



HAL
open science

Slow slip detection with deep learning in multi-station raw geodetic time series validated against tremors in Cascadia

Giuseppe Costantino, Sophie Giffard-Roisin, Mathilde Radiguet, Mauro Dalla Mura, David Marsan, Anne Socquet

► **To cite this version:**

Giuseppe Costantino, Sophie Giffard-Roisin, Mathilde Radiguet, Mauro Dalla Mura, David Marsan, et al.. Slow slip detection with deep learning in multi-station raw geodetic time series validated against tremors in Cascadia. 2023. hal-04109528

HAL Id: hal-04109528

<https://hal.science/hal-04109528>

Preprint submitted on 30 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Slow slip detection with deep learning in multi-station raw geodetic time series validated against tremors in Cascadia

Giuseppe Costantino^{1*}, Sophie Giffard-Roisin¹, Mathilde Radiguet¹, Mauro Dalla Mura^{2,3}, David Marsan¹ and Anne Socquet¹

¹ISTerre, Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, IRD, Univ. Gustave Eiffel, Grenoble, 38000, France.

²GIPSA-lab, Univ. Grenoble Alpes, CNRS, Grenoble INP, Grenoble, 38000, France.

³Institut Universitaire de France (IUF), France.

*Corresponding author(s). E-mail(s):

giuseppe.costantino@univ-grenoble-alpes.fr;

Abstract

Slow slip events (SSEs) originate from a slow slippage on faults that lasts from a few days to years. A systematic and complete mapping of SSEs is key to characterizing the slip spectrum and understanding its link with coeval seismological signals. Yet, SSE catalogues are sparse and usually remain limited to the largest events, because the deformation transients are often concealed in the noise of the geodetic data. Here we present the first multi-station deep learning SSE detector applied blindly to multiple raw geodetic time series. Its power lies in an ultra-realistic synthetic training set, and in the combination of convolutional and attention-based neural networks. Applied to real data in Cascadia over the period 2007-2022, it detects 78 SSEs, that compare well to existing independent benchmarks: 87.5% of previously catalogued SSEs are retrieved, each detection falling within a peak of tremor activity. Our method also provides useful proxies on the SSE duration and may help illuminate relationships between tremor chatter and the nucleation of the slow rupture. We find an average day-long time lag between the slow deformation and the tremor chatter both at a global- and local-temporal scale, suggesting that slow slip may drive the rupture of nearby small asperities.

Keywords: deep learning, slow slip events, transient, deformation, GPS, GNSS, geodesy, deformation, tremor, earthquakes, subduction, Cascadia, multi-station, classification, attention-based, neural network, transformer

1 Introduction

2 Slow slip events (SSEs) generate episodic deformation that lasts from a few
3 days to years. Like earthquakes, they originate from slip on faults but, unlike
4 them, do not radiate energetic seismic waves. In the mid-1990s, Global Navi-
5 gation Satellite System (GNSS) networks started to continuously monitor the
6 ground displacement, providing evidence that SSEs are a major mechanism
7 responsible for the release of stress in plate boundaries, as a complement to
8 seismic rupture [1–5]. This constituted a change of paradigm for the under-
9 standing of the earthquake cycle and of the mechanics of the fault interface.
10 Twenty years later, the characterization of the full slip spectrum and the
11 understanding of the link between slow slip and the associated seismological
12 signals are hindered by our capacity to detect slow slip events in a systematic
13 manner, more particularly those of low magnitude (typically lower than M_w
14 6), even though a systematic and complete mapping of SSEs on faults is key
15 for understanding the complex physical interactions between slow aseismic slip
16 and earthquakes. Indeed, the small deformation transients associated with an
17 SSE are often concealed in the noise [6, 7], making it difficult to precisely char-
18 acterize the slip spectrum and provide fruitful insights into the fault mechanics
19 [4, 8, 9]. Studies dealing with the detection and analysis of SSEs often rely on
20 dedicated signal analysis, involving visual inspection of the data, data selec-
21 tion, denoising, filtering, geodetic expertise, dedicated modeling methods with
22 a fine-tuning of the parameters, and also often complementary data such as
23 tremor or LFE catalogs [7, 10–13].

24 The development of in-situ geophysical monitoring generates nowadays
25 huge data sets, and machine learning techniques have been largely assimilated
26 and used by the seismological community to improve earthquake detection and
27 characterization [5, 14–16], generating catalogs with unprecedented high qual-
28 ity [17, 18] and knowledge shifts [19, 20]. However, up to now, such techniques
29 could not be successfully applied to the analysis of geodetic data and slow slip
30 event detection because of two main reasons: (1) too few true labels exist to
31 train machine learning-based methods, which we tackled by generating a real-
32 istic synthetic training data set, (2) the signal-to-noise ratio is extremely low
33 in geodetic data [21, 22], meaning that we are at the limit of detection capac-
34 ity. One possibility is to first pre-process the signals (via denoising, filtering,
35 detrending), but this is at the cost of possibly corrupting the data. Instead,
36 in this work, we assume that the information is already present in the raw
37 time series and that our deep learning model should be able to learn the noise
38 signature, and therefore to separate the noise from the relevant information
39 (here, slow slip events). In order to develop an end-to-end model capable of

40 dealing with raw geodetic measurements, it is necessary, on one hand, to set up
41 advanced methods to generate realistic noise, taking into account the spatial
42 correlation between stations as well as the large number of data gaps present in
43 the GNSS time series. On the other hand, it involves developing a specific deep
44 learning model able to treat multiple stations simultaneously, using a relevant
45 spatial stacking of the signals (driven by our physics-based knowledge of the
46 slow slip events) in addition to a temporal analysis. We address these two major
47 drawbacks in our new approach and present SSEgenerator and SSEdetector,
48 to our knowledge the first end-to-end deep learning-based detector, combin-
49 ing the spatiotemporal generation of synthetic GNSS time series containing
50 modeled slow deformation (SSEgenerator), and a Convolutional Neural Net-
51 work (CNN) and a Transformer neural network with an attention mechanism
52 (SSEdetector), that proves effective in systematically detecting slow slip events
53 in raw GNSS position time series from a large geodetic network containing
54 more than 100 stations, both on synthetic and on real data.

55 Results

56 SSEgenerator: construction of the synthetic dataset

57 We choose the Cascadia subduction zone as the target region because: (1) a
58 link between slow deformation and tremor activity has been assessed [23] and a
59 high-quality tremor catalog is available [24]; (2) a preliminary catalog of SSEs
60 has recently been proposed during the period 2007-2017 with conventional
61 methods [11]. This proposed catalog will be used for comparison and baseline
62 for our results, which are expected to provide a more comprehensive catalog
63 that will better show the link between slow deformation and tremors.

64 To overcome the scarcity of catalogued SSEs, we train SSEdetector on
65 synthetic data, consisting of simulated sets of geodetic time series for the full
66 station network. Each set of signals (60 days and 135 stations) is considered
67 as a single sample. In order to be able to detect SSEs in real raw time series,
68 several characteristics need to be present in these synthetics. First, they must
69 contain a wide range of realistic background signals at the level of the GNSS
70 network, *i.e.* spatially and temporally-correlated realistic noise time series. On
71 the other hand, while half of the samples (*negative* samples) will only consist of
72 background noise, the other half must also include an SSE signal. For this, we
73 modeled SSEs signals that are realistic enough compared to real transients of
74 aseismic deformation. Finally, the synthetics should also carry realistic missing
75 data recordings, as many GNSS stations have data gaps in practice.

76 First, we thus generated ultra-realistic synthetic time series, that repro-
77 duce the spatial and temporal correlated noise of the data acquired by the
78 GNSS network, based on the method developed by Costantino et al. [22]. This
79 database of 60,000 synthetic time series was derived from real geodetic time
80 series (details in Methods). We select data in the periods 2007-2014 and 2018-
81 2022 as sources for the noise generation, while we keep data in the period
82 2014-2017 as an independent test data set (details in Methods).

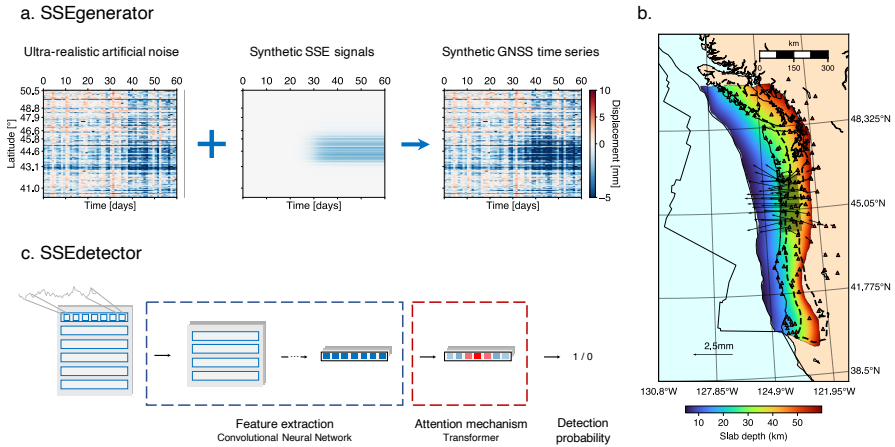
4 *Slow slip detection with deep learning in multi-station geodetic time series*

Fig. 1 Schematic architecture of SSEgenerator and SSEdetector. (a) Overview of the synthetic data generation (SSEgenerator). In the matrix, each row represents the GNSS position time series for a given station, color-coded by the value of the position. The 135 GNSS stations considered in this study are here shown sorted by latitude. The synthetic static displacement model (cf. (b) panel), due to a M_w 6.5 event, at each station is convolved to a sigmoid to model the SSE transient, and is added to the ultra-realistic artificial noise to build synthetic GNSS time series. (b) Location of the GNSS stations of MAGNET network used in this study (red triangles). An example of synthetic dislocation is represented by the black rectangle, with arrows showing the modeled static displacement field. The heatmap indicates the locations of the synthetic ruptures considered in this study, color-coded by the slab depth. The dashed black contour represents the tremor locations from the PNSN catalog. (c) High-level representation of the architecture of SSEdetector. Input GNSS time series are first convolved in the time domain, then a higher weight is assigned to slow slip transients and a probability value is provided depending on whether slow deformation has been found in the data.

83 In order to create the *positive* samples (time series containing an SSE),
 84 we modeled 30,000 dislocations (approximated as a point source) distributed
 85 along the Cascadia subduction interface (see Figure 1(b)) following the slab2
 86 geometry [25] (detailed procedure in Methods). The focal mechanism of the
 87 synthetic ruptures approximates a thrust, with rake angle following a uniform
 88 distribution (from 75 to 100°) and strike and dip defined by the geometry of
 89 the slab. The magnitude of the synthetic SSEs is drawn from a uniform prob-
 90 ability distribution (from M_w 6 to 7). Their depths follow the slab geometry
 91 and are taken down to 60 km, with further variability of ± 10 km. We fur-
 92 ther assign each event a realistic stress drop modeled from published scaling
 93 laws [26]. We use the Okada dislocation model [27] to compute static displace-
 94 ments at each real GNSS station. We scaled the amplitudes of synthetic SSE
 95 signals, modeled as sigmoidal-shaped transients, with a duration following a
 96 uniform distribution (from 10 to 30 days). Eventually, we compute a database
 97 of 30,000 synthetic SSE transients, where the amplitude was added to the
 98 positive samples (placed in the middle of the 60-day window).

99 The synthetic data set is thus made of 60,000 samples and labels, equally
100 split into pure noise (labeled as 0) and signal (labeled as 1) with different
101 nuances of signal-to-noise ratio, resulting both from different station noise
102 levels and differences in magnitude and location, so that the deep learning
103 method effectively learns to detect a variety of slow deformation transients
104 from the background noise. The data set is further split into three independent
105 training (60%), validation (20%) and test (20%) sets, with the latter being
106 used after the training phase only.

107 **SSEdetector: high-level architecture**

108 SSEdetector is a deep neural network made of a CNN [28] and a Transformer
109 network [29] that are sequentially connected (detailed structure in Methods).
110 We constructed the CNN to be a deep spatial-temporal encoder, that behaves
111 as feature extractor. The structure of the encoder is a deep cascade of 1-
112 dimensional temporal convolutional block sequences and spatial pooling layers.
113 The depth of the feature extractor guarantees: (1) a high expressive power,
114 *i.e.*, detailed low-level spatiotemporal features, (2) robustness to data gaps,
115 since their propagation is kept limited to the first layers thanks to a cascade of
116 pooling operators, and (3) limited overfit of the model on the station patterns,
117 thanks to the spatial pooling operation. The decisive component of our archi-
118 tecture is the Transformer network, placed right after the deep CNN encoder.
119 The role of the Transformer is to apply a temporal self-attention mechanism
120 to the features computed by the CNN. As humans, we instinctively focus just
121 on particular fragments of data when looking for any specific patterns. We
122 wanted to replicate such a behavior in our methodology, leading to a network
123 able to enhance crucial portions of the data and neglect the irrelevant ones.
124 This is done by assigning a weight to the data, with those weights being learnt
125 from the data itself. As a result, our Transformer has learnt (1) to precisely
126 identify the timing of the aseismic deformation transients in the geodetic time
127 series and (2) to focus on it by assigning a weight close to zero to the rest of the
128 time window. We further guide the process of finding slow deformation transi-
129 ents through a specific supervised-learning classification process. First, the
130 disclosed outputs of the Transformer are averaged and passed through a sig-
131 moid activation function. The output values are a detection probability lying
132 in the (0,1) range and can be further interpreted as a confidence measure of
133 the method. Second, we train SSEdetector by minimizing the binary cross-
134 entropy loss between the target and the predicted labels (details in Methods).
135 The combination of the two strategies allows SSEdetector to be successfully
136 applied in a real context because: (1) we can run our detector on 1-day-shift
137 windows of real data and collect an output value for each day used to build
138 a temporal probability curve, (2) thanks to the Transformer neural network,
139 such a curve will be smooth and the value of probability will gradually increase
140 in time as SSEdetector identifies slow deformation in the geodetic data.

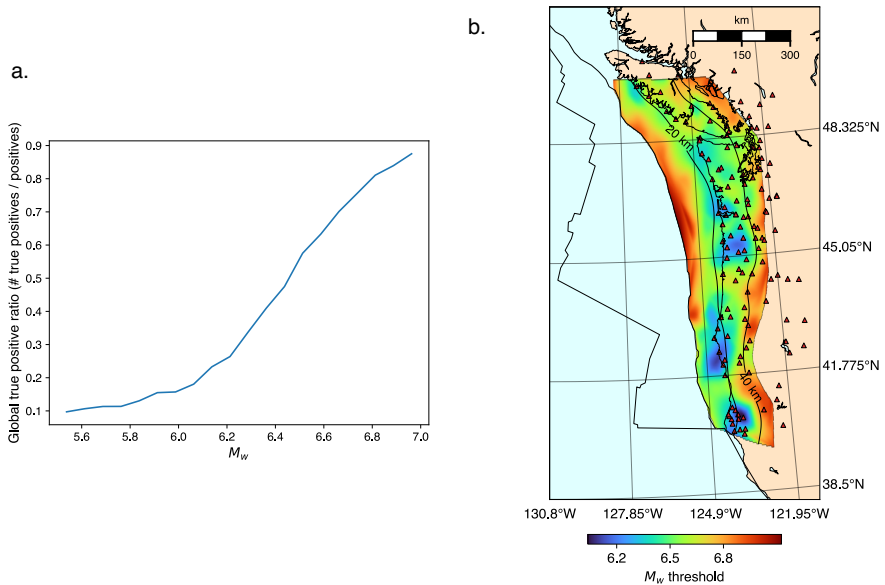


Fig. 2 Performance of SSEdetector on synthetic data. (a) The blue curve represents the true positive rate (probability that an actual positive will test positive), computed on synthetic data, as a function of the magnitude. (b) Map showing the spatial distribution of the magnitude threshold for reliable detection, computed, for each spatial bin, as the minimum magnitude corresponding to a true positive rate value of 0.7.

141 Application to the synthetic test set: detection threshold

142 We test SSEdetector against unseen synthetic samples and we analyze the
 143 results quantitatively. We generate test synthetic samples from GNSS data in
 144 the period 2018-2022 to limit the influence of data gaps (details in Method).
 145 We obtain a measure of the sensitivity of our model by computing the true
 146 positive rate (TPR, probability that an actual positive will test positive) as
 147 a function of the magnitude. On a global scale, the sensitivity is increasing
 148 with the signal-to-noise ratio (SNR), which also shows that it exists an SNR
 149 threshold limit for any SSE detection. This threshold is mainly linked to the
 150 magnitude of the event, rather than the moment rate. Thus, the ability of
 151 SSE detection is mostly influenced by the signal-to-noise ratio rather than the
 152 event duration (cf. Supplementary Figure 1). We compute the sensitivity as a
 153 function of the spatial coordinates of the SSE, by deriving a synthetic proxy as
 154 the magnitude threshold under which the TPR is smaller than 0.7 on a spatial
 155 neighborhood of approximately 50 km. We can see from Figure 2(b) that the
 156 detection power is related to the density of stations in the GNSS network, as
 157 well as to the distance between the rupture and the nearest station, and the
 158 rupture depth. When the density of GNSS stations is not high enough, our
 159 resolution power decreases as well as the reliability of the prediction. In those
 160 cases, we can only detect high-magnitude SSEs. This is also the case on the
 161 eastern side of the targeted region where the SSE sources are deeper because of

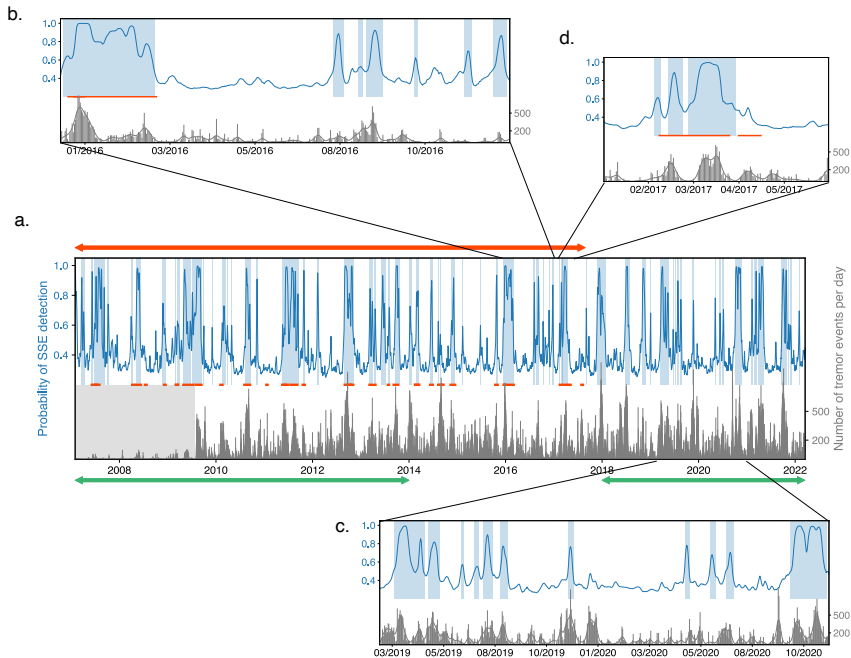


Fig. 3 Overview of the performance of SSEdetector on real raw GNSS time series. The blue curves show the probability of detecting a slow slip event (output by SSEdetector) in 60-day sliding windows centered on a given date. Grey bars represent the number of tremors per day, smoothed (gaussian smoothing, $\sigma = 2$ days) in the grey curve. Red horizontal segments represent the known events catalogued by [11]. The (a) panel shows the global performance of SSEdetector over 2007-2022. The red arrow indicates the time window analyzed by Michel et al., while the green arrows describe the two periods from which the synthetic training samples have been derived. The grey rectangle indicates the period which was not covered by the PNSN catalog. In this period, data from Ide, 2012 [30] has been used. The (b), (c) and (d) panels show zooms on 2016-2017, 2019-2021 and 2017 (January to July), respectively.

162 the slab geometry (Figure 1(b)), even in locations where the density of stations
 163 is higher. In this case, the magnitude threshold increases because these events
 164 are more difficult to detect.

165 Continuous SSE detection in Cascadia from raw geodetic 166 data during 2007-2022

167 Overall characteristics of the detected events

168 In order to evaluate how SSEdetector performs on real continuous data, we
 169 applied it to the raw GNSS time series in Cascadia for the period 2007-2022.
 170 SSEdetector scans the data with a 60-day sliding window (1-day stride), pro-
 171 viding a probability of detection for the central day in each window. Figure
 172 3(a) shows the probability of slow slip event detection (in blue) together with
 173 the tremor activity over the period 2007-2022 (in grey). We consider having a

Table 1 Comparison of the number detections from SSEdetector with respect to the catalog from Michel et al. [11]. We distinguish detections in 2007-2017 (the same period analyzed by Michel et al.), and in 2017-2022. We further discriminate, in 2007-2017, between events in common with Michel et al., and new events in the same period.

| | Period | Method | |
|-----------|-------------------------------|---------------|-------------|
| | | Michel et al. | SSEdetector |
| 2007-2017 | Common with Michel et al. | 40 | 35 |
| | Not detected by Michel et al. | 0 | 20 |
| 2017-2022 | | 0 | 23 |

reliable detection when the probability value exceeds 0.5. We find 78 slow slip events over the period 2007-2022, with durations ranging from 2 to 79 days. We find 55 slow slip events in the period 2007-2017, to be compared with the 40 detections of the catalog of Michel et al. [11] (Table 1). We detect 35 of the 40 (87.5%) catalogued SSEs. Three of the missed SSEs have a magnitude smaller than 5.5, one of them has a magnitude of 5.86. The remaining one has a magnitude of 6.03. We show their location in Supplementary Figure 2, superimposed on the magnitude threshold map derived for SSEdetector (see Figure 2(b)). Given their location, the five missed events have magnitudes that are below the magnitude resolution limit (from 6 to 6.5, see Supplementary Figure 2). The remaining 20 events may be associated with new undetected SSEs. We also find 23 new events in the period 2017-2022, which was not covered by Michel et al. [11]. We fixed the detection threshold to its default value of 0.5, *i.e.*, the model detects an event with a 50% confidence. Yet, this threshold can be modified in accordance with specific needs: if high-confidence detections are required, the threshold can be raised; conversely, it can be lowered to capture more events with lower confidence. Interestingly, the few SSE from Michel et al. that were missed with a 0.5 confidence are all detected when selecting a 0.4 threshold.

We also analyze the shape of the static displacement field in correspondence with the detected SSEs (cf. Supplementary Figure 3). We compute the static displacement field by taking the median displacement over three days and subtracting the displacement value at each station corresponding to the dates of the SSE. We find a good accordance with independent studies [11, 31, 32]. Moreover, many of the events found after 2018, as well as the new events detected in the period analyzed by Michel et al., have a displacement field suggesting that they are correct detections.

Analysis of the SSE durations

The shape of the probability curve gives insights into how SSEdetector reveals slow slip events from raw geodetic data. The probability curve in correspondence with an event has a bell shape: it grows until a maximum value, then it smoothly decreases when the model does not see any displacement associated with slow deformation in the data anymore. We use this property of the probability curve to extract a proxy on the detected SSE duration, based on the

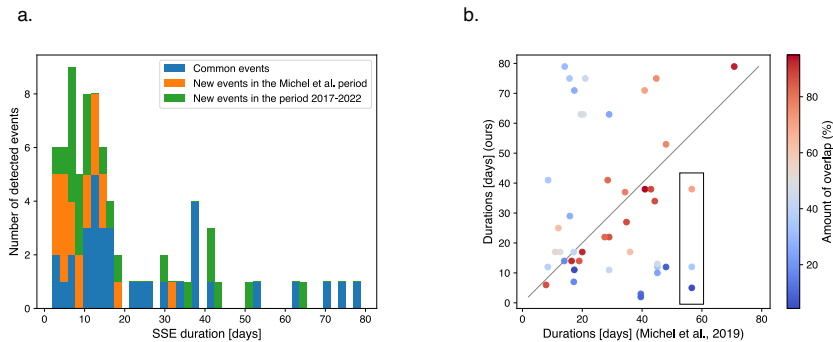


Fig. 4 Distribution of the detected SSEs and comparison with the independent catalogue from Michel et al. [11]. (a) Cumulative histogram of SSEdetector inferred durations. Blue bars represent the 35 catalogued events by Michel et al. [11] that have been successfully retrieved. Orange bars show the 20 additional events that have been discovered within the time window analyzed by Michel et al., while green ones represent the 23 events found in the time period 2017-2022, not covered by the catalog of Michel et al. (b) Event durations from Michel et al.'s catalogue with respect to the durations obtained by SSEdetector. Events are color-coded by the overlap percentage (details in Methods).

208 time span associated with the probability curve exceeding 0.5. We present the
 209 duration distribution in Figure 4. We detect most of the SSEs found by Michel
 210 et al., but we also find many more events, not only in the 2018-2022 period
 211 which was not investigated by Michel et al., but also within the 2007-2017 time
 212 window that they analyzed, suggesting that our method is more sensitive. We
 213 find potential slow slip events at all scales of durations (from 2 to 79 days).
 214 Michel et al. hardly detect SSEs that last less than 15 days, probably due to
 215 temporal data smoothing [11], while we retrieve shorter events (less than 10
 216 days) since we use raw time series, meaning that our method has a better
 217 temporal resolution. In Figure 4 (b), we show a comparison between the SSE
 218 durations of Michel et al.'s [11] catalogued events and ours. This plot is made
 219 by considering all the combinations between events in our catalog and in the
 220 Michel et al. one. Each horizontal alignment represents an event in our cata-
 221 logue that is split into sub-events in the Michel et al. catalog, while vertical
 222 alignments show events in the Michel et al. catalog corresponding to sub-events
 223 in our catalog. We find that the durations are in good accordance for a large
 224 number of events, for which the overlap is often higher than 70%, both for
 225 small- and large-magnitude ones. We can also identify, from figure 4(b), that
 226 some events are separated in one method while identified as one single SSE
 227 in the other: this is the case for the 55 day-long event from Michel et al. [11],
 228 that was paired with 3 SSEdetector sub-events (see Figure 3(d) and the rect-
 229 angle in Figure 4(b)). The majority of the points located off the identity line
 230 (the diagonal) are thus sub-events for which the grouping differs in the two
 231 catalogs. As more points are below the diagonal than above, we can see that
 232 SSEdetector tends to separate the detections more. We interpret this as a pos-
 233 sible increase in the detection precision, yet a validation with an independent
 234 acquisition data set is needed, since the separation into sub-events strongly

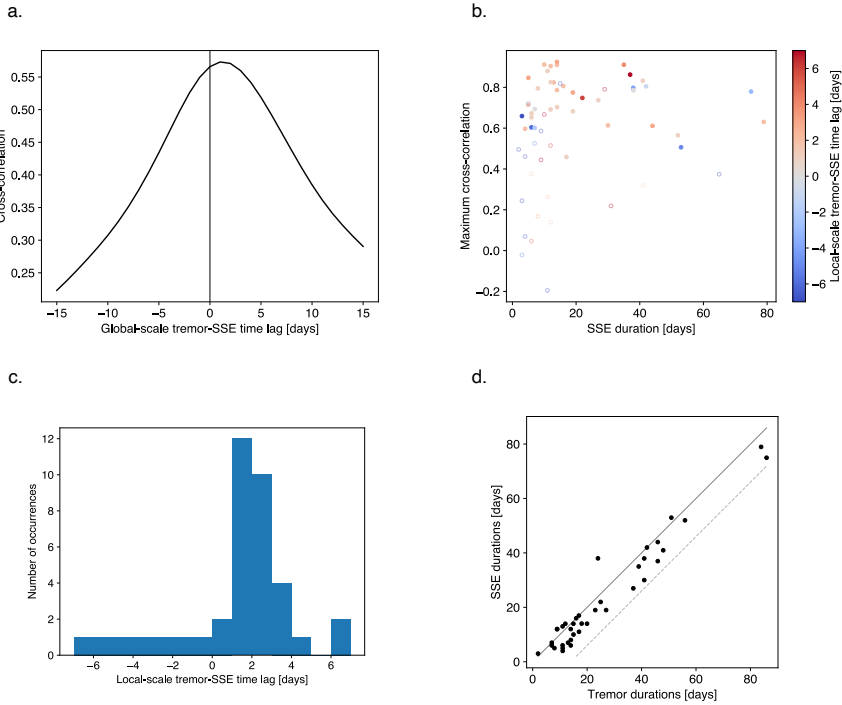


Fig. 5 Validation of SSEdetector performance against tremor activity in 2010-2022. (a) Global-scale cross-correlation between the full-length SSEdetector output probability and the number of tremors per day, as a function of the time shift between the two curves. (b) Local-scale maximum value of cross-correlation for each SSE and tremor windows, centered on the SSE duration, as a function of the SSE duration, color-coded by the associated time lag between tremor and SSE (positive lag means deformation precedes tremor). Events having a zero-point cross-correlation (correlation coefficient) lower than 0.4 are marked with an empty point. (c) Histogram of the time lags computed in the (b) panel. (d) SSE durations as a function of the tremor durations for the events in the (b) panel which have a correlation greater than 0.4. The solid black line represents the identity line, while the dashed grey line is the maximum tremor duration that can be attained for a given SSE duration, that is SSE duration + 14 days (see section section "Computation of local- and global-scale correlations").

235 depends on the threshold applied to the detection probability to define a slow
 236 slip event (0.5 in this study).

237 Validation against tremors

238 In order to have an independent validation, we compare our results with tremor
 239 activity from the Pacific Northwest Seismic Network (PNSN) catalog [24]
 240 between 2009-2022 and Ide's catalogue [30] catalog between 2006-2009, shown
 241 in grey in Figure 3. We show the location of the tremors in our catalogues with
 242 the dashed black contour in Figure 1(b). From a qualitative point of view, we
 243 can see that the detection probability curve seems to align well with the number
 244 of tremors per day, throughout the whole period. This is also true for the

245 20 possible new detected events that were not present in previous catalogs, for
246 example during the period after 2017 (see Figure 3(c)), but also in 2016-2017,
247 where we detect 11 possible events that were not previously catalogued (see
248 Figure 3(b)). The excellent similarity between tremors and our detections is
249 quantitatively assessed by computing the cross-correlation between the prob-
250 ability curve and the number of tremors per day, the latter smoothed with a
251 gaussian filter ($\sigma = 1.5$ days), as a function of the time shift between the two
252 curves (Figure 5(a)). The interval 2007-2010 has been excluded from Figure
253 5(a) in order to consider the period covered by the PNSN catalog only. The
254 maximum correlation value is around 0.58 and is obtained for a time shift
255 between 1 and 2 days. This shows that, at a global scale, the probability peaks
256 are coeval with the peaks of tremor activity.

257 We also make a further comparison at the local scale for each individual
258 detected SSE. In Figure 5(b) we observe that most of the individual detected
259 SSEs show a correlation larger than 0.4 with the coeval peak of tremor. SSE
260 and tremor signals are offset by about 2 days on average (see Figure 5(c)).
261 This result, obtained on windows of month-long scale, seems consistent with
262 the decade-long correlation shown in Figure 5(a), suggesting that the found
263 large-scale trend is also true at a smaller scale. This may suggest that the
264 slow deformation, for which the detection probability is a proxy, precedes the
265 tremor chatter by a few days, with potential implications on the nucleation of
266 the slow rupture.

267 We compare the tremor peak duration (see details in Methods) to the SSE
268 duration in Figure 5(d) for all the events that have been also considered in
269 Figure 5(b). The figure shows a correspondence between slow slip duration
270 and coeval tremor activity duration: most of the events are associated with a
271 peak of tremor activity of close duration. This is true also for large events, up
272 to 80 days. This finding gives an insight that our deep learning-based method,
273 blindly applied to raw geodetic time series, achieves reliable results. Yet, this
274 result should be taken with caution, since it is strongly dependent on the choice
275 of the window of observation (see Methods section for further details).

276 Sensitivity study

277 We analyze the sensitivity of SSEdetector with respect to the number of
278 stations. We construct an alternative test selecting 352 GNSS stations (see
279 Supplementary Figure 5), which is the number of stations used by Michel et
280 al. [11]. The 217 extra stations have larger percentages of missing data com-
281 pared to the initial 135 stations (cf. Supplementary Figure 4). We train and
282 test SSEdetector with 352 time series and we report the results in Supplemen-
283 tary Figures 6-7. We observe that the results are similar, with an excellent
284 alignment with tremors and similar correlation and lag values, although with
285 this setting the detection power slightly decreases, probably due to a larger
286 number of missing data.

287 We also test the ability of SSEdetector to identify SSEs in a sub-region
288 only (even if it is trained with a large-scale network. For that, we test SSEde-
289 tector (trained on 135 stations), without re-training, on a subset of the GNSS
290 network, situated in the northern part of Cascadia. To this end, we replace
291 with zeros all the data associated with stations located at latitudes lower than
292 47 degrees (see Supplementary Figure 8). Similarly, we find that SSEdetector
293 retrieves all the events which were found by Michel et al. [11] and the correla-
294 tion with tremors that occur in this sub-region is still high, with a global-scale
295 cross-correlation of 0.5 (cf. Supplementary Figures 9-10). This means that the
296 model is robust against long periods of missing data and, thanks to the spatial
297 pooling strategy, can generalize over different settings of stations and obtain
298 some information on the localization.

299 Finally, we test SSEdetector against other possible deep learning models
300 that could be used for detection. We report in Supplementary Figures 11-12
301 the results obtained by replacing the one-dimensional convolutional layers with
302 two-dimensional convolutions on time series sorted by latitude (as shown in
303 Figure 1(a)). This type of architecture was used in studies having similar multi-
304 station time-series data [33]. We observe that the results on real data are not
305 satisfactory because of too high a rate of false detections and a lower temporal
306 resolution than SSEdetector (in other words, short SSEs are not retrieved).
307 This suggests that our specific model architecture, handling in different ways
308 the time dimension and the station dimension, might be more suited to multi-
309 station time-series data sets.

310 Discussion

311 In this study, we use a multi-station approach that proves efficient in detecting
312 slow slip events in raw GNSS time series even in presence of SSE migrations
313 [11, 31, 32]. Thanks to SSEdetector, we are able to detect 87 slow slip events
314 with durations from 2 to 79 days, with an average limit magnitude of about
315 6.4 in north Cascadia and 6.2 in south Cascadia computed on the synthetic
316 test set (see Figure 2(b)). The magnitude of the smallest detected SSE in
317 common with Michel et al. is 5.42, with a corresponding duration of 8.5 days.
318 One current limitation of this approach is that the location information is not
319 directly inferred. In this direction, some efforts should be made in developing a
320 method for characterizing slow slip events after the detection in order to have
321 information on the location, but also on the magnitude, of the slow rupture.

322 We apply our methodology to the Cascadia subduction zone because it is
323 the area where independent benchmarks exist and it is thus possible to validate
324 a new method. However, the applicability of SSEgenerator and SSEdetector to
325 other subduction zones is possible. The current approach is, however, region-
326 specific. In fact, the characteristics of the targeted zone affect the structure
327 of the synthetic data, thus a method trained on a specific region could have
328 poor performance if tested on another one without retraining. This problem
329 can be addressed by generating multiple data sets associated with different

330 regions and combining them for the training. Also, we focus on the Casca-
331 dia subduction zone, where not much regular seismicity occurs, making it a
332 prototypical test zone when looking for slow earthquakes. When addressing
333 other regions, such as Japan, for example, the influence of earthquakes or
334 post-seismic relaxation signals could make the problem more complex. This
335 extension goes beyond the scope of this study, yet we think that it will be
336 essential to tackle this issue in order to use deep learning approaches for the
337 detection of SSEs in any region.

338 Conclusions

339 We developed a powerful pipeline, composed of a realistic synthetic GNSS
340 time-series generation, SSEgenerator, and a deep-learning classification model,
341 SSEdetector, aimed to detect slow slip events from a series of raw GNSS time
342 series measured by a station network. We built a new catalog of slow slip events
343 in the Cascadia subduction zone by means of SSEdetector. We found 78 slow
344 slip events from 2007 to 2022, 35 of which are in good accordance with the
345 existing catalog [11]. The detected SSEs have durations that range between a
346 few days to a few months. The detection probability curve correlates well with
347 the occurrence of tremor episodes, even in time periods where we found new
348 events. The duration of our SSEs, for the 35 known events, as well as for the
349 43 new detections, are found to be similar to the coeval tremor duration. The
350 comparison between tremors and SSEs also shows that, both at a local and a
351 global temporal scale, the slow deformation may precede the tremor chatter
352 by a few days, with potential implications on the link between a slow slip that
353 could drive the rupture of nearby small seismic asperities. This is the first
354 successful attempt to detect SSEs from raw GNSS time series, and we hope
355 that this preliminary study will lead to the detection of SSEs in other active
356 regions of the world.

357 Methods

358 SSEgenerator: data selection

359 We consider the 550 stations in the Cascadia subduction zone, belonging to
360 the MAGNET GNSS network, and we select data from 2007 to 2022. We train
361 SSEdetector with synthetic data whose source was affected by different noise
362 and data gap patterns. We divide the data into three periods: 2007-2014,
363 2014-2018, and 2018-2022. In order to create a more diverse training set, data in
364 the period 2007-2014 and 2018-2022 has been chosen as a source for synthetic
365 data generation. The period 2014-2018 was left aside and used as an inde-
366 pendent validation set for performance assessment on real data. Nonetheless,
367 since synthetic data is performed by applying random transformations, a test
368 on the whole sequence 2007-2022 is possible without overfitting.

369 For the two periods 2007-2014 and 2018-2022, we sort the GNSS stations
370 by the total number of missing data points and we choose the 135 stations

371 affected by fewer data gaps as the final subset for our study. We make sure
 372 that stations having too high a noise do not appear in this subset. We select
 373 135 stations since it represents a good compromise between the presence of
 374 data and the longest data gap sequence in a 60-day window. However, we also
 375 train and test SSEdetector on 352 stations (the same number used in the study
 376 by Michel et al. [11]). We briefly discuss the results in the section "Sensitivity
 377 study".

378 SSEgenerator: Generation of ultra-realistic noise time 379 series

380 Raw GNSS data is first detrended at each of the 135 stations, *i.e.* the linear
 381 trend is removed, where the slope and the intercept are computed, for each
 382 station, without taking into account the data gaps, *i.e.*, for each station the
 383 mean over time is calculated without considering the missing data points,
 384 and is removed from the series. A matrix containing all station time series
 385 $\mathbf{X} \in \mathbb{R}^{N_t \times N_s}$ is built, where N_t is the temporal length of the input time series
 386 and N_s is the number of stations. In this study, we use 2 components (N-S and
 387 E-W) and we apply the following procedure for each component independently.
 388 Each column of \mathbf{X} contains a detrended time series. We proceed as follows.
 389 The \mathbf{X} matrix is then re-projected in another vector space through a Principal
 390 Component Analysis (PCA), as follows. First, the data is centered. The mean
 391 vector is computed $\boldsymbol{\mu} \in \mathbb{R}^{N_s}$, such that μ_i is the mean of the i -th time series.
 392 The centered matrix is considered $\tilde{\mathbf{X}} = \mathbf{X} - \boldsymbol{\mu}$, and is decomposed through
 393 Singular Value Decomposition (SVD) to obtain the matrix of right singular
 394 vectors \mathbf{V} , which is the rotation matrix containing the spatial variability of
 395 the original vector space. We further rotate the data by means of this spatial
 396 matrix to obtain spatially-uncorrelated time series $\hat{\mathbf{X}} = \tilde{\mathbf{X}}\mathbf{V}$. Then, we produce
 397 $\hat{\mathbf{X}}_R$, a randomized version of $\hat{\mathbf{X}}$, by applying the iteratively-refined amplitude-
 398 adjusted Fourier transform (AAFT) method [34], having globally the same
 399 power spectrum and amplitude distribution of the input data. The number
 400 of AAFT iterations has been experimentally set to 5. The surrogate time
 401 series are then back-projected in the original vector space to obtain $\mathbf{X}_R =$
 402 $\hat{\mathbf{X}}_R\mathbf{V}^T + \boldsymbol{\mu}$. We further enrich the randomized time series \mathbf{X}_R by imprinting
 403 the real pattern of missing data for 70 % of the synthetic data. We shuffle
 404 the data gaps before imprinting them to the data, such that SSEdetector can
 405 better generalize over unseen test data for the same station, which necessarily
 406 would have a different pattern of data gaps. We leave the remaining 30% of
 407 the data as it is. We prefer not to use any interpolation method in order not
 408 to introduce new values in the data. Thus, we set all the missing data points
 409 to zero, which is a neutral value with respect to the trend of the data and the
 410 convolution operations performed by SSEdetector.

411 After this process, we generate sub-windows of noise time series as follows.
 412 Given the window length W_L , a uniformly distributed random variable is gener-
 413 ated $s \sim \mathcal{U}(-W_L/2, W_L/2)$ and the data is circularly shifted by the amount
 414 s . Then, $\lfloor N_t/W_L \rfloor$ contiguous (non-overlapping) windows are obtained. The

415 circular shift is needed in order for SSEdetector not to learn a fixed temporal
 416 pattern of data gaps. Finally, by knowing the desired number N of noise win-
 417 dows to compute, the surrogate generation $\{\mathbf{X}_R\}_i$ can be repeated $\lceil \frac{N}{\lfloor N_i/W_L \rfloor} \rceil$
 418 times. In our study, we generate $N = 60,000$ synthetic samples, by calling the
 419 surrogate data generation 1,429 times and extracting 42 non-overlapping noise
 420 windows from each randomized time series.

421 SSEgenerator: Modeling of synthetic slow slip events

422 We first generate synthetic displacements at all the 135 selected stations using
 423 Okada's dislocation model [27]. We draw random locations, strike and dip
 424 angles using the slab2 model [25] following the subduction geometry within
 425 the area of interest (see Figure 1(b)). We let the rake angle be a uniform
 426 random variable from 75 to 100 degrees, in order to have a variability around
 427 90 degrees (thrust focal mechanism). For each (latitude, longitude) couple, we
 428 extract the corresponding depth from the slab and we add further variability,
 429 modeled as a uniformly distributed random variable from -10 to 10 km. We
 430 allow for this variability if the depth is at least 15 km, in order not to have
 431 ruptures that reach the surface. We associate each rupture with a magnitude
 432 M_w , uniformly generated in the range (6, 7), and we compute the equivalent
 433 moment as $M_0 = 10^{1.5M_w+9.1}$. As for the fault geometry, we rely on the circular
 434 crack approximation [35] to compute the fault radius as:

$$R = \left(\frac{7}{16} \frac{M_0}{\Delta\sigma} \right)^{1/3} \quad (1)$$

435 where $\Delta\sigma$ is the static stress drop. We compute the average slip on the
 436 fault as:

$$\bar{u} = \frac{16}{7\pi} \frac{\Delta\sigma}{\mu} \quad (2)$$

437 where μ is the shear modulus. We assume $\mu = 30$ GPa. By imposing that
 438 the surface of the crack must equal a rectangular dislocation of length L and
 439 width W , we obtain $L = \sqrt{2\pi}R$. We assume that $W = L/2$. Finally, we model
 440 the stress drop as a lognormally-distributed random variable. We assume the
 441 average stress drop to be $\overline{\Delta\sigma} = 0.05$ MPa for the Cascadia subduction zone
 442 [26]. We also assume that the coefficient of variation c_V , namely the ratio
 443 between the standard deviation and the mean, is equal to 10. Hence, we gener-
 444 ate the static stress drop as $\Delta\sigma \sim \text{Lognormal}(\mu_N, \sigma_N^2)$, where μ_N and σ_N are
 445 the mean and the standard deviation of the underlying normal distribution,
 446 respectively, that we derive as:

$$\sigma_N = \sqrt{\ln(c_V^2 + 1)} \quad (3)$$

447 and

$$\mu_N = \ln(\overline{\Delta\sigma}) - \sigma_N^2/2. \quad (4)$$

448 We thus obtain the (horizontal) synthetic displacement vector $\mathbf{D}_s =$
 449 (D_s^{N-S}, D_s^{E-W}) at each station s . We model the temporal evolution of slow
 450 slip events as a logistic function. Let D be the E-W displacement for simplicity.
 451 In this case, we model an SSE signal at a station s as:

$$d_s(t) = \frac{D}{1 + e^{\beta(t-t_0)}} \quad (5)$$

452 where β is a parameter associated with the growth rate of the curve and
 453 t_0 is the time corresponding to the inflection point of the logistic function. We
 454 assume $t_0 = 30$ days, so that the signal is centered in the 60-day window. We
 455 derive the parameter β as a function of the slow slip event duration T . We can
 456 rewrite the duration as $T = t_{max} - t_{min}$, where t_{max} is the time corresponding
 457 to the steady-state value of the signal (*i.e.*, D), while t_{min} is associated to
 458 the minimum (*i.e.*, 0). Since these values are only asymptotically reached, we
 459 introduce a threshold γ , such that t_{max} and t_{min} are associated with $d_s(D - \gamma)$
 460 and $d_s(\gamma)$, respectively. We choose $\gamma = 0.01 \cdot D$. By rewriting the duration as
 461 $T = t_{max} - t_{min}$ and solving for β , we obtain:

$$\beta = \frac{2}{T} \ln \left(\frac{D}{\gamma} - 1 \right). \quad (6)$$

462 Finally, we generate slow slip events having uniform duration T between
 463 10 and 30 days. We take half of the noise samples (30,000) and we create a
 464 positive sample (*i.e.*, time series containing a slow slip event) as $\mathbf{X}_R + \mathbf{d}(t)$,
 465 where $\mathbf{d}(t)$ is a matrix containing all the modeled time series $d_s(t)$ for each
 466 station. We let \mathbf{X}_R contain missing data. Therefore, we do not add the signal
 467 $d_s(t)$ where data should not be present.

468 **SSEdetector: Detailed architecture**

469 SSEdetector is a deep neural network obtained by the combination of a convo-
 470 lutional and a Transformer neural network. The full architecture is shown in
 471 Supplementary Figure 13. The model takes input GNSS time series, which can
 472 be grouped as a matrix of shape (N_s, N_t, N_c) , where N_s, N_t, N_c are the number
 473 of stations, window length and number of components, respectively. In this
 474 study, $N_s = 135, N_t = 60$ days and $N_c = 2$ (N-S, E-W). The basic unit of this
 475 CNN is a Convolutional Block. It is made of a sequence of a one-dimensional
 476 convolutional layer in the temporal (N_t) dimension, which computes N_f feature
 477 maps by employing a 1×5 kernel, followed by a Batch Normalization [36] and
 478 a ReLu activation function [37]. We will refer to this unit as ConvBlock(N_f)
 479 for the rest of the paragraph (see Supplementary Figure 13). We alternate convo-
 480 lutional operations in the temporal dimension with pooling operations in
 481 the station dimension (max-pooling with a kernel of 3) and we replicate this
 482 structure as long as the spatial (station) dimension is reduced to 1. To this
 483 end, we create a sequence of 3 ConvBlock(\cdot) + max-pooling. As an example,
 484 the number of stations after the first pooling layer is reduced from 135 to 45.
 485 At each ConvBlock(\cdot), we multiply by 4 the number of computed feature maps

486 N_f . At the end of the CNN, the computed features have shape (N_t, N_f^{final}) ,
 487 with $N_f^{final} = 256$.

488 This feature matrix is given as input to a Transformer neural network. We
 489 first use a Positional Embedding to encode the temporal sequence. We do not
 490 impose any kind of pre-computed embedding, but we use a learnable matrix
 491 of shape (N_t, N_f^{final}) . The learnt embeddings are added to the feature matrix
 492 (*i.e.*, the output of the CNN). The embedded inputs are then fed to a Trans-
 493 former neural network [29], whose architecture is detailed in Supplementary
 494 Figure 14. Here, the global (additive) self-attention of the embedded CNN
 495 features is computed as:

$$\eta_{t_1, t_2} = W_a \tanh(W_1^T h_{t_1} + W_2^T h_{t_2} + b_h) + b_a, \quad (7)$$

$$a_{t_1, t_2} = \text{softmax}(\eta_{t_1, t_2}) = \frac{e^{\eta_{t_1, t_2}}}{\sum_{t_2} e^{\eta_{t_1, t_2}}}, \quad (8)$$

$$c_{t_1} = \sum_{t_2} a_{t_1, t_2} \cdot h_{t_2}, \quad (9)$$

496 where W represents a learnable weight matrix and b a bias vector. The
 497 matrices h_{t_1} and h_{t_2} are the hidden-state representations at time t_1 and
 498 t_2 , respectively. The matrix a_{t_1, t_2} contains the attention scores for the time
 499 steps t_1 and t_2 . Here, a context vector is computed as the weighted sum of
 500 the hidden-state representations by the attention scores. The context vector
 501 contains the importance at a given time step based on all the features in
 502 the window. The contextual information is then added to the Transformer
 503 inputs. Then, a position-wise Feed-Forward layer (with a dropout rate of 0.1)
 504 is employed to add further non-linearity. After the Transformer network, a
 505 Global Average Pooling in the temporal dimension (N_t) is employed to gather
 506 the transformed features and to output a vector summarizing the temporal
 507 information. A Dropout is then added as a form of regularization to reduce
 508 overfitting [38], with dropout rate $\delta = 0.2$. In the end, we use a fully-connected
 509 layer with one output, with a sigmoid activation function to express the
 510 probability of SSE detection.

511 Training details

512 We perform a mini-batch training [39] (batch size of 128 samples) by minimiz-
 513 ing the binary cross-entropy (BCE) loss between the target labels y and the
 514 predictions \hat{y} (a probability estimate):

$$\text{BCE}(y, \hat{y}) = -y \ln(\hat{y}) - (1 - y) \ln(1 - \hat{y}). \quad (10)$$

515 The BCE loss is commonly used for binary classification problems (detection is a binary classification). We use the ADAM method for the optimization [40] with a learning rate $\lambda = 10^{-3}$ which has been experimentally chosen. We schedule the learning rate such that it is reduced during training iterations

519 and we stop the training when the validation loss did not improve for 50 con-
 520 secutive epochs. We initialized the weights of SSEdetector with a uniform He
 521 initializer [41]. We implemented the code of SSEdetector in Python using the
 522 Tensorflow and Keras libraries [42, 43]. We run the training on NVIDIA Tesla
 523 A100 Graphics Processing Units (GPUs). The training of SSEdetector takes
 524 less than 2 hours. The inference on the whole 15-year sequence (2007-2022)
 525 takes a few minutes.

526 **Calculation of tremor durations**

527 We compute the durations of tremor bursts using the notion of topographic
 528 prominence, explained in the following. We rely on the software implemen-
 529 tation from the SciPy Python library [44]. Given a peak in the curve, the
 530 topographic prominence is informally defined as the minimum elevation that
 531 needs to be descended to start reaching a higher peak. The procedure is graph-
 532 ically detailed in Supplementary Figure 15. We first search for peaks in the
 533 number of tremors per day by comparison with neighboring values. In order to
 534 avoid too many spurious local maxima, we smooth the number of tremors per
 535 day with a gaussian filter ($\sigma = 1.5$ days). For each detected SSE, we search for
 536 peaks of tremors in a window given by the SSE duration ± 3 days. For each
 537 peak of tremors that is found, the corresponding width is computed as follows.
 538 The topographic prominence is computed by placing a horizontal line at the
 539 peak height h (the value of the tremor curve corresponding to the peak). An
 540 interval is defined, corresponding to the points where the line crosses either the
 541 signal bounds or the signal at the slope towards a higher peak. In this interval,
 542 the minimum values of the signal on each side are computed, representing the
 543 bases of the peak. The topographic prominence p of the peak is then defined as
 544 the height between the peak and its highest base value. Then, the local height
 545 of the peak is computed as $h_L = h - \alpha \cdot p$. We set $\alpha = 0.7$ in order to focus on
 546 the main tremor pulses, discarding further noise in the curve. From the local
 547 height, another horizontal line is considered and the peak width is computed
 548 as the intersection point of the line with either a slope, the vertical position
 549 of the bases or the signal bounds, on both sides. Finally, the total width of a
 550 tremor pulse in an SSE window is computed by considering the earliest start-
 551 ing point on the left side and the latest ending point on the right side. It must
 552 be noticed that, the derivation of the tremor duration depends on the window
 553 length. In fact, the inferred tremor duration can saturate to a maximum value
 554 equal to the length of the window. For this reason, we added in Figure 5(c)
 555 a dashed line corresponding to the window length (SSE duration + 14 days)
 556 (see section section "Computation of local- and global-scale correlations").

557 **Computation of local- and global-scale correlations**

558 We compute the time-lagged cross-correlation between the SSE probability
 559 and the number of tremors per day (Fig. 5(a) and (b)). We smooth the number

of tremors per day with a gaussian filter ($\sigma = 1.5$ days). We consider a lag between -7 and 7 days, with a 1-day stride.

In the case of Fig. 5(a), we compute the global correlation coefficient by considering the whole time sequence (2010-2022). As for Fig. 5(b), we make a local analysis. For each detected SSE, we first extract SSE and tremor slices from intervals centered on the SSE dates $[t_{SSE}^{start} - \Delta t, t_{SSE}^{end} + \Delta t]$, where $\Delta t = 30$ days. We first compute the cross-correlation between the two curves to filter out detected SSEs whose similarity with tremors is not statistically significant. For each SSE date, we discard $p''(t)$ and $n_T''(t)$ if their correlation coefficient is lower than 0.4. We build Fig. 5(b) after this process.

We compute Fig. 5(d) by comparing the SSE and tremor durations for all the events that had a cross-correlation higher than 0.4. For those, we infer the tremor duration, using the method explained in section "Calculation of tremor durations" on the $n_T(t)$ cut from an interval $[t_{SSE}^{start} - \Delta t', t_{SSE}^{end} + \Delta t']$, with $\Delta t' = 7$ days.

Overlap percentage calculation

In Figure 4 (b) we color-code the SSE durations by the overlap percentage between a pair of events, which we compute as the difference between the earliest end and the latest start, divided by the sum of the event lengths. Let E_1 and E_2 be two events with start and end dates given by (t_1^{start}, t_1^{end}) and (t_2^{start}, t_2^{end}) and with durations given by $D_1 = t_1^{end} - t_1^{start}$ and $D_2 = t_2^{end} - t_2^{start}$, respectively. We compute their overlap π as:

$$\pi = \frac{\max(0, \min(t_1^{end}, t_2^{end}) - \max(t_1^{start}, t_2^{start}))}{D_1 + D_2} \quad (11)$$

Supplementary information. This article has an accompanying supplementary file.

Acknowledgments. This work has been funded by ERC CoG 865963 DEEP-trigger. Most of the computations presented in this paper were performed using the GRICAD infrastructure (<https://gricad.univ-grenoble-alpes.fr>), which is supported by Grenoble research communities.

Declarations

- **Funding** This work has been funded by ERC CoG 865963 DEEP-trigger.
- **Competing interests** The authors declare no competing interests.
- **Materials & Correspondence** Correspondence to: Giuseppe Costantino
- **Data availability** We downloaded the data from the Nevada Geodetic Laboratory (<http://geodesy.unr.edu>).
- **Code availability** The source code of SSEgenerator and SSEdetector as well as the pre-trained model of SSEdetector are available at <https://gricad-gitlab.univ-grenoble-alpes.fr/costangi/sse-detection>.

597 • **Authors' contributions** G.C. developed SSEgenerator and SSEdetector
598 and produced all the results and figures presented here. A.S. designed the
599 study and provided expertise for the geodetic data analysis. S.G.R. provided
600 expertise for the Deep Learning aspects. G.C. wrote the first draft of the
601 paper. All the authors contributed to reviewing the manuscript.

602 References

- 603 [1] Dragert, H., Wang, K., James, T.S.: A silent slip event on
604 the deeper cascadia subduction interface. *Science* **292**, 1525–
605 1528 (2001). [https://doi.org/10.1126/SCIENCE.1060152/ASSET/
606 DD167B7B-24D4-40A4-996F-4603D42C0244/ASSETS/GRAPHIC/
607 SE1919443004.JPEG](https://doi.org/10.1126/SCIENCE.1060152/ASSET/DD167B7B-24D4-40A4-996F-4603D42C0244/ASSETS/GRAPHIC/SE1919443004.JPEG)
- 608 [2] Lowry, A.R., Larson, K.M., Kostoglodov, V., Bilham, R.: Transient fault
609 slip in guerrero, southern mexico. *Geophysical Research Letters* **28**, 3753–
610 3756 (2001). <https://doi.org/10.1029/2001GL013238>
- 611 [3] Schwartz, S.Y., Rokosky, J.M.: Slow slip events and seismic tremor at
612 circum-pacific subduction zones. *Reviews of Geophysics* **45** (2007). [https:
613 //doi.org/10.1029/2006RG000208](https://doi.org/10.1029/2006RG000208)
- 614 [4] Ide, S., Beroza, G.C., Shelly, D.R., Uchide, T.: A scaling law for slow
615 earthquakes. *Nature* 2007 447:7140 **447**, 76–79 (2007). [https://doi.org/
616 10.1038/nature05780](https://doi.org/10.1038/nature05780)
- 617 [5] Mousavi, S.M., Ellsworth, W.L., Zhu, W., Chuang, L.Y., Beroza, G.C.:
618 Earthquake transformer—an attentive deep-learning model for simulta-
619 neous earthquake detection and phase picking. *Nature Communications*
620 **11** (2020). <https://doi.org/10.1038/s41467-020-17591-w>
- 621 [6] Rousset, B., Campillo, M., Lasserre, C., Frank, W.B., Cotte, N., Walpers-
622 dorf, A., Socquet, A., Kostoglodov, V.: A geodetic matched filter search
623 for slow slip with application to the mexico subduction zone. *Journal
624 of Geophysical Research: Solid Earth* **122**, 10498–10514 (2017). [https:
625 //doi.org/10.1002/2017JB014448](https://doi.org/10.1002/2017JB014448)
- 626 [7] Frank, W.B., Radiguet, M., Rousset, B., Shapiro, N.M., Husker, A.L.,
627 Kostoglodov, V., Cotte, N., Campillo, M.: Uncovering the geodetic signa-
628 ture of silent slip through repeating earthquakes. *Geophysical Research
629 Letters* **42**, 2774–2779 (2015). <https://doi.org/10.1002/2015GL063685>
- 630 [8] Gombert, J., Wech, A., Creager, K., Obara, K., Agnew, D.: Reconsidering
631 earthquake scaling. *Geophysical Research Letters* **43**, 6243–6251 (2016).
632 <https://doi.org/10.1002/2016GL069967>

- 633 [9] Hawthorne, J.C., Bartlow, N.M.: Observing and modeling the spectrum
634 of a slow slip event. *Journal of Geophysical Research: Solid Earth* **123**,
635 4243–4265 (2018). <https://doi.org/10.1029/2017JB015124>
- 636 [10] Frank, W.B., Brodsky, E.E.: Daily measurement of slow slip from low-
637 frequency earthquakes is consistent with ordinary earthquake scaling.
638 *Science Advances* **5** (2019). [https://doi.org/10.1126/SCIADV.AAW9386/
639 SUPPL_FILE/AAW9386_SM.PDF](https://doi.org/10.1126/SCIADV.AAW9386/SUPPL_FILE/AAW9386_SM.PDF)
- 640 [11] Michel, S., Gualandi, A., Avouac, J.P.: Similar scaling laws for earth-
641 quakes and cascadia slow-slip events. *Nature* **574**, 522–526 (2019). <https://doi.org/10.1038/s41586-019-1673-6>
642
- 643 [12] Bartlow, N.M., Miyazaki, S., Bradley, A.M., Segall, P.: Space-time correla-
644 tion of slip and tremor during the 2009 cascadia slow slip event. *Geophys-
645 ical Research Letters* **38** (2011). <https://doi.org/10.1029/2011GL048714>
- 646 [13] Radiguet, M., Cotton, F., Vergnolle, M., Campillo, M., Walpersdorf, A.,
647 Cotte, N., Kostoglodov, V.: Slow slip events and strain accumulation in
648 the guerrero gap, mexico. *Journal of Geophysical Research: Solid Earth*
649 **117**, 4305 (2012). <https://doi.org/10.1029/2011JB008801>
- 650 [14] Kong, Q., Trugman, D.T., Ross, Z.E., Bianco, M.J., Meade, B.J., Ger-
651 stoft, P.: Machine learning in seismology: Turning data into insights.
652 *Seismological Research Letters* **90**, 3–14 (2019). [https://doi.org/10.1785/
653 0220180259](https://doi.org/10.1785/0220180259)
- 654 [15] Zhu, W., Beroza, G.C.: Phasenet: A deep-neural-network-based seismic
655 arrival-time picking method. *Geophysical Journal International* **216**, 261–
656 273 (2019). <https://doi.org/10.1093/gji/ggy423>
- 657 [16] Woollam, J., Münchmeyer, J., Tilmann, F., Rietbrock, A., Lange, D.,
658 Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Miche-
659 lini, A., Saul, J., Soto, H.: Seisbench-a toolbox for machine learning in
660 seismology. *Seismological Research Letters* **93**, 1695–1709 (2022). [https://doi.org/10.1785/
661 0220210324](https://doi.org/10.1785/0220210324)
- 662 [17] Ross, Z.E., Trugman, D.T., Hauksson, E., Shearer, P.M.: Search-
663 ing for hidden earthquakes in southern california. *Science*
664 (2019). [https://doi.org/10.1126/SCIENCE.AAW6888/
665 AAW6888_ROSS_SM.PDF](https://doi.org/10.1126/SCIENCE.AAW6888/SUPPL_FILE/AAW6888_ROSS_SM.PDF)
- 666 [18] Tan, Y.J., Waldhauser, F., Ellsworth, W.L., Zhang, M., Zhu, W., Michele,
667 M., Chiaraluce, L., Beroza, G.C., Segou, M.: Machine-learning-based high-
668 resolution earthquake catalog reveals how complex fault structures were
669 activated during the 2016–2017 central italy sequence. *The Seismic Record*
670 **1**, 11–19 (2021). <https://doi.org/10.1785/0320210001>

- 671 [19] Ross, Z.E., Cochran, E.S., Trugman, D.T., Smith, J.D.: 3d fault
672 architecture controls the dynamism of earthquake swarms. *Science*
673 **368**, 1357–1361 (2020). [https://doi.org/10.1126/SCIENCE.ABB0779/
674 SUPPL_FILE/ABB0779_ROSS_SM.PDF](https://doi.org/10.1126/SCIENCE.ABB0779/SUPPL_FILE/ABB0779_ROSS_SM.PDF)
- 675 [20] Tan, Y.J., Marsan, D.: Connecting a broad spectrum of transient slip on
676 the san andreas fault. *Science Advances* **6**, 2489–2503 (2020). [https://doi.
677 org/10.1126/SCIADV.ABB2489/SUPPL_FILE/ABB2489_SM.PDF](https://doi.org/10.1126/SCIADV.ABB2489/SUPPL_FILE/ABB2489_SM.PDF)
- 678 [21] Rouet-Leduc, B., Jolivet, R., Dalaison, M., Johnson, P.A., Hulbert, C.:
679 Autonomous extraction of millimeter-scale deformation in insar time
680 series using deep learning. *Nature Communications* 2021 12:1 **12**, 1–11
681 (2021). <https://doi.org/10.1038/s41467-021-26254-3>
- 682 [22] Costantino, G., Giffard-Roisin, S., Marsan, D., Marill, L., Radiguet, M.,
683 Dalla Mura, M., Janex, G., Socquet, A.: Seismic source characterization
684 from gns data using deep learning. *Authorea Preprints* (2022)
- 685 [23] Rogers, G., Dragert, H.: Episodic tremor and slip on the cascadia
686 subduction zone: The chatter of silent slip. *Science* **300**, 1942–
687 1943 (2003). [https://doi.org/10.1126/SCIENCE.1084783/ASSET/
688 DCC117CE-4CAD-4799-90C0-B9DFE257EA05/ASSETS/GRAPHIC/
689 SE2431617002.JPEG](https://doi.org/10.1126/SCIENCE.1084783/ASSET/DCC117CE-4CAD-4799-90C0-B9DFE257EA05/ASSETS/GRAPHIC/SE2431617002.JPEG)
- 690 [24] Wech, A.G.: Interactive tremor monitoring. *Seismological Research Let-
691 ters* **81**, 664–669 (2010). <https://doi.org/10.1785/GSSRL.81.4.664>
- 692 [25] Hayes, G.P., Moore, G.L., Portner, D.E., Hearne, M., Flamme, H.,
693 Furtney, M., Smoczyk, G.M.: Slab2, a comprehensive subduction zone
694 geometry model. *Science* **362**(6410), 58–61 (2018)
- 695 [26] Gao, H., Schmidt, D.A., Weldon, R.J.: Scaling relationships of source
696 parameters for slow slip events. *Bulletin of the Seismological Society of
697 America* **102**(1), 352–360 (2012)
- 698 [27] Okada, Y.: Surface deformation due to shear and tensile faults in a half-
699 space. *Bulletin of the seismological society of America* **75**(4), 1135–1154
700 (1985)
- 701 [28] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436–
702 444 (2015)
- 703 [29] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez,
704 A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in
705 neural information processing systems* **30** (2017)
- 706 [30] Ide, S.: Variety and spatial heterogeneity of tectonic tremor worldwide.

- 707 Journal of Geophysical Research: Solid Earth **117**(B3) (2012)
- 708 [31] Itoh, Y., Aoki, Y., Fukuda, J.: Imaging evolution of cascadia slow-slip
709 event using high-rate gps. *Scientific Reports* 2022 12:1 **12**, 1–12 (2022).
710 <https://doi.org/10.1038/s41598-022-10957-8>
- 711 [32] Bletery, Q., Nocquet, J.-M.: Slip bursts during coalescence of slow slip
712 events in cascadia. *Nature communications* **11**(1), 2159 (2020)
- 713 [33] Licciardi, A., Bletery, Q., Rouet-Leduc, B., Ampuero, J.-P., Juhel, K.:
714 Instantaneous tracking of earthquake growth with elastogravity signals.
715 *Nature* **606**(7913), 319–324 (2022)
- 716 [34] Schreiber, T., Schmitz, A.: Surrogate time series. *Physica D: Nonlinear
717 Phenomena* **142**(3-4), 346–382 (2000)
- 718 [35] Lay, T., Wallace, T.C.: *Modern Global Seismology*. Elsevier, ??? (1995)
- 719 [36] Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network
720 training by reducing internal covariate shift. In: *International Conference
721 on Machine Learning*, pp. 448–456 (2015). pmlr
- 722 [37] Agarap, A.F.: Deep learning using rectified linear units (relu). arXiv
723 preprint arXiv:1803.08375 (2018)
- 724 [38] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov,
725 R.: Dropout: a simple way to prevent neural networks from overfitting.
726 *The journal of machine learning research* **15**(1), 1929–1958 (2014)
- 727 [39] Bottou, L., Curtis, F.E., Nocedal, J.: Optimization methods for large-scale
728 machine learning. *Siam Review* **60**(2), 223–311 (2018)
- 729 [40] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv
730 preprint arXiv:1412.6980 (2014)
- 731 [41] He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing
732 human-level performance on imagenet classification. In: *Proceedings of
733 the IEEE International Conference on Computer Vision*, pp. 1026–1034
734 (2015)
- 735 [42] Chollet, F., et al.: Keras. <https://github.com/fchollet/keras>
- 736 [43] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Cor-
737 rado, G.S., Davis, A., Dean, J., Devin, M., et al.: Tensorflow: Large-scale
738 machine learning on heterogeneous distributed systems. arXiv preprint
739 arXiv:1603.04467 (2016)
- 740 [44] Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T.,

741 Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van
742 der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nel-
743 son, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng,
744 Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman,
745 R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro,
746 A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy
747 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature*
748 *Methods* **17**, 261–272 (2020). <https://doi.org/10.1038/s41592-019-0686-2>