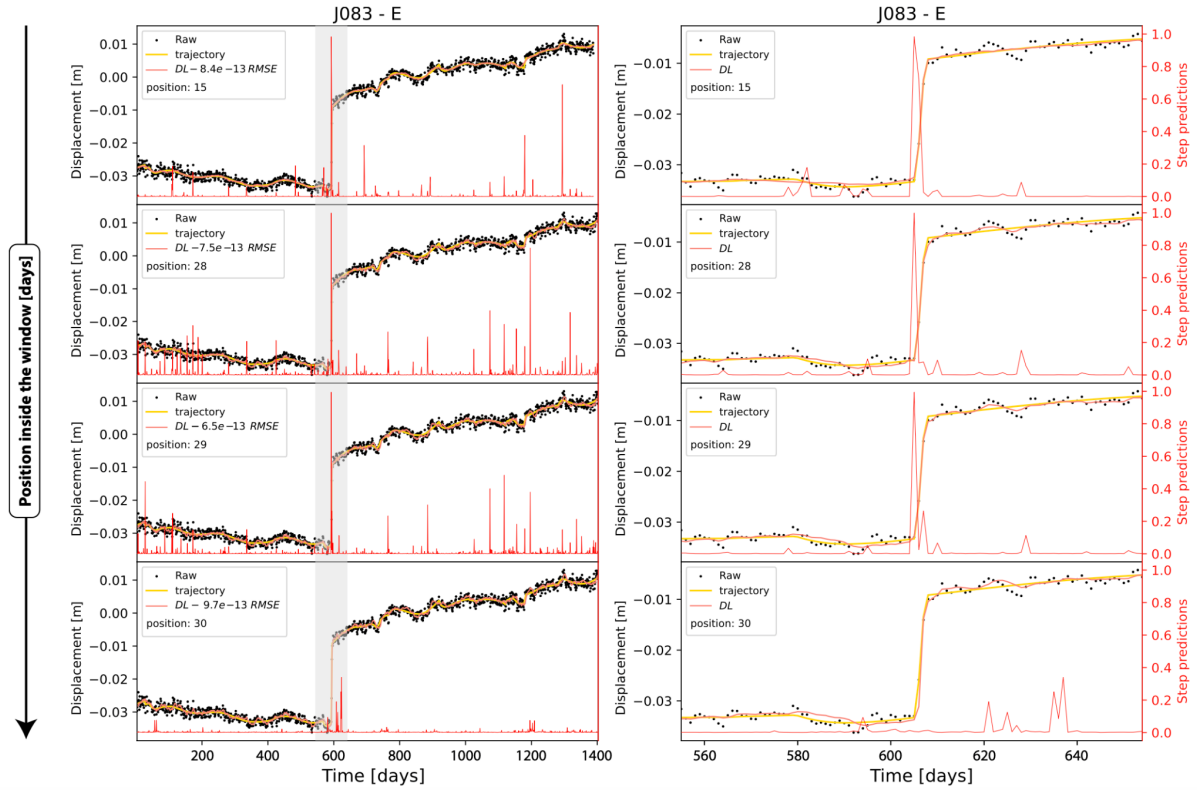

Figure 7. Denoised time-series and step predictions as a function of the input length *w* for the E component time-series of the J083 GNSS station. Shaded areas in the first column to the left represent the zoom reported in the second column to the right.

4.2. Noise predictions

The performance of the noise model is evaluated using the testing dataset of real time-series (**Figure 1B**). In line with the preprocessing procedure for training using synthetic data, the real-time series are standardized and windowed before being fed into the trained Noise model. For each time-series we calculate the Standard Deviation (STD) and the Absolute Root Mean Square error (RMSE) of the predicted noise (**Text S5**). We additionally calculate the RMSE with respect to the recovered noise isolated with GrAtSiD, that we call Relative RMSE (RRMSE). We emphasize that this metric is different from the classical relative RMSE; it simply calculates the RMSE by assuming that the noise isolated with GrAtSiD is the true noise characterizing a time-series. **Figure S3** shows these metrics as a function of the input length. Interestingly, if a larger input length is used, the RMSE and the STD increase (**S3A and S3C**, moving along the x axis), meaning that larger noise is predicted. This is clear, for example, in the first column of **Figure 7** moving from top to bottom: by increasing the length of the input window the cleaned time-series shown in orange becomes less noisy and better approximates the trajectory obtained with GrAtSiD (in yellow), which means that larger residuals are predicted by the DL approach. As for the Step model, what additionally influences the predictions is the position of the predicted point inside the input window *w* (**Figure 6** and **S3A** and **S3C,** moving along the y axis). In the case of central point framing, the RMSE and the STD increases, meaning again that larger noise is predicted. If instead the model is applied in real-time, i.e. the last point is predicted, a large amount of noise is still maintained in the filtered time-series. Differently to the other two metrics, the RRMSE exhibits a more complicated pattern. The best approximation of GrAtSiD trajectory fit (the lowest RRMSE) is obtained with an input length of 31 days and if the 29th days of the input window are predicted. For DNN architectures explored in this study, the central point framing does not always result in the best trajectory fit if compared with GrAtSiD, as evidenced by the model trained with an input length of 31 days, which shows the same RRMSE as the model trained in the last point framing. Interestingly, all the models are predicting noise relatively well even in the presence of discontinuities, albeit when the last point is predicted a small delay is caused in the denoised time-series, as evident in the last row of **Figure 6**. In general, the DL model reduces the "lift-off" from the true trend around the discontinuity, an effect that traditional filtering methods struggle to avoid. The model effectively turns down the noise in horizontal components to a greater extent compared to vertical components, primarily because the average noise level in the vertical component is 4 times higher than that in the horizontal components (average STD of 0.0016 and 0.0043 m respectively).

5. Discussion

5.1. Comparison with CMF filtering and Frequency-dependent filters

We here apply a common-mode filtering approach similar to the one applied by Márquez-Azúa and DeMets [2003] and we compare the relative denoised results with the ones obtained with our DL filtering routine. To build the CMF filter, we first calculate the Pearson correlation coefficient between each residual time-series (that we obtain after a GrAtSiD trajectory fit) as a function of the interstation distance. The results are represented in **Figure S4** where the color depicts the number of common days between each pair of time-series. As expected, the correlation coefficients between two residual time-series diminish as the interstation distance increases [e.g. Williams et al. 2004; Bogusz et al. 2015; Tian and Shen 2016; Kreemer & Blewitt, 2021]. This decrease is evident until around 3000 km, while for longer distances generally no correlation is observed. We considered this distance as the upper threshold to select the time series used to compute the common mode residuals for a given reference time-series. The common mode residuals are computed using a weighted median that takes into account the time length of each associated time series and its distance from the selected site. Essentially, longer time series that are very close to the selected one have a greater impact on the weighted median, meaning that the accuracy of the common mode residuals estimates improves in a time-dependent manner. It is important to note that the common mode median is not influenced by the trajectory model of the selected time series (the one whose noise is being reduced).

Furthermore, we conducted a comparative analysis between the filtered time-series obtained using our DL procedure and several frequency-dependent filters, including a simple moving average, a mo-

ving average with exponential weights (EMA), and a Butterworth filter. The EMA is a practical method for smoothing time series data by applying an exponential window function. Unlike the simple moving average, which assigns equal weight to all past observations, the EMA assigns progressively decreasing weights based on exponential functions that consider the passage of time, enabling prompt responsiveness to recent data fluctuations. The EMA simply computes a moving average in a given temporal window by multiplying each past observation to a weight based on an exponential distribution. To directly compare the results obtained with the DL when predicting the residuals of points located at different positions within the input window, we implement an EMA that accordingly assigns the maximum weight to the predicted point, that can be even inside the window itself and not exclusively the last one (**Figure S5**).

A Butterworth lowpass filter is a commonly used signal processing filter that allows low-frequency components to pass through while attenuating high-frequency components. It ensures a smooth and gradual transition between the passband, where signals are preserved, and the stopband, where signals are attenuated. We here apply a causal Butterworth filter.

In Figure 8A, we present a comparison between the MAHI time-series (E component) filtered using our DL approach and the CMF for the period between 2011 and 2015. For the DL approach, we employed a window length of 31 days and predicted the noise of the 29th day. Although the DL filter could potentially achieve better performance by predicting the central point in terms of RMSE (as shown in **Figure 6A**), it still outperforms the CMF (from visual inspection of the scatter of the remaining time series and from RMSE metrics). As depicted in **Figure 8C**, the residuals recovered using the DL approach exhibit significantly higher RMSE values (44%). The superior performance of the DL approach compared to the CMF is clearly evident in **Figure S6**. This figure depicts the distribution of errors for all the time-series denoised with both methods belonging to the testing set. Regarding the DL approach, the metrics are shown by considering two scenarios: one where the central point of the input window with a duration of 31 days is predicted (**Figure S6 panels A to C**), and one where the last point is predicted (**Figure S6 panels D to E**). The DL method demonstrates greater effectiveness in recovering isolated noise, as compared to the CMF. This is reflected in higher values of RMSE and STD for the predicted noise (**Figure S6 panels A-B and D-E**, respectively). The disparity in performance between the CMF and DL is particularly striking when the RRMSE is considered, which exhibits substantial differences (on average, approximately 80% higher for the CMF) (**panels C and F**). Another way to compare the predictions, is through the power spectral density of the denoised time-series, as depicted in **Figure 8D**. Particularly, the DL approach excels at reducing high-frequency scatter in the input time-series. The improvement in the signal-to-noise ratio becomes apparent at frequencies higher than 80 days, while the removal of lower-frequency noise proves challenging. When compared to the RMSE of the isolated noise obtained through the GrAtSiD trajectory fit, the RMSE achieved by the DL approach is 40% lower, indicating that the DL approach only partially mimics an approximation of a trajectory fit.

Figure 8B shows the comparison between the denoised results obtained with the frequency-dependent filters and the DL approach. Both the Moving Average and the Butterworth Filter cause an evident delay in the input time-series, even if applied considering two samples in the future (to be compared with the DL approach which predicts the point at the 29th of a moving window of 31 days). These traditional methods are unsuitable for real-time tracking of slow slip events, as depicted in **Figure 8D**. On the other hand, the EMA (Exponential Moving Average) performs relatively well in tracking the initiation of a slow slip event, as observed in **Figure 8D** around days 520-525. However, as time progresses from the beginning of the transient, the cleaned time-series accumulates more and more delay until the conclusion of the transient (days 528 to 560 in **Figure 8D**). As a consequence, this corresponds to a distortion of the investigated input signal. This demonstrates the improved performance of the DL approach in tracking the temporal evolution of tectonic transients almost in real-time.

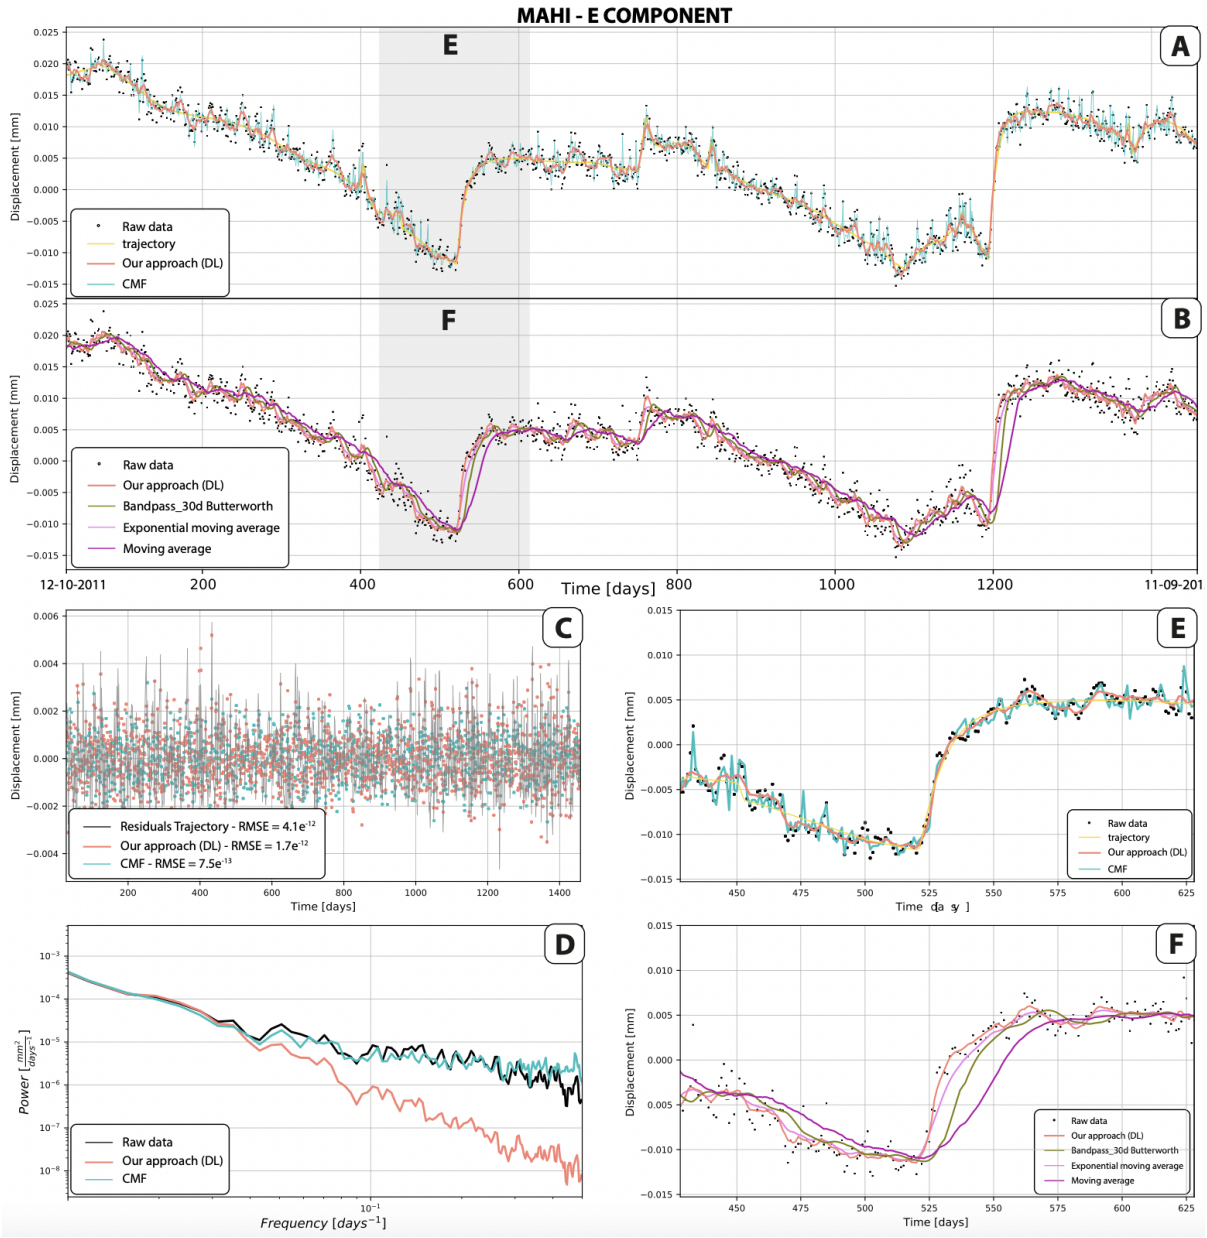

Figure 8. Comparative analysis of the denoising performances of our DL approach with respect to the CMF (**A-C-D-E**) and various frequency dependent filters (**B-F**). All the filters are applied by employing a moving window of 31 days; in specific the DL model predicts the noise of the 29th day. Panel **E** and **F** reports zoom of shaded areas in panels **A**-**B**..

5.2. Unraveling tectonic transients - Surveilling the Hikurangi subduction zone

We apply our denoising approach to the GNSS network located along the North Island of New Zealand. This region encompasses the Hikurangi subduction zone, where the Pacific Plate subducts westward beneath the North Island along the Hikurangi Trough. Significantly, the GPS displacement time-series collected in this area have revealed a range of slow-motion earthquakes, also known as slow slip events (SSEs) [Wallace and Beavan, 2010]. Due to the pronounced variations observed in SSE characteristics, such as recurrence times, durations, and depths, both along the strike and dip of the subduction margin, the Hikurangi subduction zone stands out as a globally unique natural laboratory [Wallace 2020]. As a result, the GPS time-series within this complicated system are affected by the overlapping influence of multiple processes operating at different temporal and spatial scales. We here demonstrate the efficacy of our denoising approach in effectively highlighting signals masked by noise within such a compound geodynamic system, where distinguishing between noise and signal

can be challenging since they may share similar frequencies. We employed all available GNSS stations from 2010 to 2023 and applied a trained DL model to predict noise at the third last point within a 31-day input window. Velocity is simply determined by computing the derivative of the denoised time-series. **Figure 9** displays the cumulative velocity observed during a slow slip event generated in the southern part of the margin, offshore from Porangahau and Cape Turnagain, an event probably triggered bythe Kaikōura Earthquake [Wallace et al., 2018]. This region typically experiences significant SSEs approximately every 5 years (2006, 2011, and 2016) [Wallace et al. 2012, 2017]. **Figure 9** demonstrates that our denoising method effectively reveals the onset and propagation of the investigated SSE. Compared to the results obtained after the application of the EMA filter (red quivers in **Figure S10**), our DL method allows preservation of the amplitude of slow slip velocity, whereas the EMA filter causes an excessively smoothing of the motion. Importantly, our denoising method can be applied with only a 2-day delay, without introducing any delay in the input signal. This makes it possible to track and characterize SSEs in near real-time as they propagate.

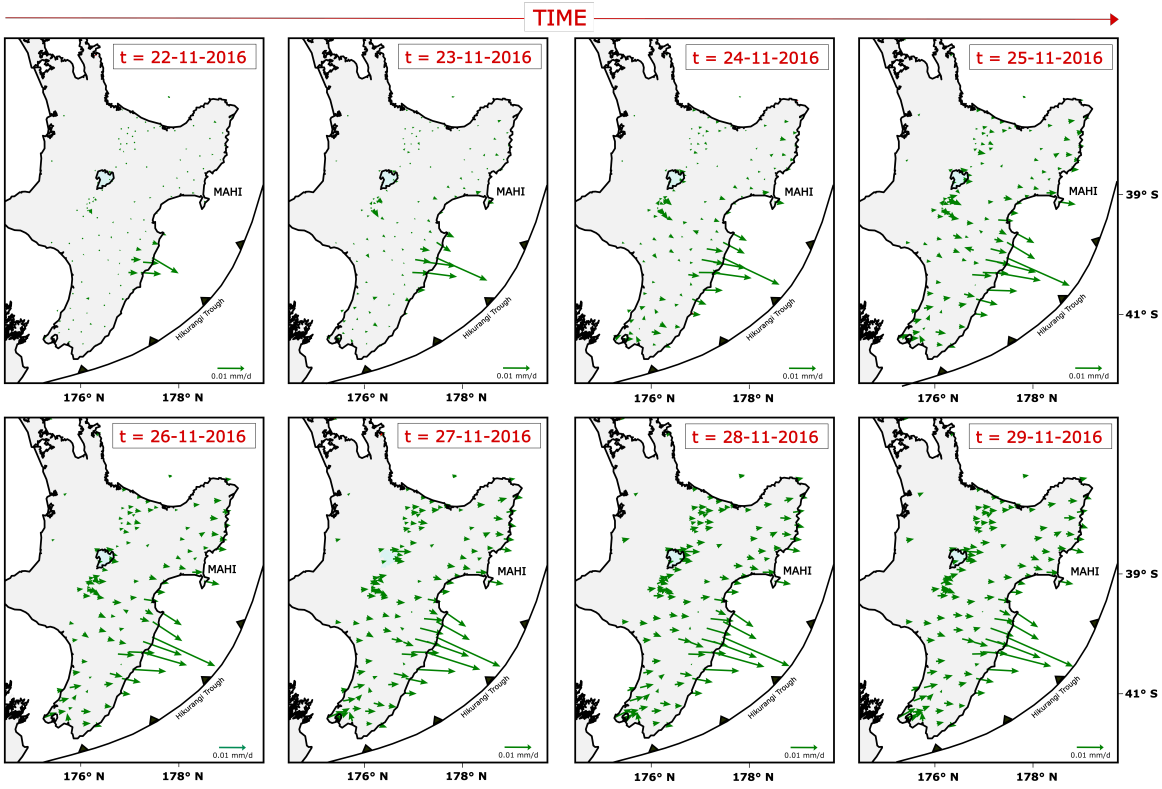

Figure 9. Cumulative time-dependent velocities of GNSS sites during the propagation of the slow-slip event at the Hikurangi subduction zone following the Kaikōura Earthquake as calculated after the application of our DL denoising method.

5.3. Approximating a trajectory fit in near real-time

Our denoising model aims to predict the residual of a trajectory model. At first glance, this approach may not appear to offer significant advantages compared to applying classical trajectory inversion routines like GrAtSiD or other similar methods [Bevis et al., 2019]. These classical methods typically assume a simpler trajectory functional model for fitting time-series, unlike the more complex model we used to generate our synthetic examples (as shown in **equation 2**). For example, GrAtSiD does not consider interannual variations of seasonal oscillations. But despite these simplifications, it has been demonstrated that GrAtSiD effectively filters out high-frequency noise in GNSS time-series allowing to track very subtle transient surface motions of tectonic plates [e.g. Bedford et al., 2020]. However, GrAtSiD has limitations. Its performance is dependent on the length of the input time-series, and it requires hyperparameter tuning to ensure appropriate convergence of the inversion routine, thereby avoiding under- or over-fitting of the data. This means that obtaining a satisfying trajectory fit is not a straightforward task and the inversion, if not run in parallel, can be computationally very expensive. In light of these concerns, our deep learning (DL) approach proposes an enhancement to GrAtSiD solutions, and in general of trajectory inversions. Essentially, the DL model learns to generalize an optimal

fit through a weakly supervised [e.g. Robinson et al., 2020] training method using millions of very different but realistic synthetic trajectories masked by the noise.
When compared with conventional filters commonly used in GNSS analysis, our method offers several notable advantages. Firstly, it effectively removes noise from time-series data, even in the presence of significant steps or jumps, without the need for explicit step detection and modeling. This streamlined approach simplifies the denoising process and enhances its efficiency. Secondly, operating on a station-by-station basis, our method can successfully denoise data from isolated stations, making it applicable even in situations where spatially dense GNSS networks are not available. On the other hand, in the case of dense networks, the use of the DL model enables to obtain more reliable trajectories, forming a basis for the application of a more robust CMF. Employing the DL denoiser model as an intermediate step effectively reduces the risk of CMFs inadvertently mapping local transients into noise, a potential consequence when the original trajectory model fits are suboptimal. Future effort will be dedicated to improve the DL model in order to explicitly consider the spatial correlation of noise within the approach itself. By incorporating the residuals of neighboring stations, the DL could become a CMF itself.
Thirdly, our method enables near real-time denoising of daily displacement time series, facilitating the accurate representation of the underlying signal without introducing noticeable temporal delays. This feature is particularly valuable in capturing time-sensitive changes and facilitating the timely detection of apparent transient features. Considering that observations from geodetic data suggested the possibility of anticipating earthquakes through a preparatory phase, consistent with the pre-slip model [e.g., Kato & Ben-Zion, 2021; Ohnaka, 1992], the ability of our method to filter geodetic data holds significant implications for time-varying seismic hazard assessment. It is worth mentioning that our filtering method was applied to high-level GNSS products, specifically PPP daily displacement time-series that have already undergone processing. This processing step typically takes some days for NGL data, and just two days for EarthScope solutions, so our method is not yet applicable in real-time. However, advancements in satellite technologies, including the advent of the Galileo constellation, are expected to improve the timeliness of obtaining accurate positions of geodetic stations. This means that in the near future, our method could be applied in real-time, further enhancing its practical utility.

6. Conclusion

We here propose a pioneering method to denoise GNSS daily displacement time-series by leveraging the power of DL algorithms. Our DL model is trained on a large dataset of synthetic but realistic GNSS measurements, enabling it to extract relevant features and filter out high-frequency noise in real time-series. This prototype effort lays a strong foundation for the development of more effective denoising methods using DL techniques. By incorporating real-time processing capabilities, it represents an initial valuable tool in unlocking the full potential of geodetic data for seismic hazard surveillance. Real-time geodetic monitoring using our approach holds the potential for early detection of anomalous plate movements and critical insights into potential earthquake occurrences.

Data and code availability

The models' implementations were developed using the Python package Tensorflow 2 [Abadi et al., 2016]. The training was conducted on a computer equipped with a single NVIDIA 457.09 GeForce RTX3090 GPU. Using a python data generator, training the Step model and the Noise model, one after the other, takes from 5 to 10 hours. Most of the figures were generated using the Python package Matplotlib v3.1.0 [Hunter, 2007] while map plots have been produced thanks to the GMT software (and its Python wrapper, pyGMT) [Wessel et al., 2019]. The GNSS data used in this work are open-access and available from the Nevada Geodetic Laboratory (NGL - http://geodesy.unr.edu). The DL models built in this work and some useful codes needed to use them will be available after publication at https://github.com/jekmasto/Daily_GNSS_DL_Denoiser. The GrAtSiD code is available at https://github.com/TectonicGeodesy-RUB/Gratsid.

Funding statement

Jonathan Bedford's ongoing development of the GrAtSiD code is funded by the European Union (ERC, Tectovison, 101042674). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., et al. (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint. https://doi.org/10.48550/arXiv.1603.04467.
- Amiri-Simkooei, A. R. (2009). Noise in multivariate GPS position time-series. Journal of Geodesy, 83(11), 1093-1106. https://doi.org/10.1007/s00190-008- 0251-8.
- Azarbad, M.R., Mosavi, M.R. A new method to mitigate multipath error in single-frequency GPS receiver with wavelet transform. *GPS Solut* 18, 189–198 (2014). https://doi.org/10.1007/s10291-013-0320-1.
- Bedford, J., & Bevis, M. (2018). Greedy automatic signal decomposition and its application to daily GPS time series. Journal of Geophysical Research: Solid Earth, 123, 6992– 7003. https://doi.org/10.1029/2017JB014765.
- Bedford, J.R., Moreno, M., Deng, Z. *et al.* Months-long thousand-kilometre-scale wobbling before great subduction earthquakes. *Nature* 580, 628–635 (2020). https://doi.org/10.1038/s41586-020-2212-1
- Bedford, J., Mastella, G. (2023): Separating seasonal oscillations from tectonic signals using advanced regression routines and denoised GNSS displacement time series, XXVIII General Assembly of the International Union of Geodesy and Geophysics (IUGG) (Berlin 2023). https://doi.org/10.57757/IUGG23-4347.
- Bevis M, Brown A, Kendrick E (2012) Devising stable geometrical reference frames for use in geodetic studies of vertical crustal motion. J Geodesy 87:311–321. doi:10.1007/s00190-012-0600-5
- Bevis, M., & Brown, A. (2014). Trajectory models and reference frames for crustal motion geodesy. Journal of Geodesy, 88(3), 283–311. https:// doi.org/10.1007/s00190-013-0685-5.
- Bevis M., Bedford J., Caccamise D. (2019) The art and science of trajectory modelling, in Montillet JP and Bos MS *Geodetic Time Series Analysis in Earth Sciences*. Springer Geophysics. https://doi.org/10.1007/978-3-030-21718-1_1.
- Blewitt G, Lavallée D (2002) Effect of annual signals on geodetic velocity. J Geophys Res 107:B72145. https://doi.org/10.1029/2001J B000570.
- Bletery Q., Nocquet J., ,The precursory phase of large earthquakes. Science 381,297-301(2023).DOI:10.1126/science.adg2565
- Bogusz J, Gruszczynski M, Figurski M, Klos A (2015) Spatio- temporal filtering for determination of common mode error in regional GNSS networks. Open Geosci. https://doi.org/10.1515/ geo-2015-0021.
- Costantino, G., Giffard-Roisin, S., Marsan, D., Marill, L., Radiguet, M., Mura, M. D., et al. (2023). Seismic source characterization from GNSS data using deep learning. *Journal of Geophysical Research: Solid Earth*, 128, e2022JB024930. https://doi.org/10.1029/2022JB024930.

- Crocetti,L.;Schartner,M.; Soja, B. Discontinuity Detection in GNSS Station Coordinate Time Series Using Machine Learning. Remote Sens. 2021,13,3906. https://doi.org/ 10.3390/rs13193906.
- D'Agostino, N., Silverii, F., Amoroso, O., Convertito, V., Fiorillo, F., Ventafridda, G., & Zollo, A. (2018). Crustal deformation and seismicity modulated by groundwater recharge of karst aquifers. Geophysical Research Letters, 45, 12,253– 12,262. https://doi.org/10.1029/2018GL079794.
- Davies, L., & Gather, U. (1993). The Identification of Multiple Outliers. *Journal of the American Statistical Association*, *88*(423), 782–792. https://doi.org/10.1080/01621459.1993.10476339.
- Davis, J. L., Wernicke, B. P., & Tamisiea, M. E. (2012). On seasonal signals in geodetic time series. *Journal of Geophysical Research*, **117**, B01403. https://doi.org/10.1029/2011JB008690.
- Dittmann, T., Liu, Y., Morton, Y., & Mencin, D. (2022). Supervised machine learning of high rate GNSS velocities for Earthquake strong motion signals. Journal of Geophysical Research: Solid Earth, 127, e2022JB024854. https://doi.org/10.1029/2022JB024854
- Dmitrieva, K.; Segall, P.; DeMets, C. Network-based estimation of time-dependent noise in GPS position time series. J. Geod. 2015, 89, 591–606. doi:10.1007/s00190-015-0801-9.
- Dong, D., Fang, P., Bock, Y., Webb, F., Prawirodirdjo, L., Kedar, S., and Jamason, P. Spatiotemporal filtering using principal component analysis and Karhunen-Loeve expansion approaches for regional GPS network analysis. Journal of geophysical research: solid earth, 111(B3), 2006. doi: 10.1029/2005JB003806.
- Ellsworth, W.L., Bulut, F. Nucleation of the 1999 Izmit earthquake by a triggered cascade of foreshocks. Nature Geosci 11, 531–535 (2018). https://doi.org/10.1038/s41561-018-0145-1
- Frank, W. Slow slip hidden in the noise: The intermittence of tectonic release. Geophysical Research Letters, 43(19):10–125, 2016. doi: 10.1002/2016GL069537.
- Fu, Y., Freymueller, J. T., and Jensen, T.: Seasonal hydrological loading in southern Alaska observed by GPS and GRACE, Geophysical Research Letters, 39, L15310, https://doi.org/10.1029/2012GL052453, 2012.
- Gaebler, P.J., Ceranna, L. Performance of the International Monitoring System Seismic Network Based on Ambient Seismic Noise Measurements. *Pure Appl. Geophys.* **178**, 2419–2436 (2021). https://doi.org/10.1007/s00024-020-02604-y.
- Gazeaux, J., et al. (2013), Detecting offsets in GPS time series: first results from the Detection of Offsets in GPS Experiment, *J. Geophys. Res. Solid Earth*, 118, 2397– 2407, doi:10.1002/jgrb.50152.
- Gobron K, Rebischung P, Van Camp M et al (2021) Influence of aperiodic non-tidal atmospheric and oceanic loading deformations on the stochastic properties of global GNSS vertical land motion time series. J Geophys Res: Solid Earth 126(9):e2021JB022370. https://doi.org/10.1029/2021JB022370.
- Gobron, K., Rebischung, P., de Viron, O. *et al.* Impact of offsets on assessing the low-frequency stochastic properties of geodetic time series. *J Geod* 96, 46 (2022). https://doi.org/10.1007/s00190-022-01634-9.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al., (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), Advances in neural information processing systems (Vol. 27, pp. 2672–2680). Curran Associates, Inc. Retrieved from http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf.
- Griffiths, J., and J. Ray (2009), On the precision and accuracy of IGS orbits, *J. Geod.*, 83, 277– 287, doi:10.1007/s00190-008-0237-6.
- Griffin D. W. and Lim J. S., "Signal estimation from modified short-time Fourier transform," IEEE Trans. ASSP, vol.32, no.2, pp.236–243, Apr. 1984.
- Gualandi, A., Serpelloni, E., & Belardinelli, M. E. (2016). Blind source separation problem in GPS time series. *Journal of Geodesy, 90*(4), 323–341.
- Jackson DA, Chen Y (2004) Robust principal component analysis and outlier detection with ecological data. Environmetrics 15(2):129–139 DOI: 10.1002/env.628.
- Johansson JM, Davis JL, Scherneck H-G et al (2002) Continuous GPS measurements of postglacial adjustment in Fennoscandia 1. Geodetic Res J Geophys Res 107:2157. https://doi.org/10.1029/2001J B000400.
- Jolivet, R & Frank, W. B. (2020). The transient and intermittent nature of slow slip. *AGU Advances*, *1*, e2019AV000126. https://doi.org/10. 1029/2019AV000126.
- He X., Hua X., Yu K., Xuan W., Lu T., Zhang W., Chen X., 2015. Accuracy enhancement of GPS time series using principal component analysis and block spatial filtering, Adv. Space Res., 55(5), 1316–1327.10.1016/j.asr.2014.12.016.

- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735.
- 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.
- Kato, A., & Ben-Zion, Y. (2021). The generation of large earthquakes. *Nature Reviews Earth & Environment*, *2*(1), 26–39. https://doi.org/10.1038/s43017-020-00108-w
- King, M. A., and C. S. Watson (2010), Long GPS coordinate time series: Multipath and geometry effects, *J. Geophys. Res.*, 115, B04403, doi:10.1029/2009JB006543.
- King, M. A., and S. D. P. Williams (2009), Apparent stability of GPS monumentation from short-baseline time series, *J. Geophys. Res.*, 114, B10403, doi:10.1029/2009JB006319.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the International conference on learning representations* (p. 141). Retrieved from https://arxiv.org/abs/1412.6980.
- Klos, A., Bos, M.S. & Bogusz, J. Detecting time-varying seasonal signal in GPS position time series with different noise levels. *GPS Solut* 22, 21 (2018). https://doi.org/10.1007/s10291-017-0686-6.
- Kreemer, C., Blewitt, G. Robust estimation of spatially varying common-mode components in GPS time-series. *J Geod* **95**, 13 (2021). https://doi.org/10.1007/s00190-020-01466-5.
- Langbein J (2008) Noise in GPS displacement measurements from Southern California and Southern Nevada. J Geophys Res 113(B05):405. doi:10.1029/2007jb005247.
- Li, Y., Xu, C., Yi, L. *et al.* A data-driven approach for denoising GNSS position time series.*J Geod* 92, 905–922 (2018). https://doi.org/10.1007/s00190-017-1102-2.
- Lin, J.-T., Melgar, D., Thomas, A. M., & Searcy, J. (2021). Early warning for great earthquakes from characterization of crustal deformation patterns with deep learning. Journal of Geophysical Research: Solid Earth, 126, e2021JB022703. https://doi. org/10.1029/2021JB022703.
- Liu, B., Dai, W., Peng, W., and Meng, X.: Spatiotemporal analysis of GPS time series in vertical direction using independent component analysis, Earth Planet. Sp., 67, 189, https://doi.org/10.1186/s40623-015-0357-1, 2015.
- Márquez-Azúa, B., and DeMets, C. (2003), Crustal velocity field of Mexico from continuous GPS measurements, 1993 to June 2001: Implications for the neotectonics of Mexico, *J. Geophys. Res.*, 108, 2450, doi:10.1029/2002JB002241, B9.
- Materna, K., Bartlow, N., Wech, A., Williams, C. & Bürgmann, R. Dynamically triggered changes of plate interface coupling in Southern Cascadia. Geoph. Res. Lett. 46, 12890–12899 (2019).
- Melgar D., LeVeque R., Dreger D., and Allen R., Kinematic rupture scenarios and synthetic displacement data: An example application to the Cascadia subduction zone. *Journal of Geophysical Research: Solid Earth*, 121(9):6658–6674, 2016. doi: 10.1002/2016JB013314.
- Michel, S., Gualandi, A. & Avouac, JP. Interseismic Coupling and Slow Slip Events on the Cascadia Megathrust. *Pure Appl. Geophys.* **176**, 3867–3891 (2019). https://doi.org/10.1007/s00024-018-1991-x
- Mousavi, S.M.; Beroza, G.C. Deep-learning seismology. Science 2022, 377, eabm4470.
- Niu, Y., Rebischung, P., Li, M. et al. Temporal spectrum of spatial correlations between GNSS station position time series. J Geod 97, 12 (2023). https://doi.org/10.1007/s00190-023-01703-7
- Nuyen, C. P., & Schmidt, D. A. (2021). Filling the gap in Cascadia: The emergence of low-amplitude long-term slow slip. Geochemistry, Geophysics, Geosystems, 22, e2020GC009477. https://doi.org/10.1029/2020GC009477.
- Obara K., Kato A., ,Connecting slow earthquakes to huge earthquakes. *Science* **353**,253-257(2016).DOI:10.1126/science.aaf1512.
- Ohnaka, M. (1992). Earthquake source nucleation: A physical model for short-term precursors. *Tectonophysics*, *211*(1–4), 149–178. https://doi.org/10.1016/0040-1951(92)90057-d.
- Pintori, F., Serpelloni, E., and Gualandi, A.: Common-mode signals and vertical velocities in the greater Alpine area from GNSS data, Solid Earth, 13, 1541–1567, https://doi.org/10.5194/se-13-1541-2022, 2022.
- Rebischung P, Altamimi Z, Ray J et al (2016) The IGS contribution to ITRF2014. J Geodesy 90(7):611–630. https://doi.org/10.1007/ s00190- 016- 0897- 6.
- Riel, B., M. Simons, P. Agram, and Z. Zhan (2014), Detecting transient signals in geodetic time series using sparse estimation techniques,*J. Geophys. Res. Solid Earth*, *119*, 5140–5160, doi:10.1002/2014JB011077.
- Robinson J., Jegelka S., and Sra S., “Strength from Weakness: Fast Learning Using Weak Supervision” (Preprint, submitted February 19, 2020), *https://arxiv.org/abs/2002.08483*.

- Ronneberger, O., Fischer, P., Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. MICCAI 2015. Lecture Notes in Computer Science(), vol 9351. Springer, Cham. https://doi.org/10.1007/978-3-319-24574-4_28.
- Rousset,B., Bürgmann,R., and Campillo,M., Slow Slip Events In the roots of the San Andreas fault. Science advances, 5(2), 2019. doi: 10.1126/sciadv.aav3274.
- Santamaría-Gómez A, Bouin M, Collilieux X, Wöppelmann G (2011) Correlated errors in GPS position time series: implications for velocity estimates. J Geophys Res 116(B1):B01405. doi:10.1029/ 2010jb007701.
- Satirapod C. & Rizos C. (2005) Multipath Mitigation by Wavelet Analysis for GPS Base Station Applications, Survey Review, 38:295, 2-10, doi: 10.1179/sre.2005.38.295.2.
- Smith KD, von Seggern D, Blewitt G et al (2004) Evidence for deep magma injection beneath Lake Tahoe, Nevada-California. Science 305:1277–1280.
- Steigenberger, P., M. Rothacher, M. Fritsche, A. Rülke, and R. Dietrich (2009b), Quality of reprocessed GPS satellite orbits, *J. Geod.*, 83, 241– 248, doi:10.1007/s00190-008-0228-7.
- Thomas, A., Inbal, A., Searcy, J., Shelly, D., and Bürgmann, R. Identification of Low-Frequency Earthquakes on the San Andreas Fault With Deep Learning. Geophysical Research Letters, 48(13): 2021 093157, 2021. doi: 10.1029/2021GL093157.
- Thomas, A., Melgar, D., Dybing, S. N., & Searcy, J. R. (2023). Deep learning for denoising High-Rate Global Navigation Satellite System data. Seismica, 2(1). https://doi.org/10.26443/seismica.v2i1.240.
- Tian Y, Shen Z-K (2016) Extracting the regional common-mode com- ponent of GPS station position time series from dense continuous network. J Geophys Res Solid Earth 121:1080–1096. https://doi. org/10.1002/2015JB012253.
- Tregoning, P., and C. Watson (2009), Atmospheric effects and spurious signals in GPS analyses, J. Geophys. Res., 114, B09403, doi:10.1029/2009JB006344.
- Xue, X., & Freymueller, J. T. (2023). Machine learning for single-station detection of transient deformation in GPS time series with a case study of Cascadia slow slip. Journal of Geophysical Research: Solid Earth, 128, e2022JB024859. https://doi. org/10.1029/2022JB024859.
- Yamaga, N., & Mitsui, Y. (2019). Machine learning approach to characterize the postseismic deformation of the 2011 tohoku-oki earthquake based on recurrent neural network. Geophysical Research Letters, 46, 11886–11892. https://doi.org/10.1029/2019GL084578.
- Yuzariyadi, M. & Heki, K. Enhancement of interplate coupling in adjacent segments after recent megathrust earthquakes. Tectonoph https://doi.org/ 10.1016/j.tecto.2021.228719 (2021).
- Wallace, L. M., and J. Beavan (2010), Diverse slow slip behavior at the Hikurangi subduction margin, New Zealand, *J. Geophys. Res.*, 115, B12402, doi:10.1029/2010JB007717.
- Wallace LM, Barnes P, Beavan RJ, Van Dissen RJ, Litchfield NJ, et al. 2012a. The kinematics of a transition from subduction to strike-slip: an example from the central New Zealand plate boundary. *J. Geophys. Res.* 117(B2):B02405.
- Wallace LM, Kaneko Y, Hreinsdóttir S, Hamling I, Peng Z, et al. 2017. Large-scale dynamic triggering of shallow slow slip enhanced by overlying sedimentary wedge. *Nat. Geosci.* 10:765–70.
- Wallace, L. M., Hreinsdóttir, S., Ellis, S., Hamling, I., D'Anastasio, E., & Denys, P. (2018). Triggered slow slip and afterslip on the southern Hikurangi subduction zone following the Kaikōura earthquake. *Geophysical Research Letters*, 45, 4710–4718. https://doi.org/10.1002/2018GL077385.
- Wallace, L. M. (2020). Slow slip events in New Zealand. *Annual Review of Earth and Planetary Sciences*, *48*(1), 175–203. https://doi.org/10.1146/ annurev-earth-071719-055104.
- Wessel, P., Luis, J., Uieda, L., Scharroo, R., Wobbe, F., Smith, W., & Tian, D. (2019). The generic mapping tools version 6. *Geochemistry, Geophysics, Geosystems*, 20(11), 5556–5564. https://doi.org/10.1029/2019gc008515.
- Williams S, Bock Y, Fang P et al (2004) Error analysis of continuous GPS position time series. J Geophys Res: Solid Earth. https://doi. org/10.1029/2003JB002741.
- Wdowinski, S., Bock, Y., Zhang, J., Fang, P., & Genrich, J. (1997). Southern California permanent GPS geodetic array: Spatial filtering of daily positions for estimating coseismic and postseismic displacements induced by the 1992 Landers earthquake. Journal of Geophysical Research, 102(B8), 18,057–18,070.
- Zhu, W., Mousavi, S. M., & Beroza, G. C. (2019). Seismic signal denoising and decomposition using deep neural networks. IEEE Transactions on Geoscience and Remote Sensing, 57(11), 9476– 9488. https://doi.org/10.1109/tgrs.2019.2926772.

- Zhu W., Beroza G.C., PhaseNet: a deep-neural-network-based seismic arrival-time picking method, Geophysical Journal International, Volume 216, Issue 1, January 2019, Pages 261–273, https://doi.org/10.1093/gji/ggy423.

Supporting Information for

Denoising Daily Displacement GNSS-Time series using Deep Neural Networks In a Near Real-Time Framing: a Single-Station Method

Mastella G.[1], Bedford J.[2], Corbi F.[3], Funiciello F.[4]

[1] Sapienza University of Rome, Earth Sciences, Italy.

[2] Institut für Geologie, Mineralogie und Geophysik, Ruhr-Universität Bochum, Bochum, 44801, Germany.

[3] Istituto di Geologia Ambientale e Geoingegneria – CNR c/o Dipartimento di Scienze della Terra, Sapienza Università di Roma, Rome, Italy.

[4] Università "Roma TRE,", Rome, Italy, Dip Scienze, Laboratory of Experimental Tectonics.

***Corresponding author: Giacomo Mastella (giacomo.mastella@uniroma1.it)**

Text S1. Time-Series Decomposition with GrAtSiD

GrAtSiD is a regression algorithm introduced by Bedford and Devis, 2018 which aims to decompose GNSS time-series signals. The algorithm uses a greedy approach to determine the minimum number of basis functions required to accurately model a given time series by evaluating the fit improvement of adding additional basis functions in a sequential manner. GrAtSiD incorporates specific basis functions in its solution, including Fourier functions for annual and semi-annual oscillations, a first-order polynomial term, and steps induced by hardware changes or when the station is within a radial cut-off R distance from a known earthquake. This cut-off distance is determined by the equation R=100.5M−0.8, where M represents the earthquake magnitude according to the approach adopted by the Nevada Geodetic Laboratory [Blewitt et al., 2018].
The regression model, *x*(*t*), of GrAtSiD is formulated as follows:

$$y(t) = mt + c + \sum_{k=1}^{n_k} \left[s_k sin(\omega_k t) + c_k cos(\omega_k t) \right] + \sum_{j=1}^{n_j} H(t - t_j) + \sum_{r=1}^{n_r} \sum_{i=1}^{n_i} A_{ri} \left(1 - e^{\frac{-t - t_r}{T_i}} \right) + \xi(t)$$

(**Eq.1**)

where t denotes time, while m and d represent the coefficients of the first-order polynomial. Annual and semiannual seasonal oscillations are captured with the sine and cosine functions, with nk = 2 (comprising coefficients sk and ck). The Heaviside functions H(t − tj) are utilized as basis functions, either predefined or automatically determined using GrAtSiD (with bj denoting the coefficients for events at time tj). The double summation term accounts for r multi-transients, which are composed of exponential functions starting at tr. Each multi-transient is characterized by Ti decay constants. These multi-transients, which encompass a combination of decay functions (in this case, exponential decay), contribute to diverse signal shapes due to the absence of constraints on the coefficients (Ai). By employing one or two multi-transients in series, a wide range of signal shapes can be generated. Notably, the multi-transient is chosen as a sparse basis function in the regression process due to its versatility in capturing transient signals of varying durations. The residual, denoted as ξ(t), represents the deviation from the modeled trajectory. For comprehensive details of the algorithm, we refer to the GrAtSiD methods paper [Bedford and Devis, 2018].

Text S2. From time-series to spectrogram and back to time-series

We represent displacement residuals time-series in the time-frequency domain by applying the Short-time Fourier transform (STFT). This operation computes a spectrogram of the input time signal by performing consecutive discrete Fourier transforms (DFT) over short overlapping windows. The com-

puted spectrogram highlights the change of nonstationary signal's frequency content over time. We indeed notice that residuals time series sometimes have time-dependent frequency content. Without speculating in detail on their nature, these variations may arise both from an underfitting of the original displacement time series as well as to some specific signals that derive from noise sources that the trajectory model is not able to fit. These time-frequency representations of residuals are the input of the GAN model that is trained to generate synthetic spectrograms. The synthetically created spectrograms are subsequently reverted in the time-domain by employing the Griffin-Lim algorithm [Griffin and Lim, 1994]. This algorithm is an iterative method commonly used in audio signal processing to reconstruct a signal from its short-time Fourier transform (STFT) representation. In practice, the Griffin-Lim algorithm involves initializing a complex-valued spectrogram and iteratively applying the inverse STFT (iSTFT) to the magnitude of the spectrogram, while retaining the phase information from the previous iteration. This process is repeated until the reconstructed signal converges to a stable solution. We here use the librosa python package [McFee et al., 2015] to implement the STFT spectrogram derivation and its conversion through the Griffin-Lim method.

Text S3. Performance evaluation of GANs for generating synthetic residuals time series

We utilized various statistical tests to evaluate the quality of the generated data. For the generated samples to be considered reliable, they must meet both the diversity and fidelity criteria. This means that the generated data should have a distribution that encompasses the real data, while also being indistinguishable from it. The first test we applied is the application of the Principal Component Analysis (PCA) to synthetic and real datasets. The PCA is a linear dimensionality reduction algorithm that converts the data into a smaller new coordinate system preserving as much as feasible of the variability of the original data [Tipping and Bishop, 1999]. PCA is often used for image compression, a method exploited to reduce the size of an image while maintaining the highest possible level of the image's clarity. We here fit the PCA decomposition model to real spectrograms employing the first two components. This means that the real spectrograms are converted in a 2-D space represented by the directions of their maximum variance. The obtained dimensionality reduction model is subsequently applied on the synthetic dataset as well as to a randomly generated dataset of white noise spectrograms. The white noise dataset is randomly drawn from a normal (Gaussian) distribution characterized by an average and a standard deviation equal to the average and standard deviation of real residuals spectrograms. **Figure S8** shows the distributions of synthetic, real and white noise datasets for horizontal (**S8A**) and vertical components (**S8B**) in the 2-D coordinate space represented by the two first PCA components of real data. The synthetic and the real data points show an almost perfect overlap, despite the former distribution appearing slightly less spread. On the other hand, the white noise representation is confined in a very small region, demonstrating that real residuals time-series encompass temporal correlations and variations of the frequency content. This visualization provides a qualitative evaluation of the GAN performance satisfying the diversity criteria by showing how closely the distribution of generated spectrograms matches that of the original in the reduced 2-dimensional space.

To quantitatively assess the similarity and fidelity between the original and generated datasets, we trained three post-hoc classification models using 3-layer CNNs in a standard supervised framework [e.g. Yoon, et al., 2019]. These models were designed to differentiate between spectrograms from the original dataset and those generated by the GAN. The first classification model was trained by labeling real data as "real" and white noise as "fake." It was then tested on GAN-generated spectrograms as well as different examples of white noise. Conversely, the second model was trained to label GAN-generated data as "real" and white noise as "fake," and its testing focused on classifying real spectrograms as "real" and different white noise examples as "fake". Lastly, the third CNN model was trained to classify real spectrograms as "real" and generated samples as "fake". To evaluate the performance of each model, we calculated the classification error using held-out test sets. This provided a quantitative measure of the effectiveness of our generative model. We perform the same exercise for the residuals from horizontal and vertical components separately. The results, presented in **Figure S9**, include a confusion matrix for each model. Notably, the first two models achieved high accuracy, from 0.90 to 0.99, indicating their ability to accurately distinguish between white noise spectrograms and real or synthetic spectrograms respectively. Differently, for the last model, the accuracy was 0.5, suggesting that real and GAN-generated spectrograms were indistinguishable. These findings highlight the good capability of our generative model to produce synthetic spectrograms that closely resemble real data. The high accuracy achieved by the first two models demonstrates the effectiveness of our approach in generating realistic synthetic data, while the last model's accuracy indicates the model's ability to generate synthetic samples that are virtually indistinguishable from real data.

Text S4. Networks Hyperparameters optimization: the Hyperband algorithm

To speed up the tuning of hyperparameters, we used the Machine Learning platform Weight and Biases (https://docs.wandb.ai), which includes different sweep optimization strategies. In particular, in this work we rely on the Random Search tuner along with the Hyperband early-stopping algorithm. Hyperband is a powerful and efficient adaptive algorithm that combines elements of random search and successive halving to intelligently allocate resources and explore the hyperparameter space. The key idea behind Hyperband is to strike a balance between exploration and exploitation. It explores a wide range of hyperparameter configurations initially, allowing for diverse exploration. Then, it focuses on the best-performing configurations, exploiting their potential by allocating more resources. In specific, the algorithm begins by randomly sampling a set of hyperparameter configurations and training them for a fixed number of iterations. These configurations are then evaluated based on their performance metric, the loss. Next, Hyperband applies the successive halving technique to focus on the top-performing configurations. It selects a subset of configurations and allocates more training iterations to each of them. This allows for a more thorough evaluation of their potential. The poorly performing configurations are discarded, reducing computational waste. The process of selection and resource allocation continues iteratively, with the number of configurations and iterations progressively decreasing. By selectively retaining and allocating more resources to the most promising configurations, the algorithm converges towards promising regions of the hyperparameter space quickly and effectively.

Text S5. Dealing with data gaps

GNSS displacement time-series are frequently characterized by the presence of data gaps. These gaps can be due to large earthquakes but generally are associated with inconsistently daily measurements of GNSS antennas, for example related to electricity blackouts. Data gaps limit the application of our DL approach, which requires temporal contiguous data to build the sequences of input examples. To fill the majority of data gaps, we interpolate the time-series using a linear spline interpolation assuming a maximum admitted number of temporal contiguous gaps equals five days. This means that if the gap is shorter than five days, then gaps are filled, otherwise not. This interpolation routine is applied before building the inputs and the targets of DL models. When applied to real-time series the interpolation routine is applied causally, namely without looking into the future, but considering only data from the past. This prevents any leakage.

Text S6. Assessing Denoising Performance

To evaluate the effectiveness of our denoising method, we employ various metrics on each GNSS daily displacement time-series from the testing dataset. These metrics provide insights into the performance of the method and its comparison to the trajectory fitting obtained with GrAtSiD. The calculated metrics are as follows:

a) Standard Deviation of Predicted Noise: this metric quantifies the variability of the predicted noise in the denoised time-series.

b) Root Mean Square Error (RMSE) of predicted noise, calculated by taking the square root of the mean of noise values.

c) Relative Root Mean Square Error (RRMSE) of predicted noise: This metric compares the performance of our method to the trajectory fitting achieved with GrAtSiD.

The RRMSE measures the average magnitude of the differences between the noise values predicted by the DL approach and the noise values isolated with GrAtSiD. It is calculated by taking the square root of the mean of the squared differences.

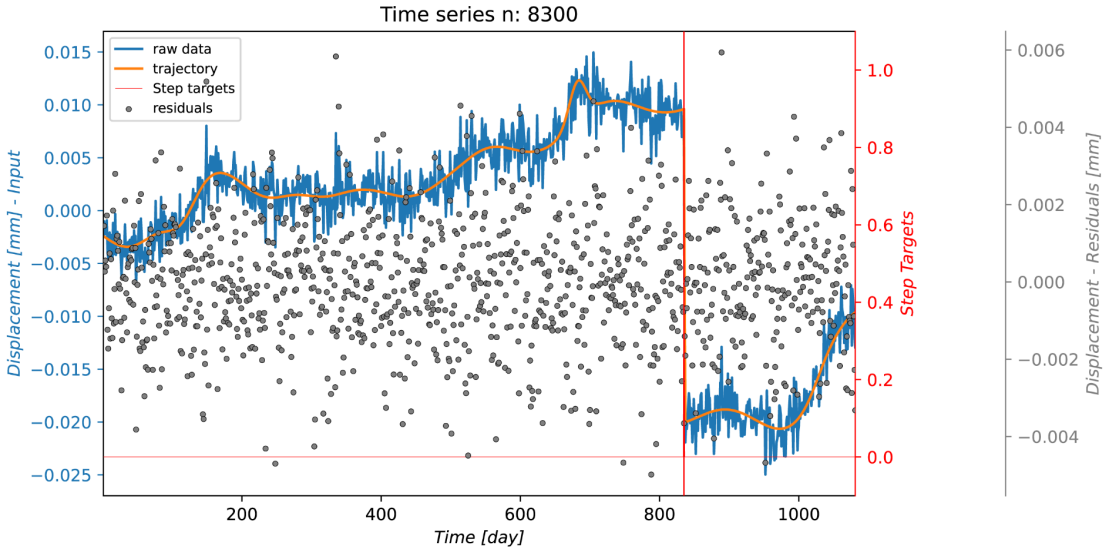

Figure S1. Example of a GNSS displacement synthetic time-series, defined as the sum of a synthetic trajectory (in orange) and synthetic noise (i.e. the residuals, gray scatters). In red the prediction targets of the Step model are depicted while the residuals in gray represent the prediction targets of the Noise model.

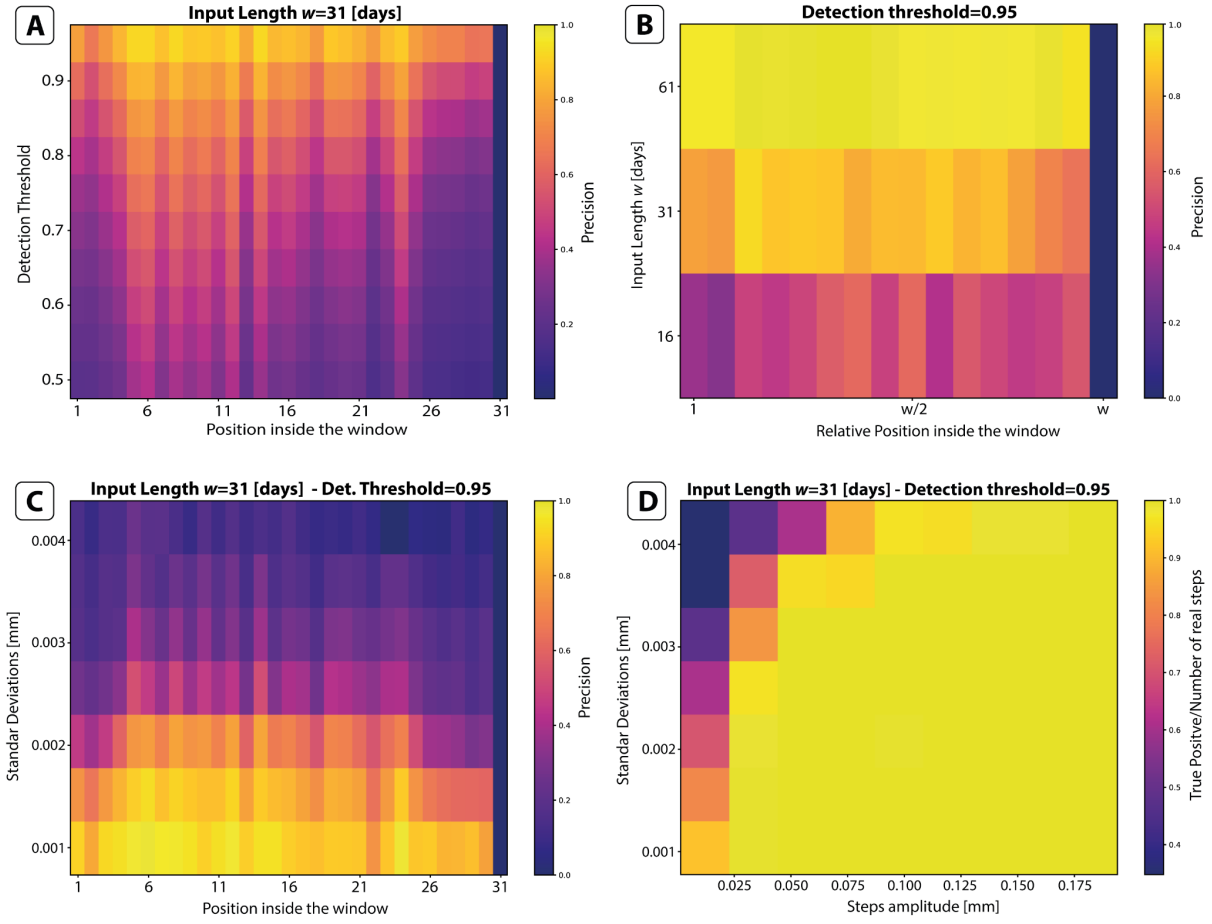

Figure S2. Classification results of the Step model evaluated by using the synthetic time-series testing dataset, for which we exactly know the locations of steps.

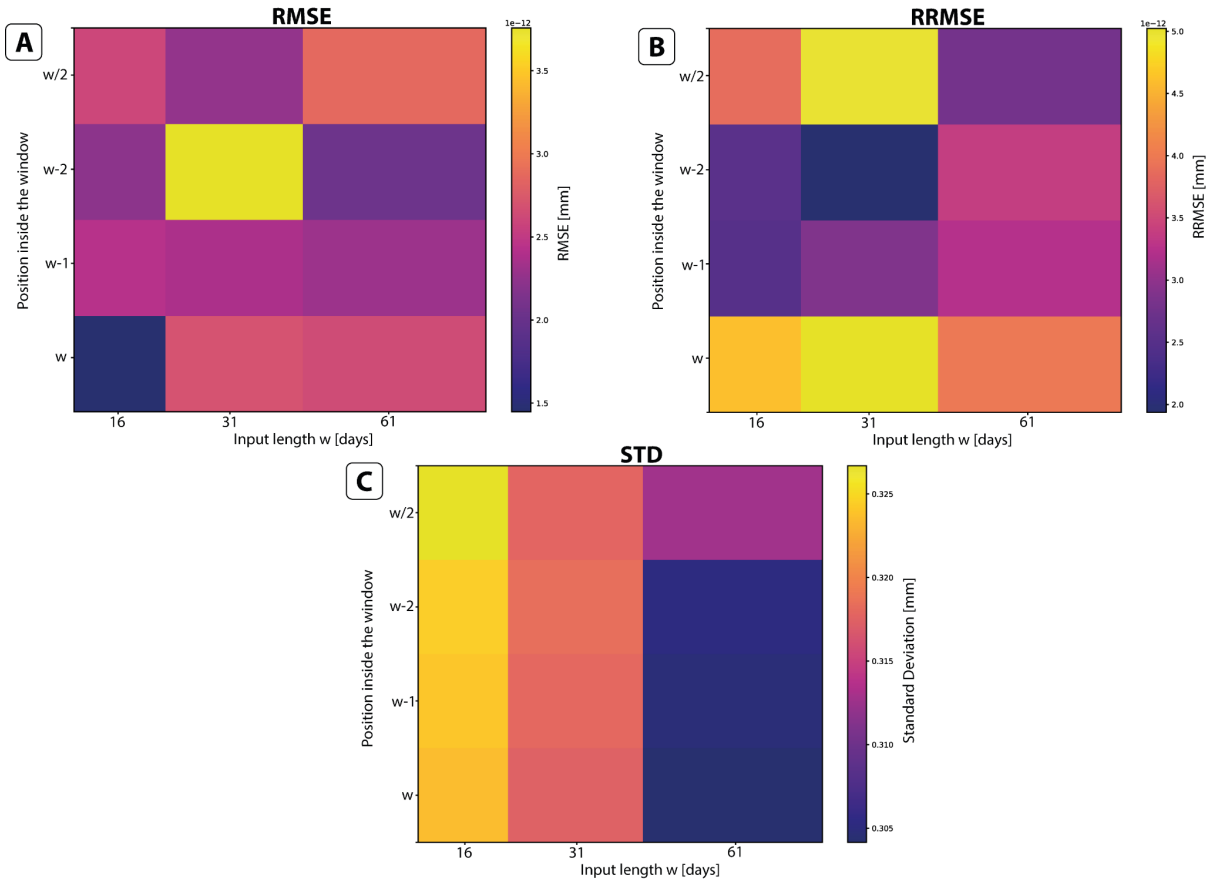

Figure S3. Denoising results of the Noise model evaluated by using the real time-series testing dataset.

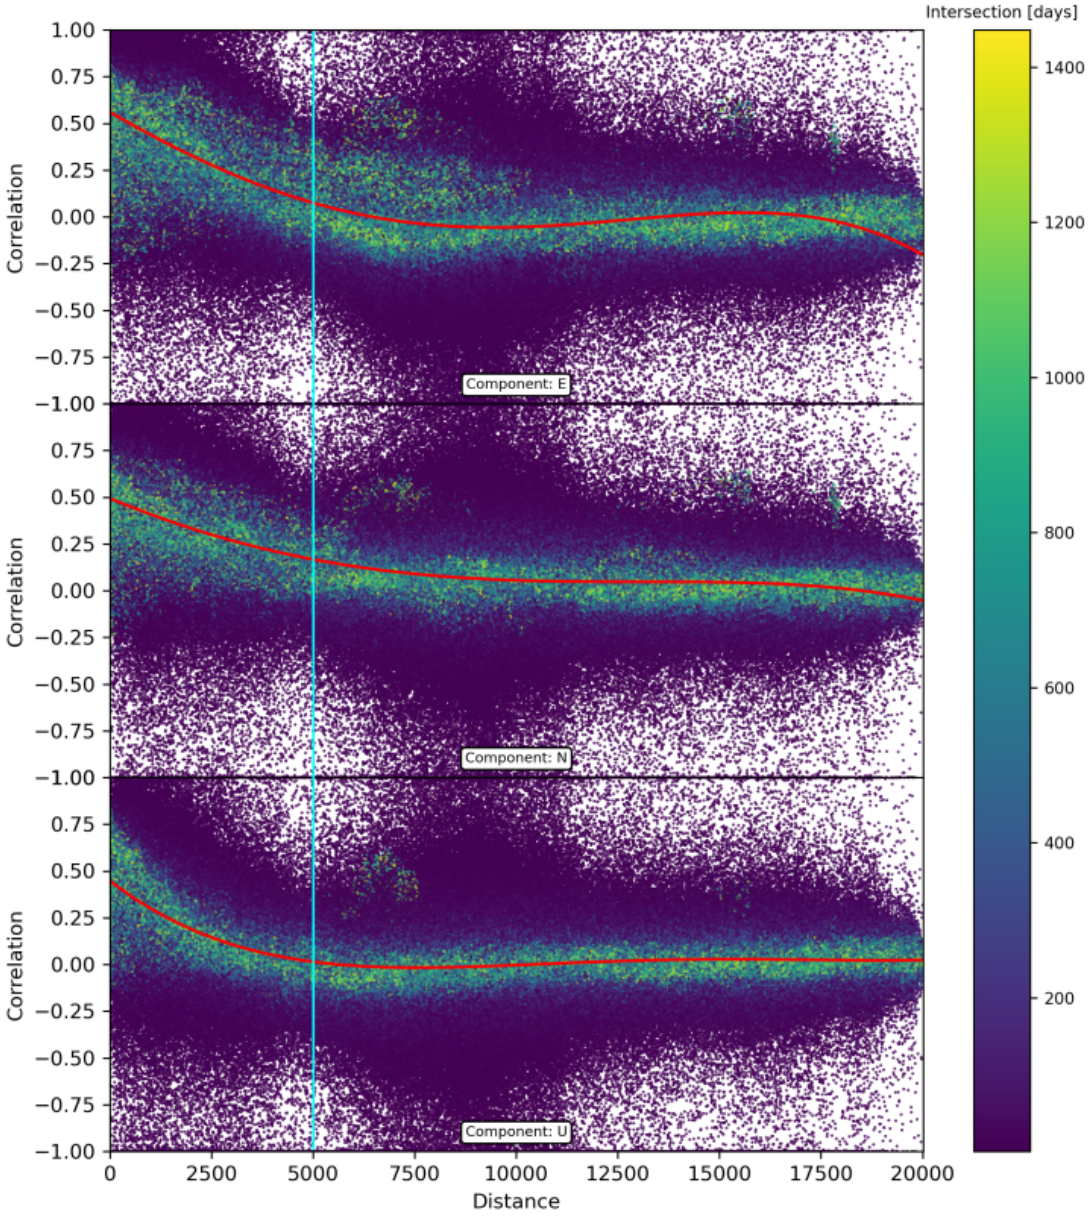

Figure S4. Correlation coefficients between residuals of each pair of GNSS coordinate time-series as a function of distance between sites in km. The red line depicts a third-degree polynomial fitting obtained by incorporating weights into the inversion process based on the number of shared days between each pair of time series (represented by the color of the scatter points). Vertical cyan line depicts the distance threshold considered to select the time-series used to compute the common mode residuals for a given reference time-series.

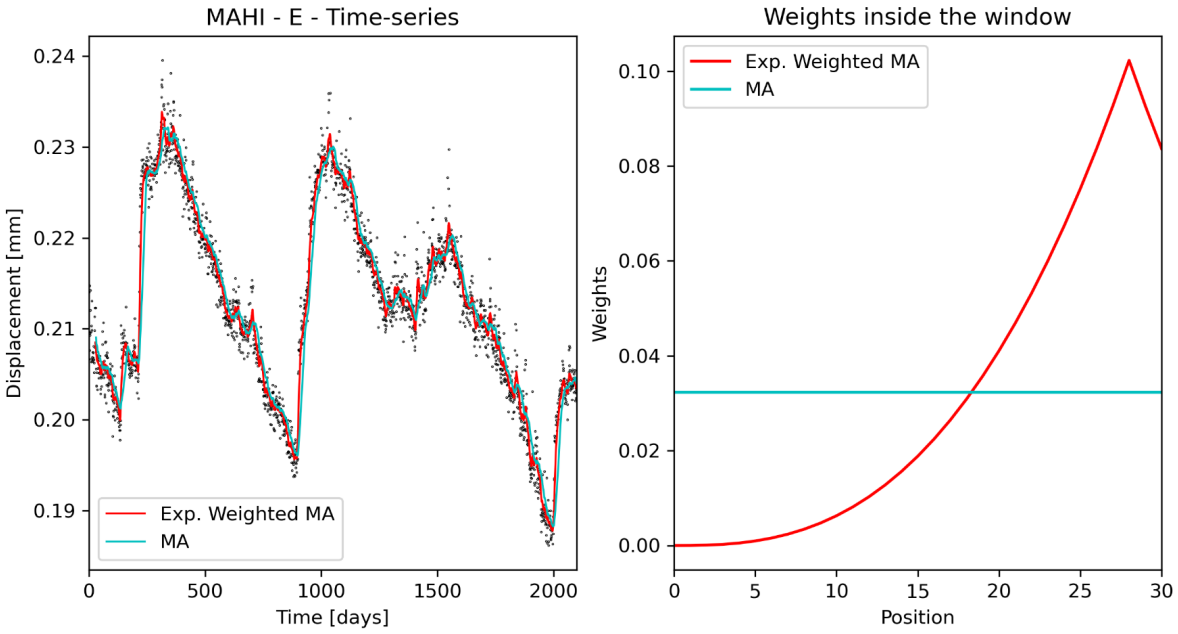

Figure S5. Example of the application of the Exponential Moving Average model using a moving window of 31 days and the maximum weight assigned to the 29th observation.

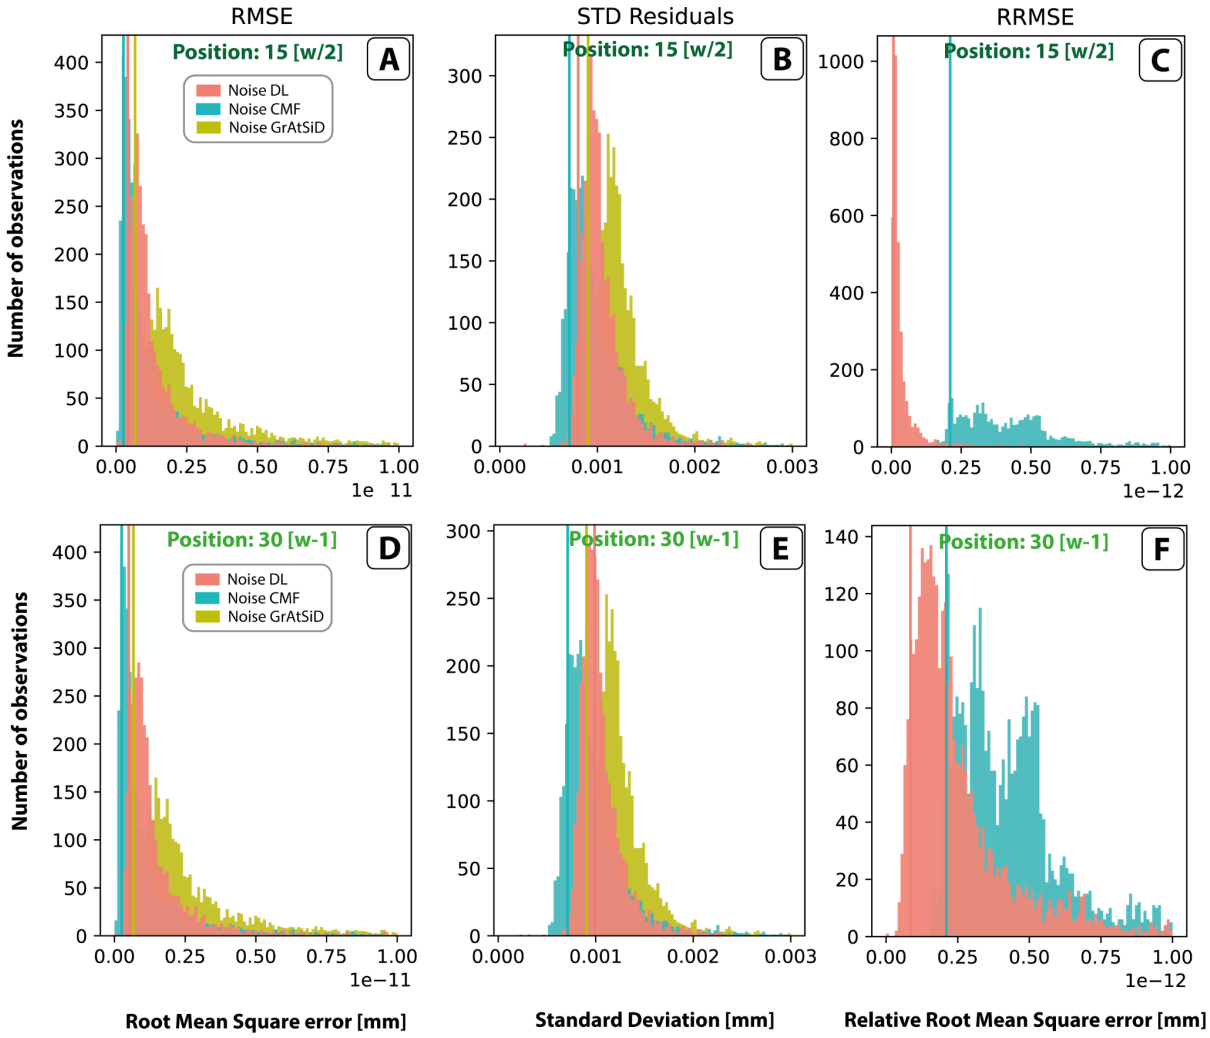

Figure S6. Distribution of the absolute Root Mean Square Error (RMSE), Standard Deviation (STD) and Relative Root Mean Square Error (RRMSE) of the noise predicted with our DL approach (in orange) and the noise recovered with the CMF filter (in cyan). The Relative Root Mean Square Error is calculated comparing the values predicted with these two methods with respect to the noise isolated with GrAtSiD. In yellow, the distribution of the noise isolated with GrAtSiD is reported. DL models are trained with an input length of 31 days, to predict the central point of the window (panels A-C) and the last point of the window (panels D-F).

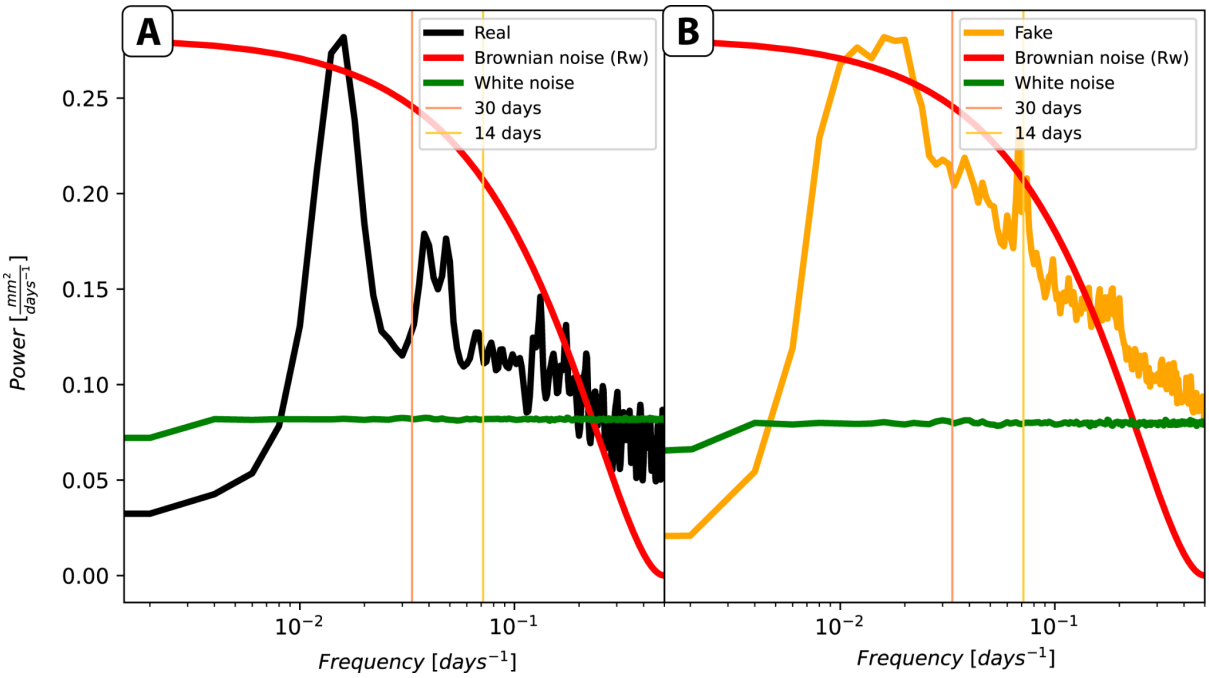

Figure S7. Cumulative Power Spectral Density of the real dataset of residuals time-series (**A**) and of the fake residuals time-series generated with the GAN (**B**). In green it depicts the distribution of a dataset composed by white noise examples while in red the spectral signature of Brownian (random walk - red) noise. White noise and Brownian noise datasets are randomly drawn from a normal (Gaussian) distribution characterized by an average and a standard deviation equal to the average and standard deviation of real residuals.

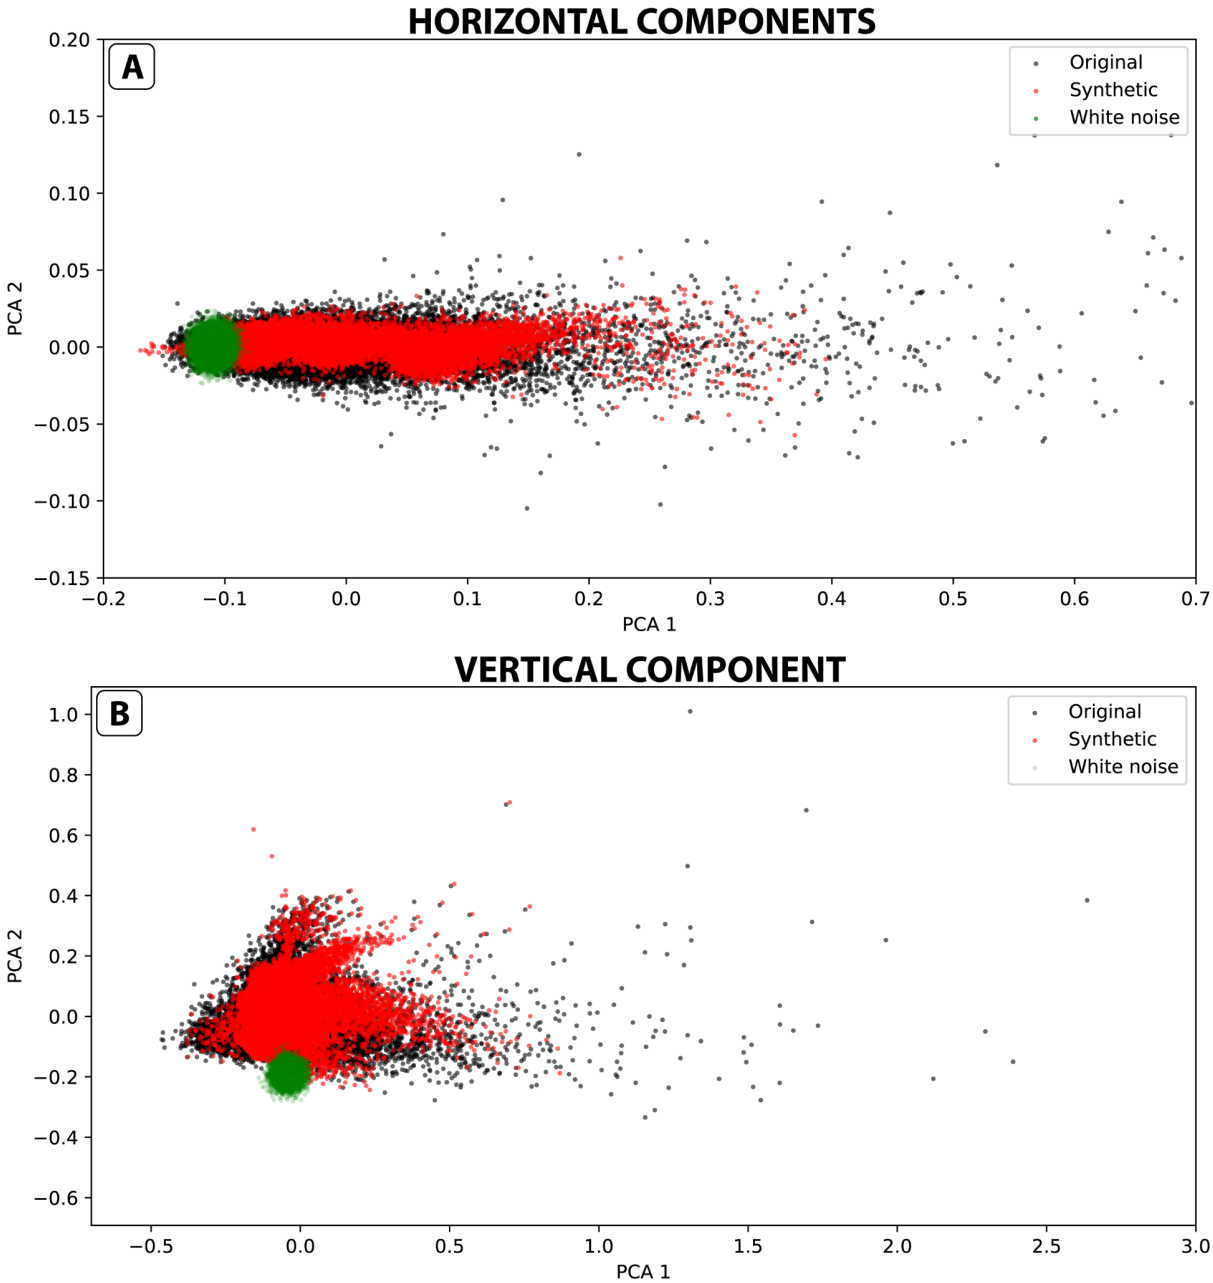

Figure S8. Distributions of synthetic, real and white noise datasets of residuals spectrograms of horizontal components (**A**) and of the vertical component (**B**) in the 2-D coordinate space represented by the two first PCA components of real data.

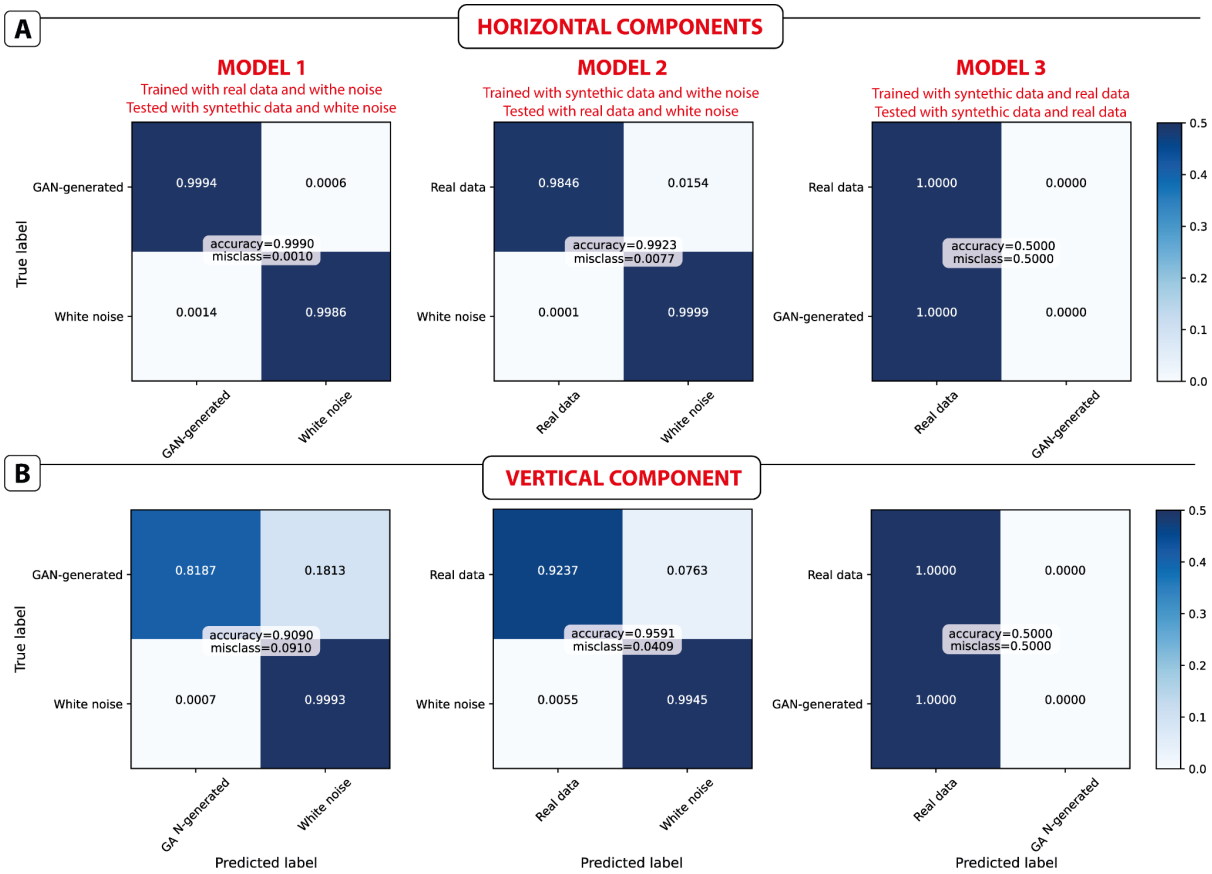

Figure S9. Performances of the classification models trained to assess the similarity and fidelity between the real and GAN-generated spectrograms.

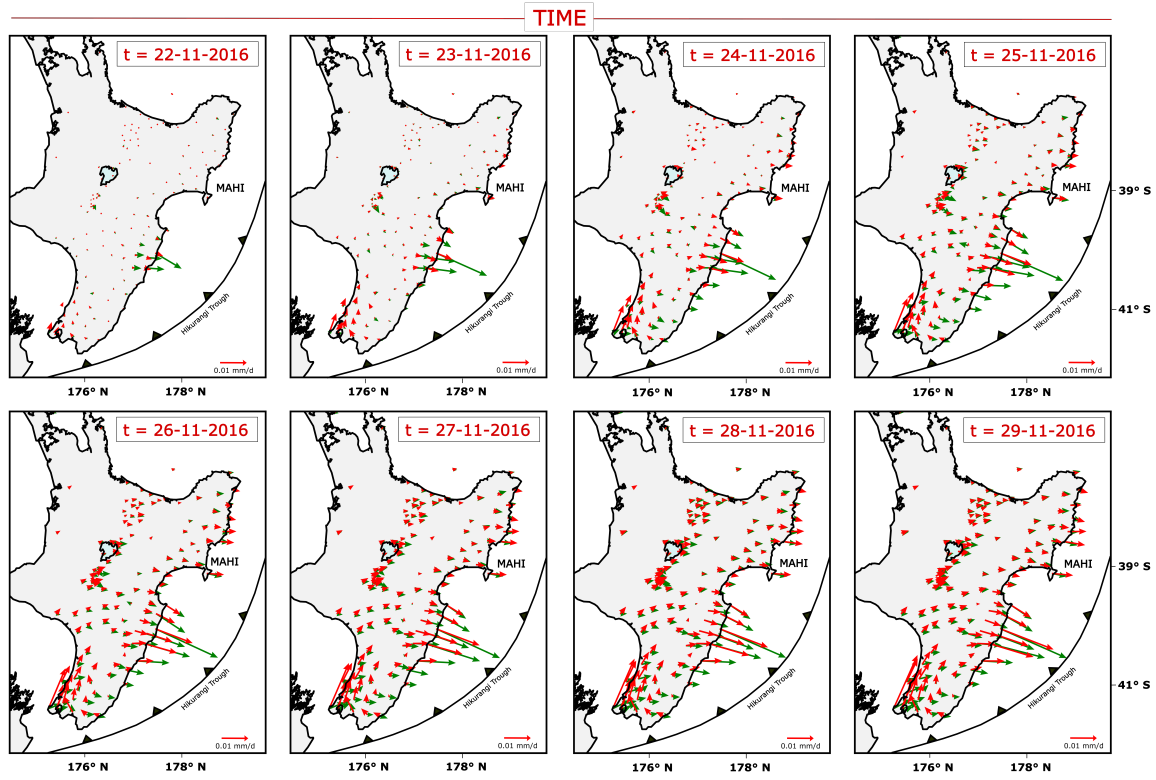

Figure S10. Same as Figure 9 in the main text, but with EMA results represented. Cumulative time-dependent velocities of GNSS sites during the propagation of the slow-slip event at the Hikurangi

subduction zone following the Kaikōura Earthquake as calculated after the application of our DL denoising method (green quivers) and the EMA filter (red quivers).

SI References

- Bedford, J., & Bevis, M. (2018). Greedy automatic signal decomposition and its application to daily GPS time series. *Journal of Geophysical Research: Solid Earth*, 123, 6992– 7003. https://doi.org/10.1029/2017JB014765.
- Blewitt, G., Hammond, W. C. & Kreemer, C. Harnessing the GPS data explosion for interdisciplinary science. *Eos* 99, https://doi.org/10.1029/2018EO104623 (2018).
- Griffin D. W. and Lim J. S., "Signal estimation from modified short-time Fourier transform," IEEE Trans. ASSP, vol.32, no.2, pp.236–243, Apr. 1984.
- Li L., Jamieson K., DeSalvo G., Rostamizadeh A., Talwalkar A., Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization, in Machine Learning Research 18, 2018.
- McFee B., Raffel, C., Liang, D., Ellis, D. P., McVicar, M., Battenberg, E., & Nieto, O. (2015). librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference* (Vol. 8)
- Tipping, M. E., and Bishop, C. M. (1999). "Probabilistic principal component analysis". Journal of the Royal Statistical Society: Series B (Statistical Methodology), 61(3), 611-622.
- Yoon J., Jarrett D., van der Schaar M., "Time-series Generative Adversarial Networks," Neural Information Processing Systems (NeurIPS), 2019.