





Monitoring urban construction and quarry blasts with low-cost seismic sensors and deep learning tools in the city of Oslo, Norway

Andreas Köhler  * 1,2, Erik Myklebust  1, Anna Maria Dichiarante  1, Volker Oye  1

¹NORSAR, Kjeller, Norway, ²UiT - The Arctic University of Norway, Tromsø, Norway

Author contributions: *Conceptualization:* Andreas Köhler. *Methodology:* Erik Myklebust, Andreas Köhler. *Software:* Erik Myklebust, Andreas Köhler. *Formal Analysis:* Andreas Köhler. *Writing - Original draft:* Andreas Köhler. *Writing - Review & Editing:* Anna Maria Dichiarante. *Visualization:* Anna Maria Dichiarante. *Project administration:* Volker Oye. *Funding acquisition:* Volker Oye.

Abstract The aim of this study is to collect information about events in the city of Oslo, Norway, that produce a seismic signature. In particular, we focus on blasts from the ongoing construction of tunnels and under-ground water storage facilities under populated areas in Oslo. We use seismic data recorded simultaneously on up to 11 Raspberry Shake sensors deployed between 2021 and 2023 to quickly detect, locate, and classify urban seismic events. We present a deep learning approach to first identify rare events and then to build an automatic classifier from those templates. For the first step, we employ an outlier detection method using auto-encoders trained on continuous background noise. We detect events using an STA/LTA trigger and apply the auto-encoder to those. Badly reconstructed signals are identified as outliers and subsequently located using their surface wave (Rg) signatures on the seismic network. In a second step, we train a supervised classifier using a Convolutional Neural Network to detect events similar to the identified blast signals. Our results show that up to 87% of about 1,900 confirmed blasts are detected and locatable in challenging background noise conditions. We demonstrate that a city can be monitored automatically and continuously for explosion events, which allows implementing an alert system for future smart city solutions.

Non-technical summary Monitoring infrastructures and operations in cities relies on different kinds of sensors providing information for local authorities and the general public. In this study we collect information about events in the city of Oslo, Norway, that produce ground shaking. We focus on blasts from the ongoing construction of tunnels and under-ground storage facilities under populated areas in Oslo. We use data from sensors in the city, deployed between 2021 and 2023 for example in schools, to identify these blasts by means of machine learning methods. We are able to detect up to 87% of about 1,900 confirmed blasts.

1 Introduction

Global estimates for future growth indicate that the population of cities will continue to increase (Brockhoff, 1999). This growth has caused many cities to upgrade their infrastructures and to embrace the vision of a “smart-city” (McKinsey, 2018). Data collection through different types of sensors represents the base layer for such solutions. Large data sets are being produced and need to be automatically processed so that relevant information can be extracted and transferred to local authorities and the general public to facilitate decisions and to optimize the performance of cities in areas such as transport, safety and supply of water and energy (Fischer et al., 2013; Chang et al., 2014; Al Nuaimi et al., 2015).

Integrating seismic data into the data collection of such systems is currently not a common and widespread approach, although the potential of urban seismology using seismometers or Distributed Acoustic

Sensing (DAS) has already been recognized in previous studies (Ritter et al., 2005; Díaz et al., 2017; Spica et al., 2020). To date, this approach is routinely used mainly for earthquake early warning and fast response in urban areas with substantial seismic hazard (Kong et al., 2016), or for monitoring geothermal or other reservoirs in proximity to cities (Kraft et al., 2009; Hillers et al., 2020; Fiori et al., 2023). Advantages of using seismic data to monitor other urban activities compared to, for example, optical and acoustic systems include better compliance with General Data Protection Regulations (GDPR) (Zhang et al., 2017), efficient propagation of signals in the ground, independence of visibility, and in general a new type of sensor data not provided by the other methods.

This study focuses on the city of Oslo, Norway, addressing common needs of two departments of the municipality, i.e., the Agency for Emergency Planning and the Water and Sewage Department. The Agency for Emergency Planning is interested in obtaining quick information about any kind of unusual event that pro-

Production Editor:
Gareth Funning
Handling Editor:
Lise Retailleau
Copy & Layout Editor:
Hannah F. Mark

Signed reviewer(s):
René Steinmann

Received:
December 15, 2023
Accepted:
May 27, 2024
Published:
June 17, 2024

*Corresponding author: andreas.kohler@norsar.no

duces a seismic signature, e.g., explosions or sudden mass movements, to facilitate fast emergency response. An example of such an event was the bombing of a government building in the city center of Oslo during the terrorist attack on 22 July 2011 which was recorded on seismometers in and around Oslo (Bergen University, 2012). The Water and Sewage Department is concerned with monitoring ongoing construction activity to secure the freshwater supply of the city of Oslo. The construction of tunnels and under-ground water storage facilities under populated areas started in 2021 and is planned to be finished in 2028. Furthermore, due to population growth in Oslo, public transport infrastructure is currently extended, i.e., new metro tunnels are being constructed below or close to residential areas. Finally, a tunnel for a main electric power line under the city has been under construction since 2023. All these construction activities are accommodated by blasts which are partly felt by citizens, which have raised concerns in the population during a few documented incidents when the explosion yield was higher than anticipated.

Explosion monitoring with seismic sensors is a well-established technique for observing mining and quarry activities on a regional scale (Gibbons and Ringdal, 2006) or for verifying the Comprehensive Nuclear Test Ban Treaty (CTBTO) on a global scale (Kalinowski and Mialle, 2021). More recently it has also been used for identifying military attacks (Dando et al., 2023). A challenge with pursuing such an approach in urban areas on a very local scale and preferably in real-time, is the presence of a multitude of other seismic sources and high background noise levels. Such complex records require advanced processing methods which may be found in machine (ML) or deep learning (DL), fields which have made great advances within seismology in recent years (Kong et al., 2019; Bergen et al., 2019; Mousavi et al., 2019; Mousavi and Beroza, 2023; Zhu and Beroza, 2018; Mousavi et al., 2020; Provost et al., 2017).

In this context, there are two main possible approaches we can pursue: (1) Identification of so far unidentified seismic events of interest in an unsupervised manner or (2) using a sufficiently large number of already identified events of interest to train a classifier in a supervised manner. Approach (1) will be required in most cases as an initial step for urban monitoring purposes. It can be further divided into clustering, where the outcome are groups of signals or time windows of potential interest, or outlier detection, where the target is only a particular group of infrequent events. Clustering can be either done by automatically grouping the continuous seismic records (Köhler et al., 2010; Johnson et al., 2020; Chamarczuk et al., 2020; Seydoux et al., 2020; Steinmann et al., 2022a,b) or by grouping pre-detected transient signals (Sick et al., 2015; Jenkins et al., 2021). The features a clustering algorithm utilizes must be either extracted beforehand (e.g., Köhler et al., 2010) or are extracted automatically by a DL architecture (e.g., Mousavi et al., 2019). In a broader sense, simple non-machine learning methods, such as the well-known Short-Term Average over Long-Term Average (STA/LTA) trigger or trigger algorithms based on

other characteristic functions of the seismic waveforms (kurtosis, spectral amplitudes in different bands, etc.), may be considered to belong to approach (1). They can be used directly or combined with clustering for outlier detection. Hence, we can consider the STA/LTA method to be the baseline which ML or DL methods must outperform. In other words, under certain conditions STA/LTA may still be the most efficient way to detect events of interest.

In this study, we use passive seismic records acquired with the objective to quickly detect, classify and locate urban seismic events, particularly blasts. The use cases for detecting those events in near real-time include, but are not limited to, quickly informing the public in case of ground shaking felt by citizens or quickly identifying large blasts from construction works or attacks that can impact public safety and infrastructure integrity due to potential damage caused to structures (Shallan et al., 2014; Dowding, 2016; Naveen et al., 2021) or mobilization of unstable ground (Bouchard et al., 2018). For this purpose, a seismic network of low-cost sensors was deployed in target areas within the city of Oslo from spring 2021 onwards. We present a DL approach to first identify target events and then, if target events are sufficient in number, to build an automatic classifier from those templates. For the first step, we suggest an outlier detection method for automatic identification of rare events. These events are then located using their short-period fundamental-mode Rayleigh wave (Rg) signatures on the seismic network by means of stacking the observed travel-time corrected waveform envelopes. We then identify blasts inside and close to the city limits of Oslo and use them to train a supervised deep learning classification method to detect more of these events missed by the outlier detector.

2 The seismic network

We deployed three-component Raspberry Shake 3D sensors (Nugent, 2018) at different locations within the city of Oslo (Fig. 1, Table 1). The network was extended gradually starting in May 2021, with up to 11 stations recording simultaneously from June 2022 to July 2023. The sensors were connected to mobile network modems for real-time data transmission and remote maintenance. GPS antennas were deployed where possible. However, we found that the timing provided through NTP (Network Time Protocol) was sufficient at a few sites where free view to the sky could not be established. Sensor locations were in the basement of private businesses, private houses, and public school buildings. The first batch of sensors (ALNN1-4, ALNN7) was deployed with a dense layout in an area with quick clay in the sub-surface in the Eastern part of Oslo (Alna area) to allow for near-surface structural monitoring using ambient noise and detection of possible ground movements, a task which will be described in a future study. ALNN2 was removed after a few months in November 2021 due to construction activities in the host building. More sensors were deployed to the North of that area (ALNN5, ALNN6, ALNN8) and towards the city center and the Western part of Oslo (EKBG1, OSLN1-OSLN5).

The latter batch of sensors was located closer to the area of activity related to the construction of tunnels and an underground cavity for freshwater storage (Fig. 1), where the excavation is mostly done by blasting. All data are recorded continuously with a 100 Hz sampling rate, and the corner frequency of the sensors is at about 0.5 Hz. In addition, we use two permanent seismic stations equipped with broadband seismometers, one located on the university campus (OSL, part of Norwegian National Seismic Network) (Ottemöller et al., 2018, 2021) and the other one outside the city, to the Southeast of Oslo (OFNS2, not on map in Fig. 1).

3 Methods

Our aim is to detect rare or unusual seismic events observed on the deployed network using an outlier detection method. In contrast to a standard STA/LTA trigger approach, we do not want to simply detect all transient signals in the data stream. Frequently and regularly occurring urban events or noise bursts only recorded at single stations are not the focus of this study, although for example traffic monitoring with seismic data might be another topic of interest in urban seismology. Outlier events in our definition are singular or repeating events, but the latter not dominating the record, i.e., occurring not more than a few times per day. Hence, here we do not pursue a full clustering of all occurring signal and noise classes using ML, DL or other big data methods as done in previous studies (Köhler et al., 2010; Yoon et al., 2015; Seydoux et al., 2020; Steinmann et al., 2022a). However, it should be noted that clustering can be used for outlier detection. It would require to identify the event cluster of interest, i.e., the rare events. However, we decided to not pursue this approach further since we want to avoid the manual step to identify the outlier cluster. Furthermore, rare events may not necessarily be caught up in a distinct cluster.

Our workflow begins with identifying and collecting these repeating outlier events based on their origin locations and, if a sufficient number of observations has been collected, building a supervised classifier with labeled training data to more reliably detect those events in continuous data. Doing so, we have an outlier detection method available shortly after the start of the measurements which is flexible enough to pick up new events, while the supervised classifier can be gradually enhanced during the course of the seismic deployment by training it with newly identified events.

3.1 Outlier detection

Auto-encoder neural networks are popular methods for dimensionality reduction (Wang et al., 2016) and to identify anomalies or outliers in time series data (Yin et al., 2020; Thill et al., 2021). The idea is to use a Convolutional Neural Network (CNN) to reduce the dimension of the input time series, here three-component seismic waveforms with T time samples per component, using a series of convolutional layers or filters, and then use the resulting latent features to reconstruct the signal with a mirrored model neural network structure (Fig. 2).

In seismology this approach has been mostly adapted for data compression and interpolation (Navarro et al., 2019; Zheng and Zhang, 2020; Nuha et al., 2020). Furthermore, Valentine and Trampert (2012) highlighted the potential of auto-encoders for various applications in seismology. Mousavi et al. (2019) used an auto-encoder model to extract features suitable for unsupervised clustering. If data compression is the goal, the number of latent features should be low. However, our objective here is not primarily to compress data, and therefore we tested different dimensions from no compression at all, i.e., number of latent features is equal to $3T$, down to a latent dimension of T . Our final choice with best performance is a number of $2T$ latent features, i.e., a data compression by 33% (see supplementary Figure S1 for a comparison).

If the auto-encoder is trained using a continuous (unlabeled) seismic record, which is representative for a particular station, waveforms of regularly occurring signals and background noise should be reconstructed well by the model. If a signal is not reconstructed well enough, it can be considered to be an outlier. This approach has some relation to an auto-regressive model, which predicts a time series based on previously observed data and which is well known in seismology for its ability to detect the onset of seismic arrivals (Leonard and Kennett, 1999). However, similar to the STA/LTA method, an auto-regressive model is sensitive to all (including frequently occurring) transient signals with different characteristics than the background noise, a property which is not desired in our case.

We train auto-encoder models for single stations using the vertical and both horizontal components. Here, we use two stations with comparable low background noise levels as trigger stations: EKBG1 southeast of the city center and OSLN2 to the west of the city center (Fig. 1). The auto-encoder is trained for the two sites and then applied to time windows including STA/LTA detections obtained from the continuous data. Doing so, we aim to select only those STA/LTA triggered signals that can potentially be of interest. A future extension would ideally include outlier triggers on all stations. However, since in this study we are only interested in locatable events observed simultaneously on several stations of our network, we found it to be sufficient to trigger only on these two stations, since all locatable events are observed on at least one of these.

The auto-encoder input differs slightly for both stations, and also the selection of training data is done in a different manner. For OSLN2 the training data is a continuous record of eight consecutive days (02/06/2022–09/06/2022) band-pass filtered between 0.3–12.5 Hz. The size of the input time segment fed into the auto-encoder is $T = 512$ samples (see Fig. 2) for each component. For EKBG1 we use a higher number of samples ($T = 1024$), partly because this record visually appeared a bit more complex (frequent transients). A band-pass filter between 0.3 and 25 Hz is used to potentially also capture outliers with higher frequency content. The training data for EKBG1 are 90 time periods of continuous data of 6 hours' duration each, selected between 02/11/2021 and 14/03/2022. The

Name	Longitude	Latitude	Recording time	Location
ALNN1	10.8582	59.9336	08.06.2021-30.09.2023	Alfaset graveyard
ALNN2	10.8497	59.9282	09.06.2021-15.11.2021	private business
ALNN3	10.8452	59.9300	21.05.2021-30.09.2023	private business
ALNN4	10.8480	59.9314	09.06.2021-30.09.2022	private business
ALNN5	10.8336	59.9409	25.09.2021-30.09.2023	Linderud public school
ALNN6	10.8353	59.9405	29.09.2021-30.09.2023	Linderud public school
ALNN7	10.8464	59.9302	30.09.2021-30.09.2023	private business
ALNN8	10.8373	59.9411	15.11.2021-22.01.2023	Bjerke public school
EKBG1	10.7581	59.8974	03.11.2021-30.09.2023	Kongshavn public school
OSLN1	10.7694	59.9552	27.04.2022-30.09.2023	private house
OSLN2	10.7062	59.9415	01.06.2022-30.09.2023	Vinderen public school
OSLN3	10.7328	59.9425	01.06.2022-13.09.2023	Ullevål public school
OSLN4	10.6548	59.9415	24.10.2022-24.06.2023	Hovseter public school
OSLN5	10.7670	59.9650	10.06.2023-30.09.2023	private house
OSL	10.7227	59.9372		permanent NNSN station
OFSN2	10.9108	59.8401		permanent NORSAR station

Table 1 Seismic stations used in the study.

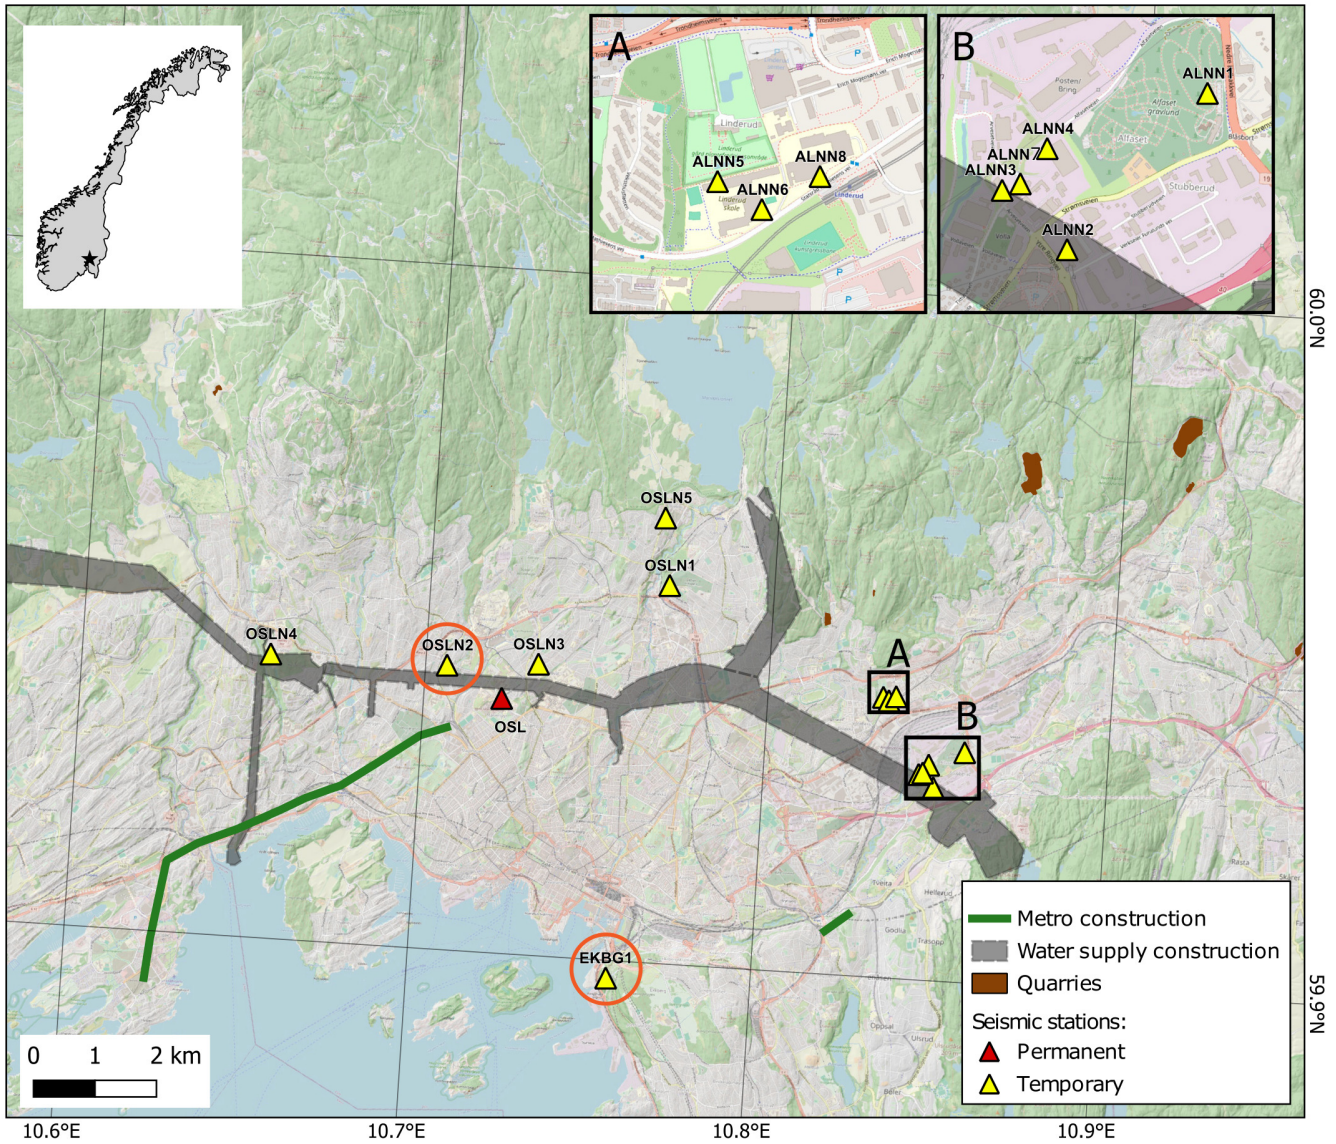


Figure 1 Map of the city of Oslo (OpenStreetMap contributors, 2017) draped on the Digital Elevation Model (DEM) at 10 m resolution, with infrastructure, potential seismic source areas and seismic stations. Map location is shown on the top left inset. (A, B) Close-ups of two areas in Eastern Oslo with denser seismic deployment. Stations which were used for STA/LTA triggering and outlier detection are marked with orange circles.

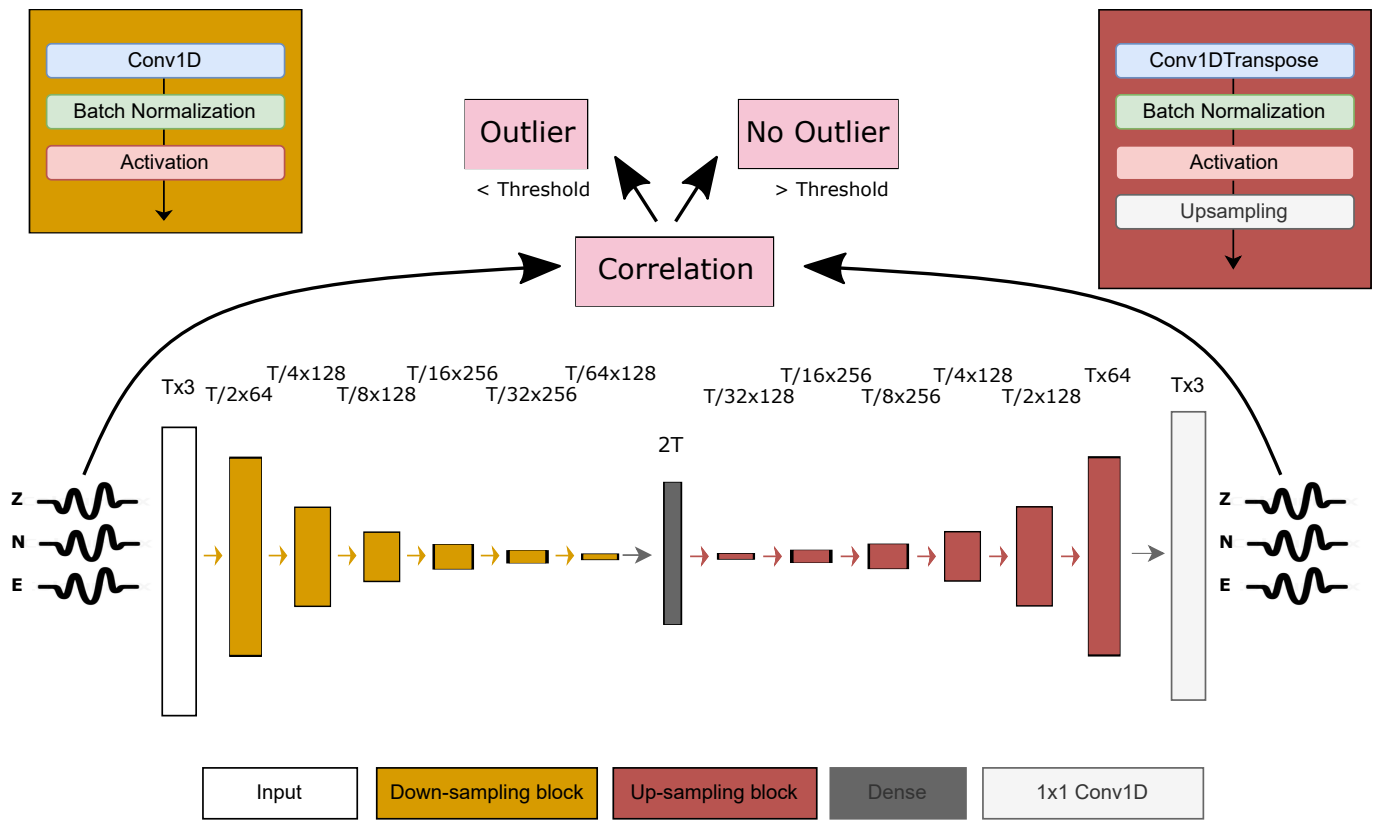


Figure 2 Auto-encoder architecture consisting of several down-sampling and up-sampling blocks shown in detail above the neural network. T represent the number of time-steps of one component. After down-sampling, the output is flattened and a dense layer is used to create the latent dimension ($2T$). The latent dimension is reshaped and used as input to the series of up-sampling blocks. Finally, a single convolutional layer is used to construct the output of the network. Each convolutional layer uses between 64 and 256 filters with a kernel size of 7. Input and output signals are cross-correlated to detect badly reconstructed events with low correlation coefficients, i.e., outliers.

motivation for this selection was to exclude visually detected blast signals from the training. This was achieved by manually screening a selection of days, from which the 6 hour-long time windows were then chosen. We did not pursue the same approach for OSLN2, which was implemented later, since we found that keeping the rare outlier events in the training data did not impair detection performance. In general, we found it to be more important that representative noise records are included (i.e., day and night, weekday and weekend), rather than making sure that outlier events of interest are excluded from the auto-encoder training. All waveform time windows are normalized with minimum and maximum amplitude before being fed into the auto-encoder.

We apply an STA/LTA detector with a low threshold (STA length = 0.5 s, LTA length = 10 s, STA/LTA threshold = 4) to continuous three-component data, and if all three components exhibit a coincident trigger, a detection is declared. We then apply the auto-encoder method to a time window around each detection (same duration as training data). One way of evaluating the quality of the auto-encoder signal reconstruction, i.e., to identify an outlier signal, is to compute the normalized correlation coefficient of the original and reconstructed seismic traces (Fig. 2). We also tested using the reconstruction loss, i.e., RMS value of the observed and reconstructed waveforms, for outlier detection. However,

we found the correlation value to be more suitable to identify outlier events which are mostly blasts in our case. A comparison of RMS and correlation values of all STA/LTA detections and confirmed blast signals at station OSLN2 is provided in the supplementary Figure S4, showing that the RMS is not a good discriminant for blasts in our case.

Figure 3 shows examples of one outlier event and three STA/LTA detections not recognized as outliers at each station (OSLN2 (a-d) and EKBG1 (e-h)). The correlation between original and reconstructed traces for most data is in general very high, i.e., larger than about 0.8. Using a lower/higher latent dimension would compress the seismic data more/less, which would increase/decrease the construction loss and decrease/increase the correlation. We found the current dimension of the auto-encoder to be optimal for our task. However, it must be tuned for each new data set. Figure 3a and e show seismic signals which were later confirmed as blasts. Particularly, the vertical component waveforms are not well reconstructed. Hence, both exhibit comparably low correlation coefficients (0.6 and 0.84) in relation to the other STA/LTA detected transient signals. The latter are manually identified as regularly occurring noise bursts, including signals originating from very close to the sensors.

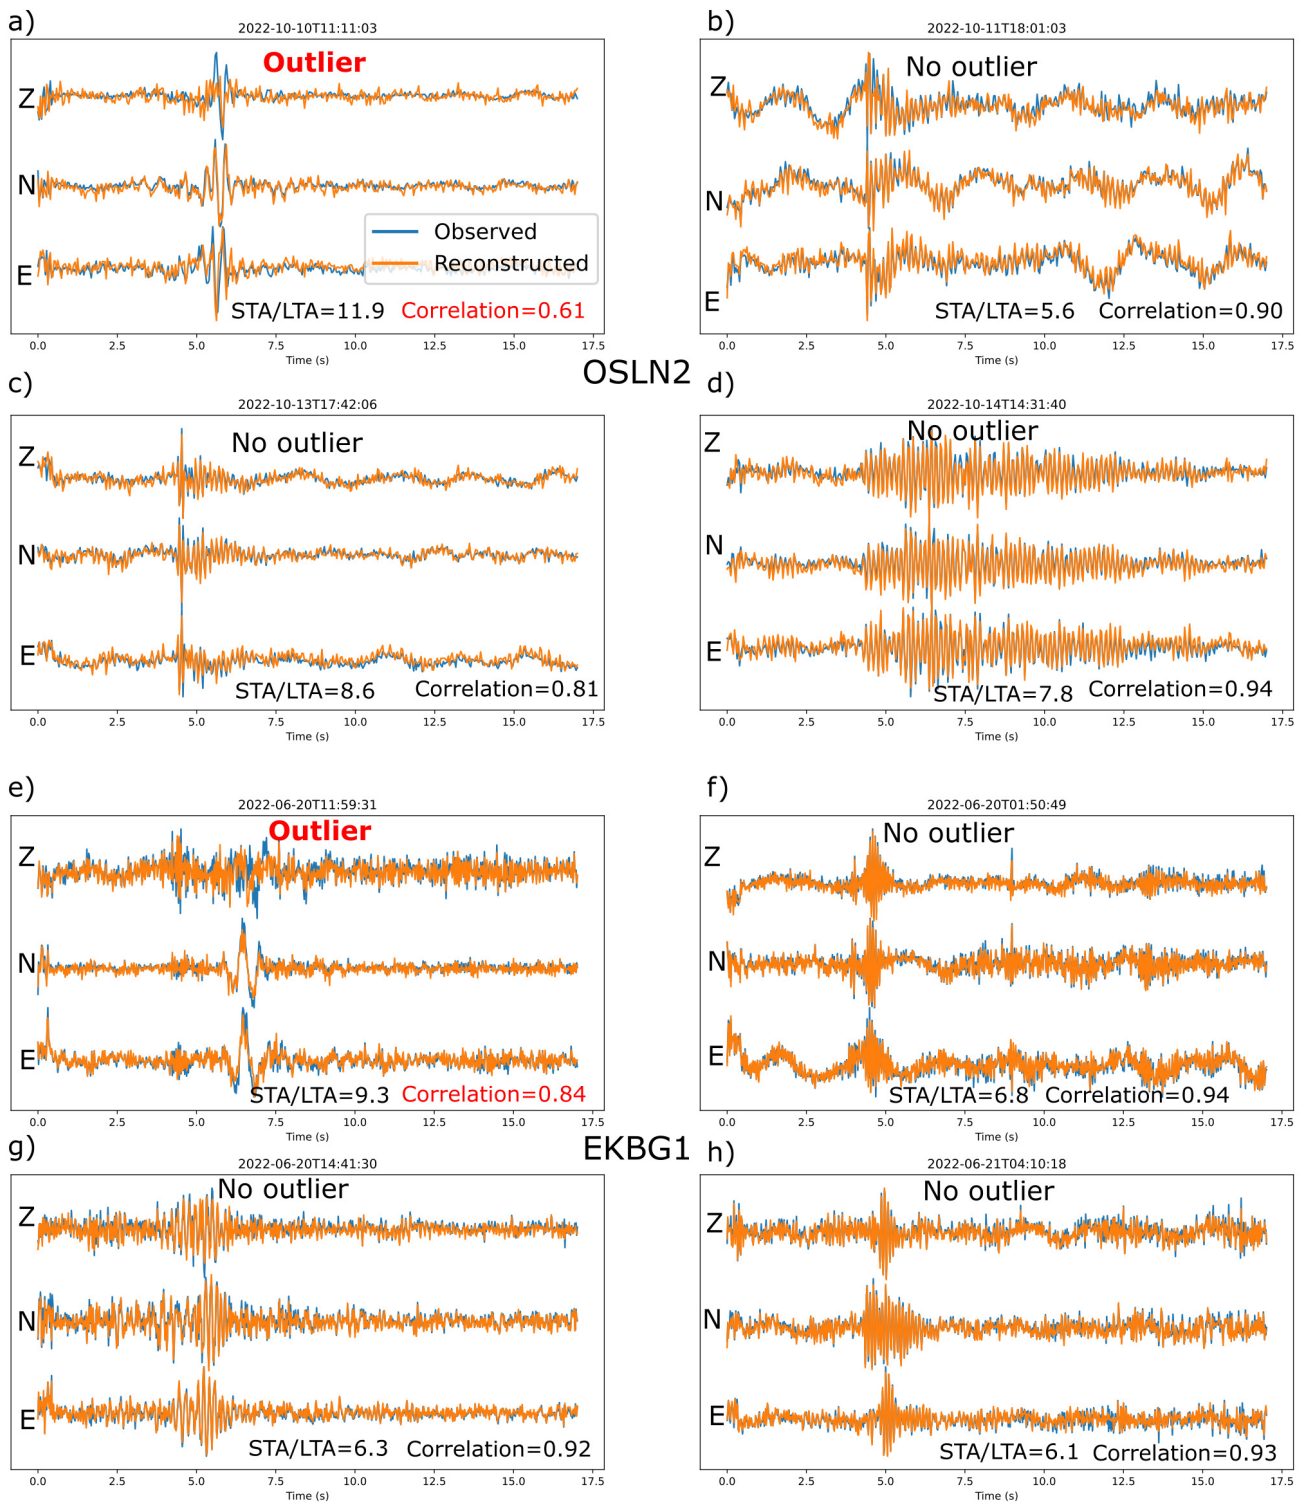


Figure 3 Comparison of recorded three-component waveforms (blue) and waveforms predicted by auto-encoder (orange) for station OSLN2 and EKBG1. STA/LTA value and correlation coefficient between traces are given in each panel. Two detected outlier events, marked red in a) and e), are confirmed blast signals. The other transient signals triggered by the STA/LTA method were not detected as outliers due to higher correlation.

3.2 Locations based on Rg waves

Once an outlier is detected at station OSLN2 or EKBG1, the remaining stations are used in addition to attempt an automatic location. P and S wave arrivals would be needed for traditional event location based on onset time readings, but are not observed for the majority

of events due to high noise levels. However, we found the Rg wave, which is a short-period Rayleigh wave in the Earth crust typically observed for seismic sources close to the surface, to be well recorded over the entire network between 0.8 and 2 Hz. In order to use Rg for event location, we first compute envelopes of

the band-pass filtered vertical component data and discard stations with Signal-to-Noise Ratios (SNR) below 7.0 (OSLN2) and 6.2 (EKBG1). If a minimum of four stations are left, we perform a 2D gridsearch to find the event source location maximizing the stacked time-shifted envelopes of all stations, an approach similar to stacking and migration methods developed for seismic event localization (Gharti et al., 2010). For Rg travel time computation we assume a velocity of 2.0 km/s which we found to fit observed Rg waves with plane wave fronts travelling over the network from a known source at a large distance. The extent of the gridsearch is from 10.55 to 11.0 degrees longitude and 59.86 to 60.0 degrees latitude. The step width is 0.01 degrees in longitude and 0.005 degrees in latitude. In addition to maximizing envelope stacks, we estimate the Rg back-azimuth from three-component data for the stations above the SNR threshold. This is done by finding the rotation angle maximizing the amplitude on the radial component. The 180 degree ambiguity is avoided by selecting the direction whose correlation coefficient between the vertical and Hilbert-transformed radial component of the Rg wave is positive. The weighted back-azimuth residual of each location grid point is subtracted from the stacking amplitude. The grid point maximizing this value is taken as the source location. Based on the location, we assign to each confirmed and locatable outlier a label corresponding to different construction areas in the city of Oslo.

3.3 CNNs for blast classification

The supervised classifier is a Convolutional Neural Network (CNN) which takes three-component waveform data of a single seismic station as input (Fig. 4). The method uses the well-established AlexNet architecture (Krizhevsky et al., 2012) and is loosely based on the model we used in Köhler et al. (2022) to classify calving events in the Arctic. We train a two-class model distinguishing STA/LTA detections of blasts in Oslo and all other detections (noise and other events). Here, we only use station OSLN2 to train and test the classifier. The model consists of a layer to randomly crop the input waveforms, five convolutional layers with batch normalization and max-pooling, and finally two dense layers which process the flattened output of the convolutional layers and generate the output probabilities. We use 26 s as input time window duration around each blast which is cut to 22 s by random cropping. For the noise class we use a time window of waveform data of the same duration before each blast detection such that the classes are balanced. The hyper-parameters controlling size of convolutional filters and type of pooling are tuned with KerasTuner (Chollet et al., 2015). In the tuning process, the number of filters in each of the five convolutional layers was kept constant. KerasTuner uses ranges of hyper-parameters (filter length between 3 and 49) and different options (max vs. average pooling) as input and searches the parameter space to optimize the classification accuracy. The final hyper-parameters are shown in Figure 4.

For the final classifier, we use a stratified 5-fold cross

validation, i.e., five different models are trained, each using 80% of the shuffled data (confirmed blasts from the outlier detection and the noise class examples) for training and 20% for validation. When applying the classifier for prediction, the averaged probabilities for blast and not blast of these five models are used. Finally, we have to set a probability threshold for detecting a blast. We can either use a threshold of 0.5, i.e., select the winning class, or require a higher confidence for a blast to be detected using a higher threshold.

3.4 Reference blast detections based on STA/LTA method

For evaluating the outlier detection method and the blast classifier, we would need complete ground-truth data about blast occurrence in the city of Oslo, which turned out to be difficult to obtain. Alternatively, we can visually screen all potential seismic events observed on the network. This will not allow us to assess the network's detection sensitivity, but rather the detection method's ability to recognize all recorded blasts. Since our methods use STA/LTA detections as input, we create our reference event catalog by processing all these STA/LTA detections in the same way we process the outliers, i.e., attempt a Rg wave based localization and label the recognized blast signals in the different parts of the city. This resulted in 1,870 blasts located within the study area between November 2021 and October 2023.

We use the recall and precision metrics to evaluate all deep learning models with respect to this data set. We want to achieve a high recall (high number of recognized blasts) and high precision (low number of false triggers). In contrast to a conventional event detector, the decision on what is a false and what is a true positive is a moving target for an outlier detector. Additional events not being part of the reference data could still be events of interest. Nevertheless, we can still use the recall-precision metrics as a proxy to compare model performance relative to each other. To decide on decision thresholds for the outlier detector (correlation coefficient) and blast classifier (blast probability), we evaluate recall-precision curves provided in the supplementary material. We compute the performance metrics before event localization since we want to include unlocatable events.

4 Results of outlier event detection

After visual inspection of many signals detected on multiple days as well as evaluating recall and precision, we finally set the outlier detection threshold to 0.86 for EKBG1 and 0.78 for OSLN2. Supplementary Figure S1 shows recall-precision curves for outlier detectors at OSLN2. The optimal detection threshold corresponds to correlation coefficients producing a recall-precision point closest to the upper-left corner. Note that precision will increase after applying the automatic location procedure because false alarms are usually unlocatable events. The distribution of correlation coefficients at stations OSLN2 and EKBG1 in relation to the




Figure 4 Architecture of the deep Neural Network of the blast classifier.

chosen thresholds are provided in supplementary Figure S5.

Figures 5, 6 and 7 show maps of stacked envelope amplitudes and the best location estimates for several blast signals detected as outliers. The events in Figure 5a–c are blasts related to the construction of underground water tunnels and storage facilities at the Stubberud, Oset, and Huseby sites (see also Fig. 1). Figure 6 shows a blast at the Losby quarry East of Oslo, a blast at the metro tunnel construction site at Bryn, and a blast from a construction site for a new electric power line tunnel in Sogn. For the latter we have a ground-truth confirmation of location and blasting time. Figure 7 shows three blast signals for which we also have ground-truth times and locations. They originate from the construction of the new metro tunnel for the Fornebu line in the West of Oslo, about 3 km south of the underground water storage construction site. The precise locations were provided to us and the public with an uncertainty of about 100 m.

The spatial distribution of stacking amplitudes shows as expected that the resolution strongly depends on the event location. Outside of the network, resolution is poorer and consequently the blast locations are not well-constrained. This can be observed as broader amplitude maxima to the West and East of Oslo, as well as biases with respect to the ground-truth locations. Note that since OSLN4 was not deployed before October 2022 and produced corrupted data from summer 2023 onwards, poor resolution is also expected for locations at the Huseby construction site for events outside this time interval. As a consequence, blasts from the Fornebu metro tunnel and the water storage site cannot always be discriminated. However, as the comparison with ground-truth location shows, the accuracy is very good when the entire network was in operation (see Fig. 5, 6 and 7).

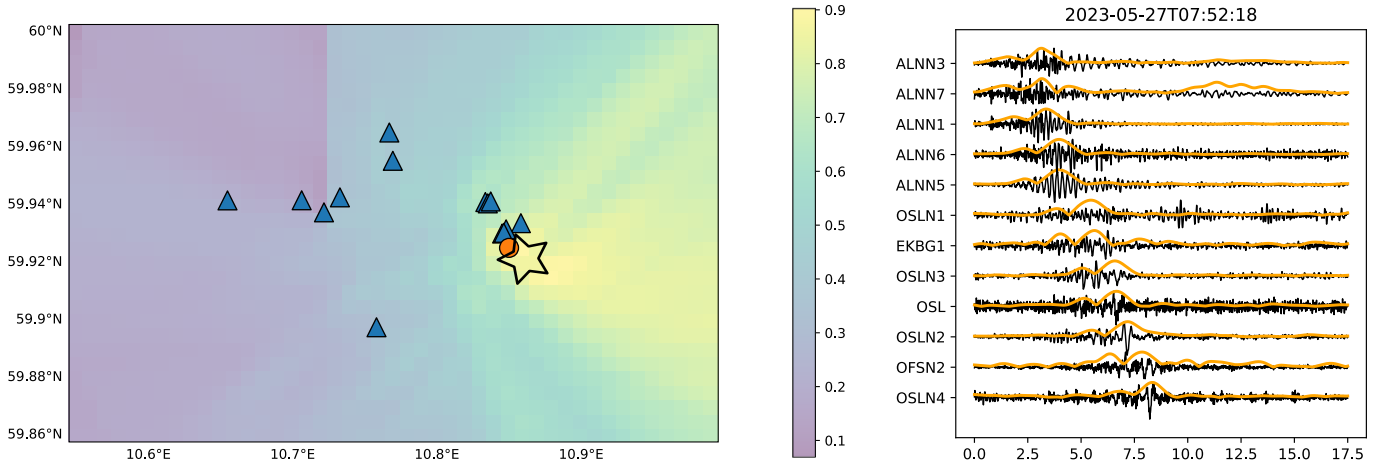
Figure 8a shows all locatable outliers triggered at OSLN2 and EKBG1 in the study period (1,272 events). Two different symbol colors are used to distinguish the time period of complete and incomplete station coverage (see Table 1). Almost all events are located inside or close to areas of known construction or quarry blast activities. Clusters of events are indicated, of which we have already seen examples above. By far the most

blasting activity is observed to the West of the city center, i.e., the Huseby and Fornebu constructions sites. Events are well-located during complete station coverage, while a number of blasts tend to be falsely located westwards from Huseby when the westernmost station OSLN4 was not in operation. Note that we have filtered out distant events (blasts from distant quarries and regional and teleseismic earthquakes) since they are usually falsely located at the edge of the grid search region or in the center of the network if the incidence angles are steeper.


Figure 8b shows the corresponding time line of locatable outliers. Before July 2022 only stations with high noise levels in the east of Oslo were in operation and consequently only a few events are observed. For the rest of the study period, there is a lot of blasting activity with up to 10 events per day and about 4 per day on average. Pauses in the blasting during public holidays (Easter and Christmas break) and school holidays during summer are clearly visible. Figure 9 presents more detection statistics for all located blasts. Time of day and day of week distribution are consistent with blasting which usually ceases on Sundays and during nighttime (Figure 9d and e). Local seismic magnitudes of blasts are between -0.5 and 1.5 (Figure 9c).

The temporal distribution of blasts in the reviewed reference data set is shown in the background of Figure 8b. A high percentage of these events are recognized as outliers (69%). However, 31% of visually identified blasts are not found (i.e., false negatives; 520 events). A closer look at the distributions of STA/LTA ratios of all events in Figure 9a and b, as a proxy for SNRs, gives an explanation for those results. The distribution of ratios are shown for all locatable outlier events, for all detected outliers (including those that were not locatable), and for all STA/LTA detections. Note that logarithmic scales for number of detections are used. The comparison shows that STA/LTA detections at OSLN2 with STA/LTA ratios above 10 are almost all classified as outliers, the majority being also locatable as blasts (Fig. 8a). In other words, there would be no need for an outlier detection method for those events, and we could simply use the STA/LTA detections directly to monitor blasts. However, towards lower SNRs the picture changes. We obtain an increasing number of STA/LTA detections, the

a) Construction blast Stubberud



b) Construction blast Oset



c) Construction blast Huseby

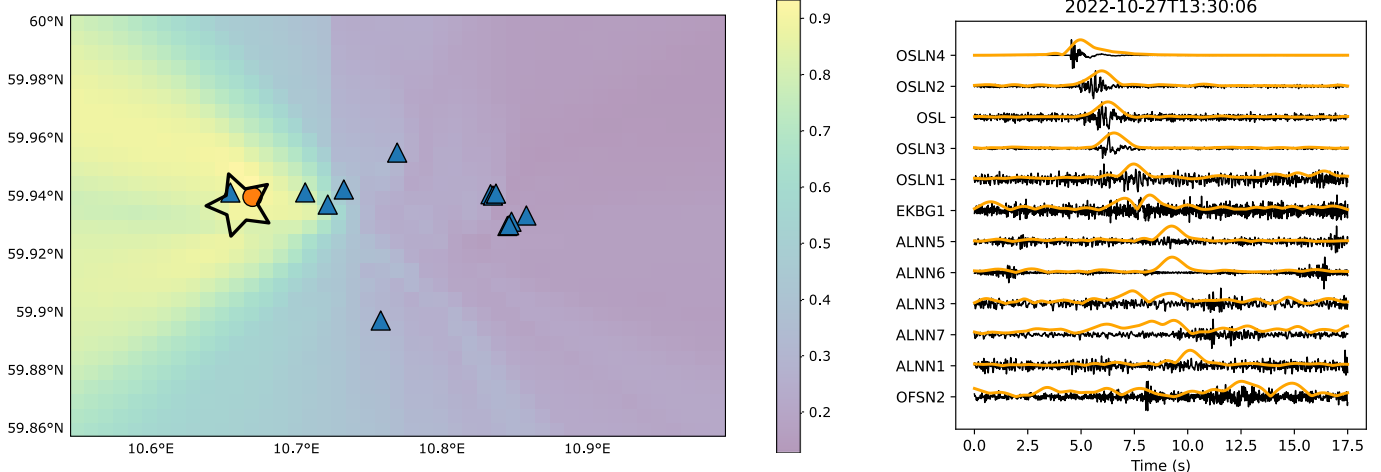


Figure 5 Map of Oslo region overlaid with stacked envelope amplitudes (normalized). High amplitudes indicate more likely event location. The best location (orange circle) and seismic stations (triangles) are shown. The star symbols and their extents indicate the areas of known blasting activity. On the right-hand side the vertical component seismic data for all stations are shown. Orange data indicate low-pass filtered envelopes enhancing Rg arrivals.

majority not being locatable events, probably mostly local noise bursts. The outlier detection method allows us to reduce the number of detections to be screened for location considerably. For the lowest SNRs we obtain a number of locatable STA/LTA detections which turned out to be blasts not detected as outliers (red bars). This indicates that our outlier detection method has limi-

tations in recognizing weak events. We will deal with these missed events when applying the supervised blast classifier. It is worth emphasizing that visual inspection of the unlocatable outliers at OSLN2 revealed clear blast signals which were not observed on more than three stations. Four examples are included in supplementary Figure S6. This shows that the outlier detector com-

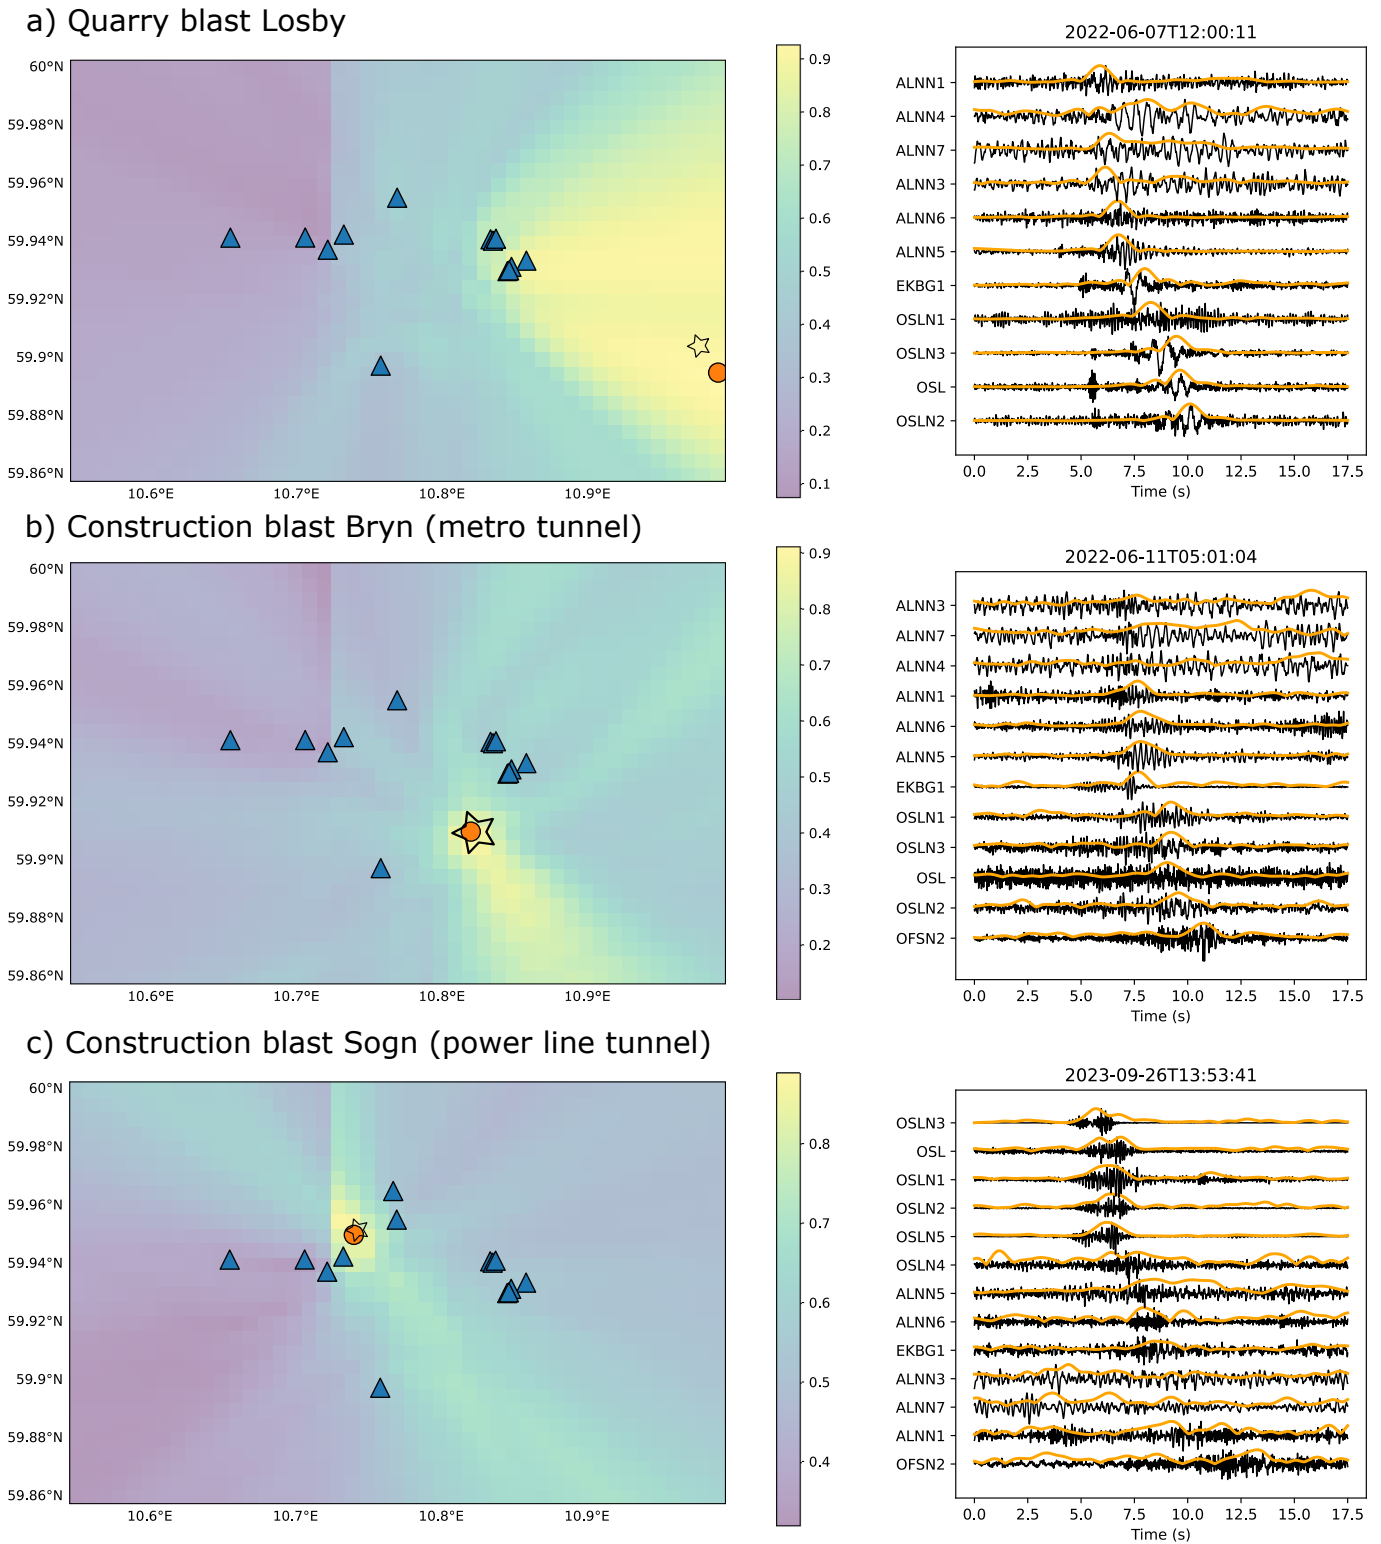


Figure 6 Same as Figure 5. Note that smaller size of star in (c) indicates more certain (ground-truth) blast location.

combined with a denser station network would have recognized even more blasts.

For EKBG1 (Fig. 8b) there are a few more detections with high SNRs that are not classified as outliers. The major difference to OSLN2 is that fewer outliers turned out to be locatable events which could be identified as blasts. In other words, a lot of outliers are seismic events at EKBG1 which are not observed on other stations. Given that we know that blasts are usually picked up by at least two more stations, it is likely that these are

mostly local noise bursts around EKBG1 that cannot be explained by normal background noise fluctuations and hence are not well-reconstructed by the auto-encoder. However, it is worth emphasizing that many STA/LTA are still removed by outlier detection, which reduces the amount of detections to be screened for possible localization.

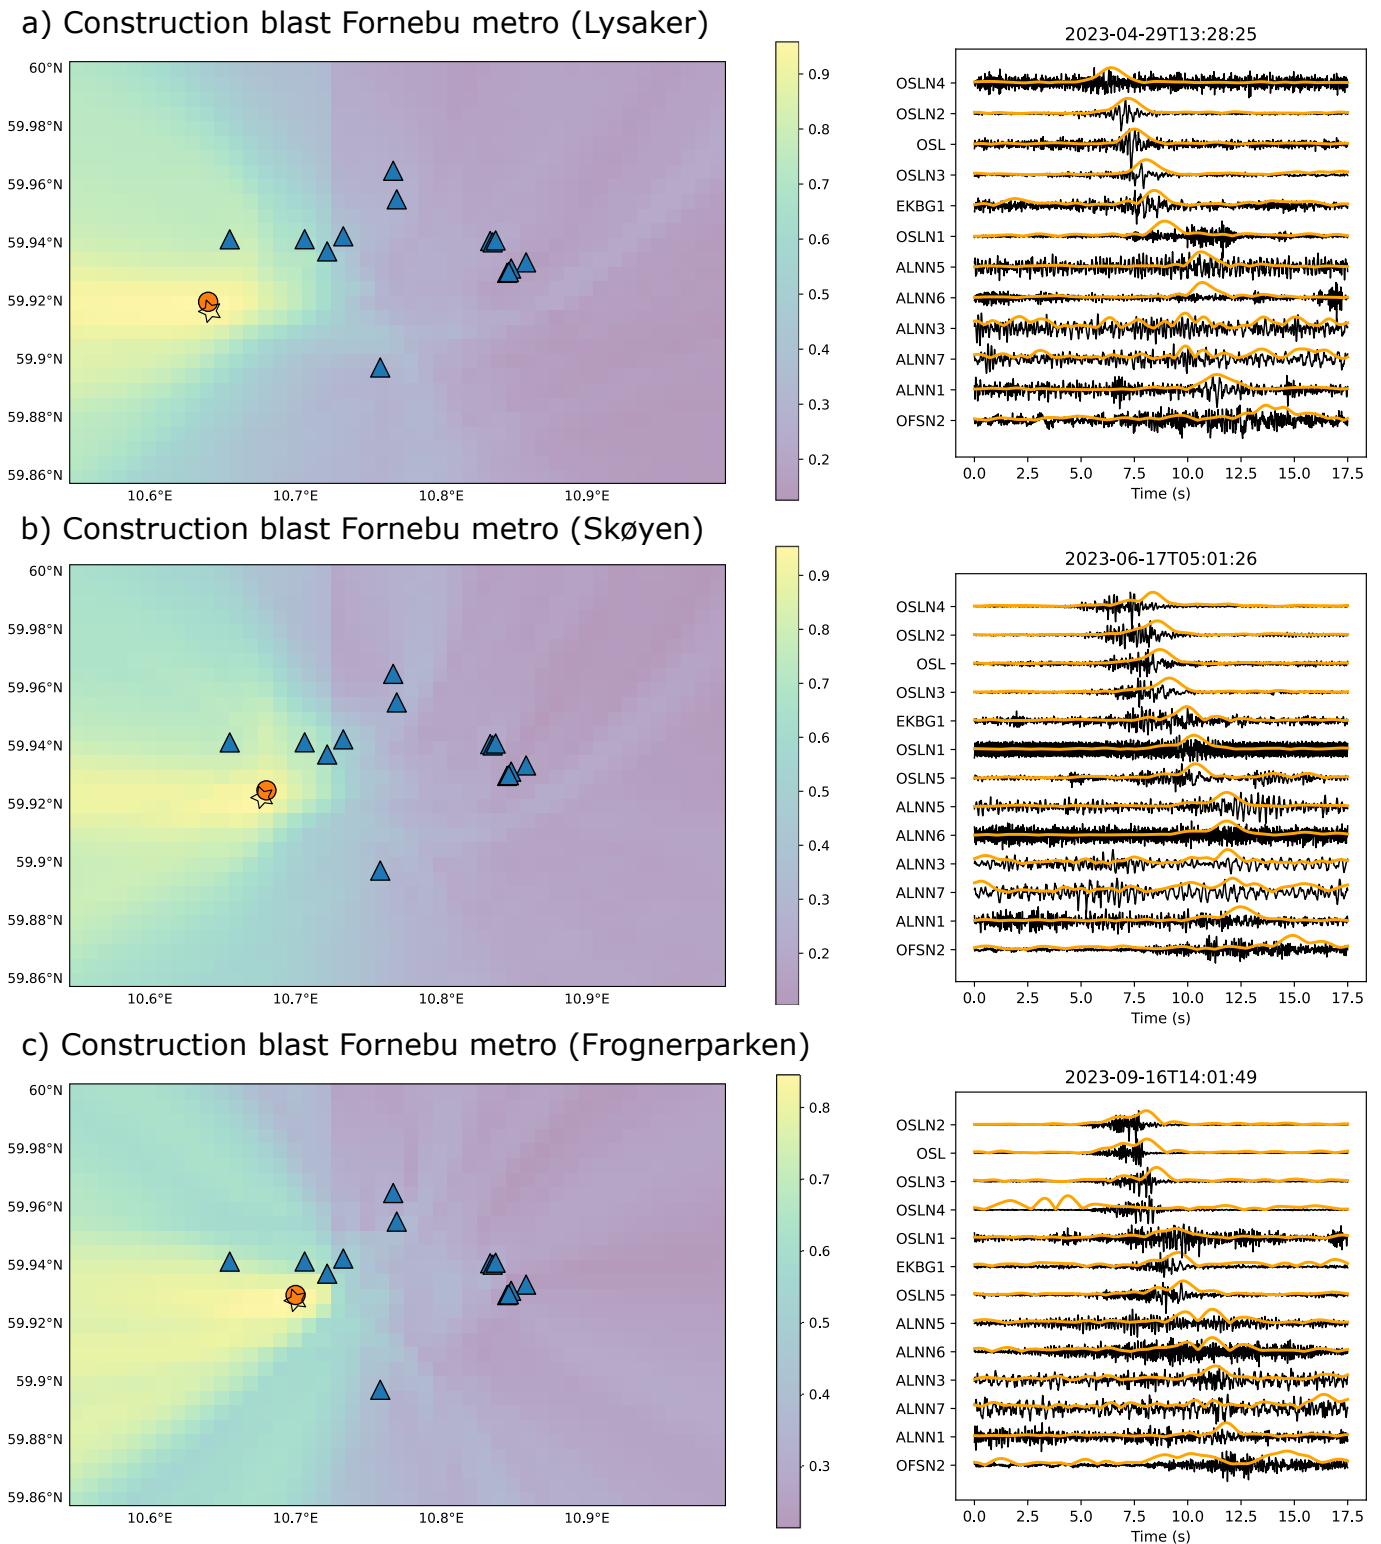


Figure 7 Same as Figure 5. Note that smaller size of stars in all panels indicate more certain (ground-truth) blast locations.

5 Results of blast classifier

Ideally, we need to train a classifier with data from several stations so that it generalizes well enough and can identify blast signals in the data from different stations. As also discussed below, the reason for starting with a classifier trained for a single station (OSLN2) is the limited number of blasts observed on all stations which results in an unbalanced training data set with respect to event location and observing stations. However, a few

stations close to OSLN2 have a comparable number of observations and could be included in the training. We will come back to this in the discussion section.

We train the blast classifier with waveforms from OSLN2 including all 1,272 blasts from different areas in Oslo detected as outliers at OSLN2 and with the same number of noise examples. The reason for not using all 1,870 signals in the reference data set is that we want to simulate a workflow where we only train with blasts previously detected by the outlier detector, excluding those

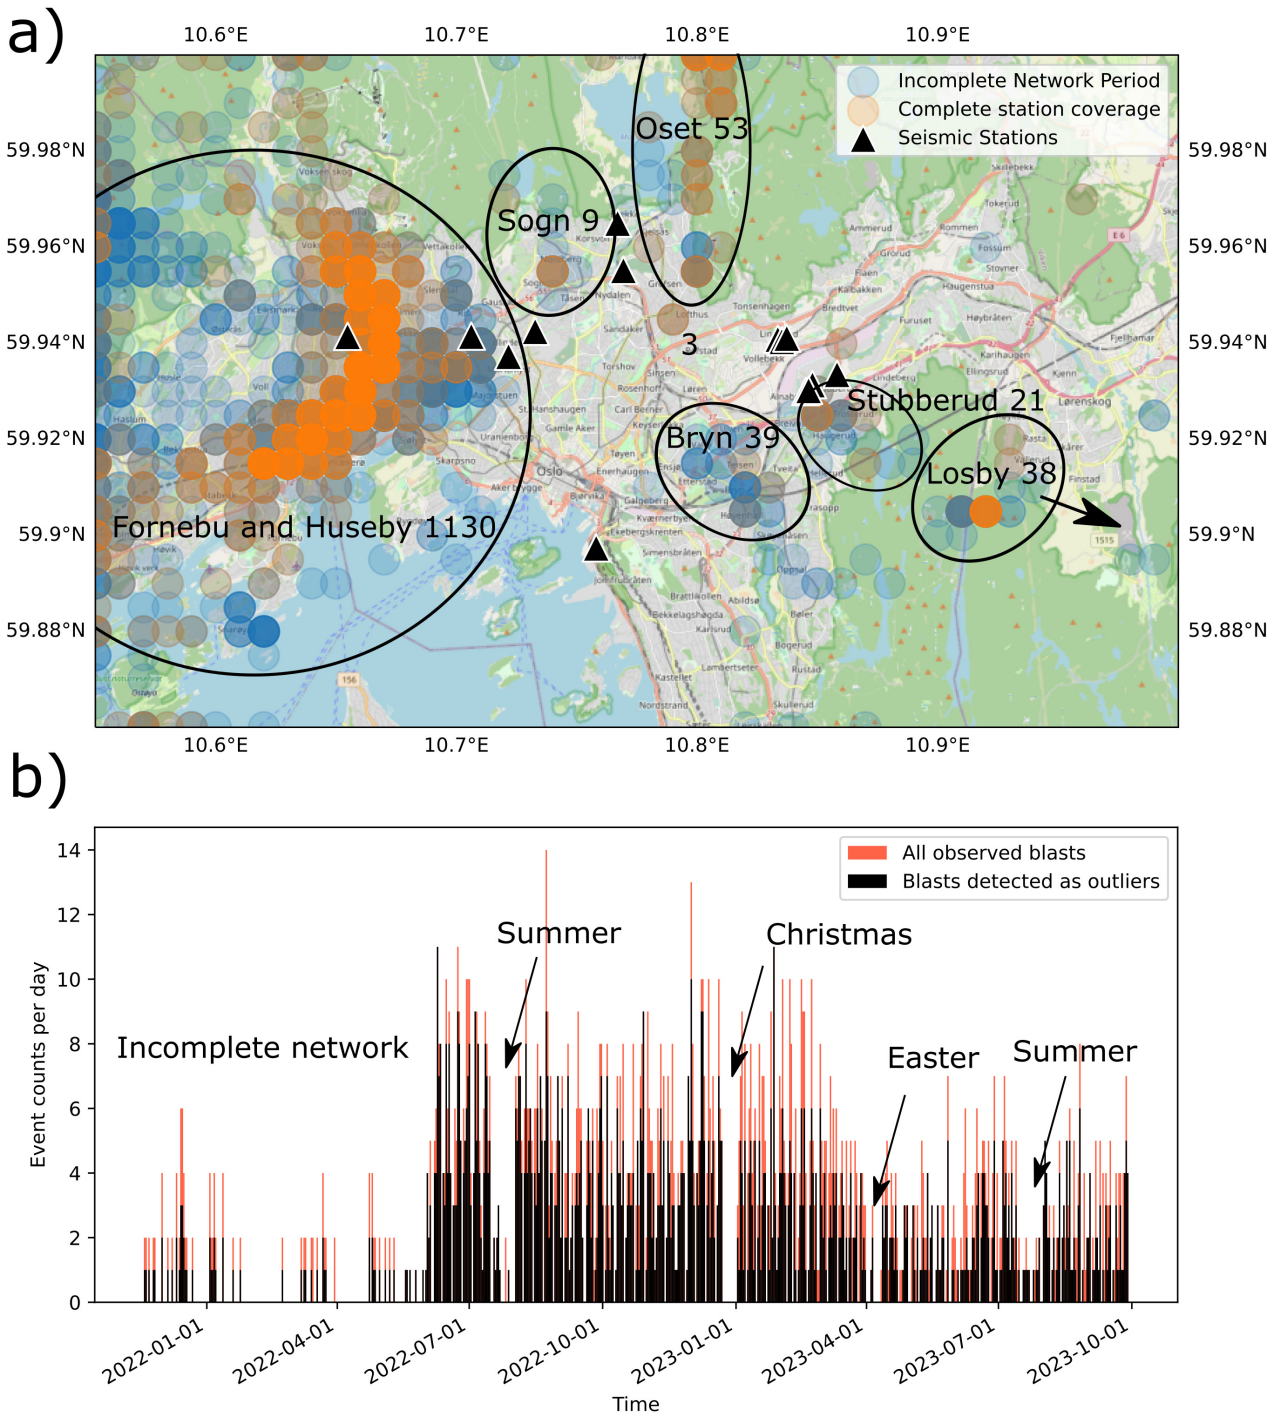


Figure 8 a) Located seismic events from outlier detection. Clusters and number of blasts are indicated. b) Time line of those events together with all observed and manually confirmed blast STA/LTA detections.

that we only identified after screening all STA/LTA detections. The classification performance metrics of the best model of all five folds using data not used for training are shown in Table 2. We achieve high values for both precision and recall.

Next, we apply the classifier to all 29,058 STA/LTA detections at OSLN2 in the time period from 01/06/2022 until 28/09/2023. We use a probability threshold of 0.5, i.e., picking the winning class (blast vs. not blast) as well as 0.7 to test different confidence levels. Figure 10a shows the time line of 1,385 classified and locatable blasts using a threshold of 0.7 together with the refer-

Class	Blast	Not blast
Precision	0.92	0.95
Recall	0.94	0.93
F1	0.93	0.94
Accuracy	0.93	

Table 2 Performance of blast classifier on validation data.

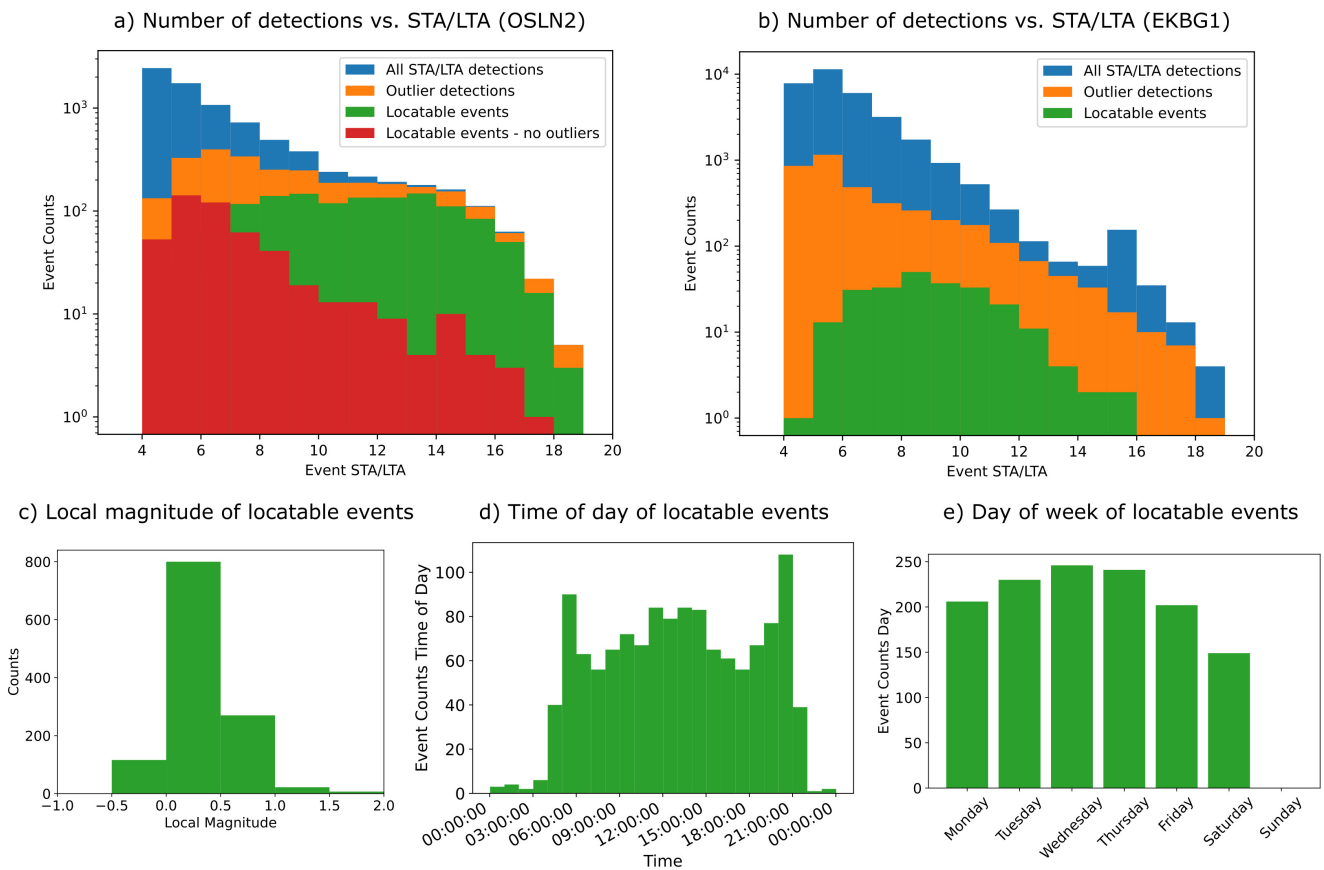


Figure 9 Statistics of STA/LTA and outliers detections, and locatable blasts.

ence data set. In comparison with the outlier detector (Fig. 8b), more blasts are recognized. When we use our reference data set, which includes 1,870 blast manually identified by screening all locatable STA/LTA detections, as ground-truth for evaluation, the recognition rate increases from 69 to 80%, and the number of False Negatives decreases to 371 events. A number of 22 False Positives are either blasts outside the study area, which we did not include in the reference data set of locatable blasts, or are local signals at OSLN2 which are randomly associated with Rg wave-like signals at other stations. Of all STA/LTA detections, 424 events are classified as blasts but are not locatable. As with the outlier detector, these events are not necessarily false, but are simply not observed on more than three stations, which would allow for a reliable location and for being included in the reference data set (supplementary Figure S6 shows examples). In fact, we checked a selection of these detections manually and found that almost all of these show clear blast-like signatures at station OSLN2 and OSLN3. Hence, in relation to the high number of tested STA/LTA detections (29,058), the actual number of false classification is negligible if the goal is to provide real-time information about ongoing blasting at construction sites. However, if the goal is early warning in case of unusual events (accidents, attacks), any false detection should be avoided and other data have to be included before issuing an alert.

For a probability threshold of 0.5 the blast recognition rate increases further to 87%. As the maps in Figure

10b–c show, this is partly due to more of the underrepresented blasts in the north and east of Oslo being correctly classified (compare Fig. 10b and c). It is expected that those events yield a lower blast probability since the training data is unbalanced with respect to event location. However, there are also more False Positives (77) and about 100 more unlocatable events (519) compared to using the high threshold. Visually inspecting those 100 additional events revealed that about 30% look like blasts, but are not locatable because they are observed on only one or two stations. However, the rest (70%) are now actual false detections which we would avoid with a higher threshold.

6 Discussion

We present a prototype for an automatic urban seismic monitoring system which identifies any potentially interesting event as well as routinely detects previously identified events. Our system is based on a low number of low-cost seismic sensors and was running for almost two years. We demonstrate that with comparably low effort when it comes to upgrades of the sensor infrastructure, a city can be monitored continuously for explosion activity. In our case this includes construction blasts, but there is no reason for other types of explosions, such as accidents or deliberate attacks, not to be detected as long as there is sufficient coupling with the ground.

An alternative approach to identify events of inter-

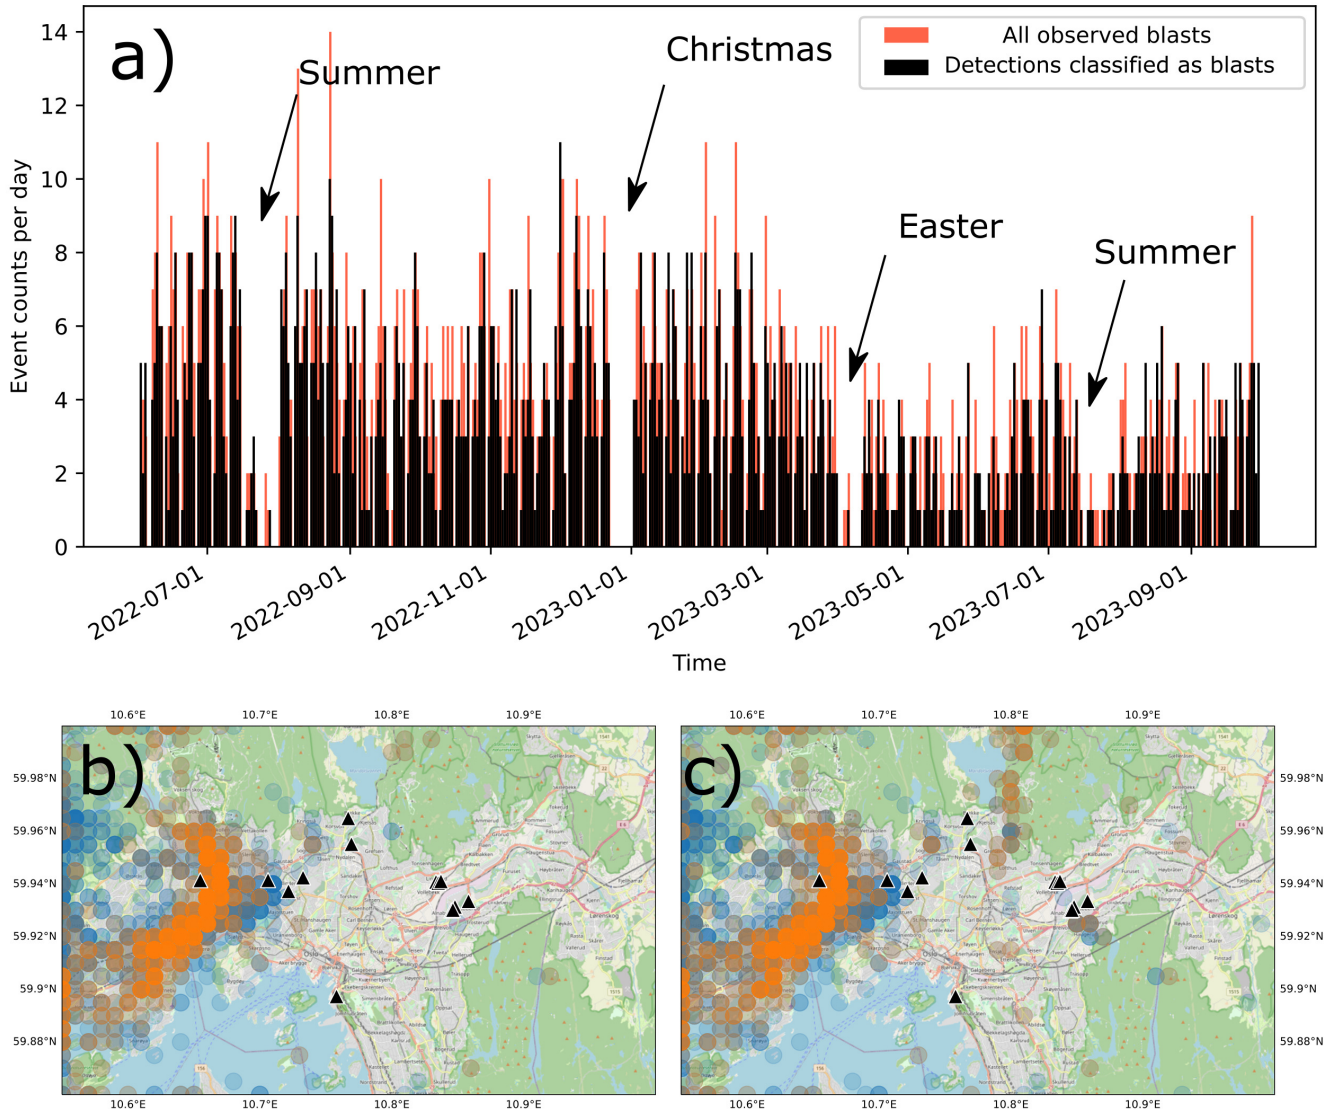


Figure 10 a) Time line of locatable STA/LTA detections classified as blasts by the CNN model in comparison to reference blast data set. b–c) Locations of classified blasts using a probability threshold of 0.7 (a) and 0.5 (c). Orange symbols are for complete and blue symbols for incomplete network operation. Black triangles show seismic stations.

est in an unsupervised fashion would be to apply recently developed (deep) clustering methods (e.g., Seydoux et al., 2020). While adapting such methods and comparison of the results with our method would go beyond the scope of the current paper, we encourage further studies to compare both approaches. The main reason why we did not choose a method to identify outlier events via clustering, is the additional need to identify the cluster(s) of interest and the risk that outlier might be too rare to form separate clusters. We found that an auto-encoder is a relatively easily implemented and trained alternative (compared to more sophisticated deep learning models), and is not difficult to tune for each station. The only tuning we did for each new station was the selection of the training data, the number of samples in the input time window, and the latent dimension if the auto-encoder. We identified the sample number and latent dimension to be the most important parameters, while the rest of the model archi-

ture, hyper-parameters, and the choice of training data (except that it should cover different noise conditions at different times) does not need to be adapted for each station.

There are different possibilities for improving the system. First and most important, the seismic station network can be extended by covering a larger area and by increasing station density. This will improve location accuracy, especially outside the current network area beyond the city limits. A denser sensor deployment will also enable locating more events that so far are only observed on a single or two stations and are, therefore, not locatable with our Rg wave stacking approach. This would also allow us to potentially run the outlier detector on more than two stations simultaneously and to locate the detected events with more stations that are close-by.

Secondly, the detection process could be further improved. The outlier detector could be retrained regu-

larly since local noise conditions around each station may have changed over time. Furthermore, more work can be done to further tune the neural network architecture of the auto-encoder to optimize outlier detection at different seismic stations. One could also further investigate if the same outlier detector model can be applied to different stations despite of station-specific background noise conditions and sub-surface-related site effects. To test this we already trained an additional auto-encoder model with combined data from stations OSLN2 and OSLN3. We then applied this and the original model for OSLN2 to both stations and compared the results. The outlier detector performance confirmed our choice to make the detector station-specific (see recall-precision curves in supplementary Figure S2). While we found best performance with the current model, we acknowledge the high number of latent features. However, this is not an issue as such since our goal is not data compression but outlier detection. A more systematic study of outlier detector performance for decreasing the number of latent features beyond the tests we presented here may help to further optimize the detection rate.

We train the supervised blast classifier with a comparably low number of data points. Longer recording periods will increase the available data and, thus, most likely improve classifier performance. Alternatively, it is possible to augment the existing training data with noise or other seismic signals (Köhler et al., 2022). This is of particular importance for areas with infrequent blasts which are currently not well-represented in the training data and are therefore less likely to be correctly classified (events in the east of Oslo). We only used a single station to train the classifier. Consequently, the model learns station-specific features and does not generalize well. However, ideally we would need a classifier which generalizes well enough to detect blasts on all stations. We started to explore training a single classifier with waveform data from all stations, either using independent input data (three-component waveforms from different stations) or using multiple channels from different stations in one data sample as input. However, we found that both approaches require more and better balanced training data when it comes to blast detection. We trained a model using data from two stations in the west of Oslo (OSLN2 and OSLN3). Supplementary Figure S3 shows that the generalization ability is not satisfactory. The classifier trained on OSLN2 and applied to OSLN3 does not perform well at all. The model trained on both stations and applied separately to both stations performs better at OSLN3, but still clearly worse than our preferred model. The model trained on both stations and applied to OSLN2 performs slightly worse than the original model we used above. With longer time series of blasts being available in future, we would like to generalize the blast classifier for more stations, and most important, for other source areas.

The SNRs of many detected blasts in our study are rather low (see Fig. 9a and b) which is expected for an urban environment. This naturally impacts the performance of the outlier detector, as discussed above, as well as to some extent the blast classifier. Again, the best

way to deal with this issue is a larger training data set to better represent noisy waveform data.

For evaluating our methodology we used a manually compiled data set of locatable blasts in the city of Oslo. We have shown that the outlier detector detects about 70% of those events. The classifier specifically trained to detect blasts increases recognition rate to 80%. If the goal of monitoring is to detect as many real blasts as possible while accepting a number of false alarms, the classifier performance can be improved further up to 87% by using a lower probability threshold. In general, our outlier event detector and blast classifier generate a very low number of false alarms. We encountered randomly associated Rg waves producing false events in less than 1% of the locatable outlier events. However, this is partially due to the combination with the Rg-based locations procedure which sorts out many unlocatable events. Nevertheless, even without event location, the amount of events to be processed is reduced considerably compared to simply applying an STA/LTA detector and attempting to locate all those events. Furthermore, we found also many real blast signals among the outliers being unlocatable due to limited station density. With larger training data and denser seismic networks, we therefore expect the benefit of our methodology to become even more evident.

A system such the one we have proposed has to be adapted and modified when deployed in another city. The most critical parts are the existence of a station network with sufficient resolution for event location and the deployment of a functional outlier detector. If all signals of interest have a high SNR, it may be sufficient to simply use an STA/LTA detector combined with Rg-wave based location for outlier detection. However, the supervised blast classifier would still be part of such a system. Before deploying the network and selecting stations for the outlier detector, potential source locations for blasts, e.g., construction sites, or infrastructure to be monitored should be identified. From our experience, the Rg waves of blasts needed for localization can be observed at up to about 12 km distance, whereas blast signals detected as outliers require stations at not more than about 6 km distance. Furthermore, adaption to other stations in another city requires retraining of the DL models. The outlier detector does not require a large data set for this. After 2-3 weeks there should be sufficient wave field variability captured to train the auto-encoder(s). However, the blast classifier would require rather frequent blasting to gather enough events for training the CNN; from our experience around 1,000 events are needed. This might be a limitation of our work flow in areas with infrequent blasting.

The final question is how our system can be integrated into a smart-city solution. The simplest objective could be to provide the general public with real-time information about event locations on a public web-dashboard or through a mobile app. If citizens felt ground shaking, they can easily check if it was related to any known construction site. If the goal is early warning in case of unusual events (accidents, attacks) and the seismic monitoring system is supposed to automatically alert the city authorities and possibly the pub-

lic, combination with other data sources available to the stakeholders may be needed to avoid false alarms. All data layers combined could then provide automatic alerts and initiate further actions. For example, we envision that the recorded ground motion could be utilized to predict potential damage of infrastructure and buildings categorized on the European Macroseismic Scale (EMS-98). If most detected events are construction blasts as in our case, automatic monitoring can still be useful in a smart-city application, for example to alert about blasts above a certain magnitude or amplitude threshold in areas with unstable ground such as quick clay.

7 Conclusions

The objective of this study was to detect events in the city of Oslo, Norway, that generate seismic signals. To this end, we have successfully developed a prototype of an automatic urban seismic monitoring system using input data from low-cost seismic sensors deployed between 2021 and 2023. The work flow of our system includes two deep learning methods: the first one identifies rare events using an event outlier detector based on an auto-encoder model and the second one classifies events of interest using a CNN model trained in a supervised manner. Both methods used waveforms of signals as input, which were pre-detected using the traditional STA/LTA trigger method. For both evaluating the outlier detector and training the event classifier with events of interest, we relied on locating the seismic signals using Rg waves observed on the seismic network.

The results of our approach impressively reveal ongoing construction activity and their temporal variation in the city of Oslo. From about 1,870 construction blasts in different areas during 22 months of monitoring generating locatable seismic signals on our network, the outlier detector recognized 69%. The classifier trained on these blasts was able to detect between 80 and 87% of those, many with low Signal-to-Noise ratios. At the same time, the false detection rate is very low. In absolute numbers, the automatic system was able to retrieve 1,271 of the manually identified blasts in the initial outlier detection step, and between 1,385 and 1,627 blasts, depending on the detection threshold, using the blast classifier.

The performance of our prototype system could be improved by expanding and densifying the seismic network as well as increasing the training data with more blast records. However, we demonstrated that even with a low number of seismic sensors, a city can be monitored automatically and continuously for explosion events. This opens up new possibilities to include seismic records into the sensor data stream of future smart city solutions. We are therefore confident that the outcome of our pilot study represents a robust prototype system for urban explosion monitoring in the city of Oslo and possibly elsewhere.

Acknowledgements

Seismic data processing was done using Obspy (Beyreuther et al., 2010) and DL models were imple-

mented in Keras (Chollet et al., 2015). This study was funded by The Research Council of Norway (project number 311596, GEObyIT). We thank all seismic station hosts for participating in this study: Oslo municipality, Møller Medvind AS, Tore Ligård AS, Gravco AS, Alfaset graveyard, Bjerke VGS, Linderud skole, Kongshavn VGS, Ullevål skole, Vinderen skole, Hovseter skole, and the Gleiss family. Special thanks go to the NORSAR technicians who helped deploying and maintaining the network, in particular Sindre Stokkan and Bjørn Christian Freitag. We thank Karoline Aastrup-Köhler for spellchecking the manuscript. We thank René Steinmann and one anonymous reviewer for their very helpful comments and suggestions.

Data and code availability

All data of the temporary network and station OSL are openly available at the Norwegian EIDA node (Ottemöller et al., 2021). The temporary stations have the network code 4X (Köhler, 2021). Information about blasting activity in Oslo was retrieved from <https://nabovarsling.no>. Maintained code is available at <https://github.com/NorwegianSeismicArray/urban-seismic-monitoring>. The version of the code at the time of publication of this study with scripts accessing the data on the Norwegian EIDA node is available at <https://doi.org/10.5281/zenodo.10777734> (Köhler and Myklebust, 2024).

Competing interests

The authors have no competing interests.

References

- Al Nuaimi, E., Al Neyadi, H., Mohamed, N., and Al-Jaroodi, J. Applications of big data to smart cities. *Journal of Internet Services and Applications*, 6(1):1–15, 2015. doi: 10.1186/s13174-015-0041-5.
- Bergen, K. J., Chen, T., and Li, Z. Preface to the Focus Section on Machine Learning in Seismology. *Seismological Research Letters*, 90(2A):477–480, 2019. doi: 10.1785/0220190018.
- Bergen University, E. S. Operation of the Norwegian National Seismic Network 2011. Technical report, Department of Earth Science University of Bergen, 2012. https://nnsn.geo.uib.no/reports/2011/All_reports_2011.pdf.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J. ObsPy: A Python toolbox for seismology. *Seismological Research Letters*, 81(3):530–533, 2010. doi: 10.1785/gssrl.81.3.530.
- Bouchard, S., L’Heureux, J.-S., Johansson, J., Leroueil, S., and LeBoeuf, D. Blasting induced landslides in sensitive clays. In *Landslides and Engineered Slopes. Experience, Theory and Practice*, pages 497–504. CRC Press, 2018.
- Brockhoff, M. Urban growth in developing countries: a review of projections and predictions. *Population and Development Review*, 25(4):757–778, 1999. doi: 10.1111/j.1728-4457.1999.00757.x.
- Chamarczuk, M., Nishitsuji, Y., Malinowski, M., and Draganov, D. Unsupervised learning used in automatic detection and

- classification of ambient-noise recordings from a large-N array. *Seismological Research Letters*, 91(1):370–389, 2020. doi: 10.1785/0220190063.
- Chang, S. E., McDaniels, T., Fox, J., Dhariwal, R., and Longstaff, H. Toward disaster-resilient cities: Characterizing resilience of infrastructure systems with expert judgments. *Risk analysis*, 34(3):416–434, 2014. doi: 10.1111/risa.12133.
- Chollet, F. et al. Keras. <https://github.com/fchollet/keras>.
- Dando, B. D., Goertz-Allmann, B. P., Brissaud, Q., Köhler, A., Schweitzer, J., Kværna, T., and Liashchuk, A. Identifying attacks in the Russia–Ukraine conflict using seismic array data. *Nature*, pages 1–6, 2023. doi: 10.1038/s41586-023-06416-7.
- Díaz, J., Ruiz, M., Sánchez-Pastor, P. S., and Romero, P. Urban seismology: On the origin of earth vibrations within a city. *Scientific reports*, 7(1):15296, 2017. doi: 10.1038/s41598-017-15499-y.
- Dowding, C. H. *Blast Vibration Monitoring for Engineering*, volume 4. Elsevier, 2016.
- Fiori, R., Vergne, J., Schmittbuhl, J., and Zigone, D. Monitoring induced microseismicity in an urban context using very small seismic arrays: The case study of the Vendenheim EGS project. *Geophysics*, 88(5):WB71–WB87, 2023. doi: 10.1190/geo2022-0620.1.
- Fischer, J., Redlich, J.-P., Scheuermann, B., Schiller, J., Günes, M., Nagel, K., Wagner, P., Scheidgen, M., Zubow, A., Eveslage, I., et al. From earthquake detection to traffic surveillance—about information and communication infrastructures for smart cities. In *System Analysis and Modeling: Theory and Practice: 7th International Workshop, SAM 2012, Innsbruck, Austria, October 1-2, 2012. Revised Selected Papers 7*, pages 121–141. Springer, 2013. doi: 10.1007/978-3-642-36757-1_8.
- Gharti, H. N., Oye, V., Roth, M., and Kühn, D. Automated microearthquake location using envelope stacking and robust global optimization. *Geophysics*, 75(4):MA27–MA46, 2010. doi: 10.1190/1.3432784.
- Gibbons, S. J. and Ringdal, F. The detection of low magnitude seismic events using array-based waveform correlation. *Geophysical Journal International*, 165(1):149–166, 2006. doi: 10.1111/j.1365-246X.2006.02865.x.
- Hillers, G., T. Vuorinen, T. A., Uski, M. R., Kortström, J. T., Mäntyniemi, P. B., Tiira, T., Malin, P. E., and Saarno, T. The 2018 geothermal reservoir stimulation in Espoo/Helsinki, southern Finland: Seismic network anatomy and data features. *Seismological Research Letters*, 91(2A):770–786, 2020. doi: 10.1785/0220190253.
- Jenkins, W. F., Gerstoft, P., Bianco, M. J., and Bromirski, P. D. Unsupervised deep clustering of seismic data: Monitoring the Ross Ice Shelf, Antarctica. *Journal of Geophysical Research: Solid Earth*, 126(9):e2021JB021716, 2021. doi: 10.1029/2021JB021716.
- Johnson, C. W., Ben-Zion, Y., Meng, H., and Vernon, F. Identifying different classes of seismic noise signals using unsupervised learning. *Geophysical Research Letters*, 47(15):e2020GL088353, 2020. doi: 10.1029/2020GL088353.
- Kalinowski, M. B. and Mialle, P. Introduction to the topical issue on nuclear explosion monitoring and verification: scientific and technological advances. *Pure and Applied Geophysics*, 178(7):2397–2401, 2021. doi: 10.1007/s00024-021-02783-2.
- Köhler, A., Ohrnberger, M., and Scherbaum, F. Unsupervised pattern recognition in continuous seismic wavefield records using self-organizing maps. *Geophysical Journal International*, 182(3):1619–1630, 2010. doi: 10.1111/j.1365-246X.2010.04709.x.
- Köhler, A., Myklebust, E., and Mæland, S. Enhancing seismic calving event identification in Svalbard through empirical matched field processing and machine learning. *Geophysical Journal International*, 230(2):1305–1317, 2022. doi: 10.1093/gji/ggac117.
- Kong, Q., Allen, R. M., Schreier, L., and Kwon, Y.-W. MyShake: A smartphone seismic network for earthquake early warning and beyond. *Science advances*, 2(2):e1501055, 2016. doi: 10.1126/sciadv.1501055.
- Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., and Gerstoft, P. Machine learning in seismology: Turning data into insights. *Seismological Research Letters*, 90(1):3–14, 2019. doi: 10.1785/0220180259.
- Kraft, T., Mai, P. M., Wiemer, S., Deichmann, N., Ripperger, J., Kästli, P., Bachmann, C., Fäh, D., Wössner, J., and Giardini, D. Enhanced geothermal systems: Mitigating risk in urban areas. *Eos, Transactions American Geophysical Union*, 90(32):273–274, 2009. doi: 10.1029/2009EO320001.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. In Pereira, F., Burges, C., Bottou, L., and Weinberger, K., editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.
- Köhler, A. GEObyIT seismic network in Oslo, Norway, 2021. doi: 10.7914/3dms-sj84.
- Köhler, A. and Myklebust, E. B. Code for monitoring urban construction and quarry blasts with low-cost seismic sensors and machine learning tools in the city of Oslo, Norway, 2024. doi: 10.5281/zenodo.10777734.
- Leonard, M. and Kennett, B. Multi-component autoregressive techniques for the analysis of seismograms. *Physics of the Earth and Planetary Interiors*, 113(1-4):247–263, 1999. doi: [https://doi.org/10.1016/S0031-9201\(99\)00054-0](https://doi.org/10.1016/S0031-9201(99)00054-0).
- McKinsey. *Smart cities: Digital solution for a more liveable future*. NY: McKinsey Global Institute, 2018.
- Mousavi, S. M. and Beroza, G. C. Machine Learning in Earthquake Seismology. *Annual Review of Earth and Planetary Sciences*, 51, 2023. doi: 10.1146/annurev-earth-071822-100323.
- Mousavi, S. M., Zhu, W., Ellsworth, W., and Beroza, G. Unsupervised clustering of seismic signals using deep convolutional autoencoders. *IEEE Geoscience and Remote Sensing Letters*, 16(11):1693–1697, 2019. doi: 10.1109/LGRS.2019.2909218.
- Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., and Beroza, G. C. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nature communications*, 11(1):1–12, 2020. doi: 10.1038/s41467-020-17591-w.
- Navarro, J., Schiavon, A., Vieira, M., and Silva, P. Deep seismic compression. In *81st EAGE Conference and Exhibition 2019*, pages 1–5. European Association of Geoscientists & Engineers, 2019. doi: 10.3997/2214-4609.201901620.
- Naveen, G., Sastry, V., and Chandar, K. R. Assessment of Structural Damage Due to Blasting in Hydro Power Tunnel. In *International Conference on Geotechnical Challenges in Mining, Tunneling and Underground Infrastructures*, pages 229–240. Springer, 2021.
- Nugent, J. Raspberry Shake: Watch the Earth move under your feet. *Science Scope*, 42(4):22–25, 2018. <https://www.jstor.org/stable/26611884>.
- Nuha, H. H., Balghonaim, A., Liu, B., Mohandes, M., Deriche, M., and Fekri, F. Deep neural networks with extreme learning machine for seismic data compression. *Arabian Journal for Science and Engineering*, 45(3):1367–1377, 2020. doi: 10.1007/s13369-019-03942-3.
- OpenStreetMap contributors. Planet dump retrieved from <https://planet.osm.org>, 2017. <https://>

//www.openstreetmap.org.

- Ottmøller, L., Strømme, M. L., and Storheim, B. M. Seismic monitoring and data processing at the Norwegian National Seismic Network. *Summary of the Bulletin of the International Seismological Centre*, 52(1):27–40, 2018.
- Ottmøller, L., Michalek, J., Christensen, J.-M., Baadshaug, U., Halpaap, F., Natvik, Ø., Kværna, T., and Oye, V. UiB-NORSAR EIDA node: Integration of seismological data in Norway. *Seismological Society of America*, 92(3):1491–1500, 2021. doi: 10.1785/0220200369.
- Provost, F., Hibert, C., and Malet, J.-P. Automatic classification of endogenous landslide seismicity using the Random Forest supervised classifier. *Geophysical Research Letters*, 44(1):113–120, 2017. doi: 10.1002/2016GL070709.
- Ritter, J. R., Balan, S. F., Bonjer, K.-P., Diehl, T., Forbriger, T., Mărmureanu, G., Wenzel, F., and Wirth, W. Broadband urban seismology in the Bucharest metropolitan area. *Seismological Research Letters*, 76(5):574–580, 2005. doi: 10.1785/gssrl.76.5.574.
- Seydoux, L., Balestrieri, R., Poli, P., Hoop, M. d., Campillo, M., and Baraniuk, R. Clustering earthquake signals and background noises in continuous seismic data with unsupervised deep learning. *Nature communications*, 11(1):3972, 2020. doi: 10.1038/s41467-020-17841-x.
- Shallan, O., Eraky, A., Sakr, T., and Emad, S. Response of building structures to blast effects. *International journal of engineering and innovative technology*, 4(2):167–175, 2014. https://www.humanitarianlibrary.org/sites/default/files/2018/10/IJEIT1412201408_30.pdf.
- Sick, B., Guggenmos, M., and Joswig, M. Chances and limits of single-station seismic event clustering by unsupervised pattern recognition. *Geophysical Journal International*, 201(3):1801–1813, 2015. doi: 10.1093/gji/ggv126.
- Spica, Z. J., Perton, M., Martin, E. R., Beroza, G. C., and Biondi, B. Urban seismic site characterization by fiber-optic seismology. *Journal of Geophysical Research: Solid Earth*, 125(3):e2019JB018656, 2020. doi: 10.1029/2019JB018656.
- Steinmann, R., Seydoux, L., Beaucé, E., and Campillo, M. Hierarchical exploration of continuous seismograms with unsupervised learning. *Journal of Geophysical Research: Solid Earth*, 127(1):e2021JB022455, 2022a. doi: 10.1029/2021JB022455.
- Steinmann, R., Seydoux, L., and Campillo, M. AI-based unmixing of medium and source signatures from seismograms: ground freezing patterns. *Geophysical Research Letters*, 49(15):e2022GL098854, 2022b. doi: 10.1029/2022GL098854.
- Thill, M., Konen, W., Wang, H., and Bäck, T. Temporal convolutional autoencoder for unsupervised anomaly detection in time series. *Applied Soft Computing*, 112:107751, 2021. doi: 10.1016/j.asoc.2021.107751.
- Valentine, A. P. and Trampert, J. Data space reduction, quality assessment and searching of seismograms: autoencoder networks for waveform data. *Geophysical Journal International*, 189(2):1183–1202, 2012. doi: 10.1111/j.1365-246X.2012.05429.x.
- Wang, Y., Yao, H., and Zhao, S. Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242, 2016. doi: 10.1016/j.neucom.2015.08.104.
- Yin, C., Zhang, S., Wang, J., and Xiong, N. N. Anomaly detection based on convolutional recurrent autoencoder for IoT time series. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 52(1):112–122, 2020. doi: 10.6038/pg2020DD0013.
- Yoon, C. E., O'Reilly, O., Bergen, K. J., and Beroza, G. C. Earthquake detection through computationally efficient similarity search. *Science advances*, 1(11):e1501057, 2015. doi: 10.1126/sciadv.1501057.
- Zhang, K., Ni, J., Yang, K., Liang, X., Ren, J., and Shen, X. S. Security and privacy in smart city applications: Challenges and solutions. *IEEE Communications Magazine*, 55(1):122–129, 2017. doi: 10.1109/MCOM.2017.1600267CM.
- Zheng, H. and Zhang, B. Intelligent seismic data interpolation via convolutional neural network. *Progress in Geophysics*, 35(2):721–727, 2020. doi: <https://doi.org/10.6038/pg2020DD0013>.
- Zhu, W. and Beroza, G. C. PhaseNet: a deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International*, 216(1):261–273, 2018. doi: 10.1093/gji/ggy423.

The article *Monitoring urban construction and quarry blasts with low-cost seismic sensors and deep learning tools in the city of Oslo, Norway* © 2024 by Andreas Köhler is licensed under CC BY 4.0.