



Deep learning for laboratory earthquake prediction and autoregressive forecasting of fault zone stress



Laura Laurenti ^{a,*}, Elisa Tinti ^{a,b}, Fabio Galasso ^a, Luca Franco ^a, Chris Marone ^{a,c}

^a Università La Sapienza, Rome, Italy

^b Istituto Nazionale Geofisica e Vulcanologia, Rome, Italy

^c The Pennsylvania State University, University Park, PA, USA

ARTICLE INFO

Article history:

Received 21 March 2022

Received in revised form 28 August 2022

Accepted 17 September 2022

Available online 6 October 2022

Editor: R. Bendick

Dataset link: <https://github.com/lauralaurenti/DNN-earthquake-prediction-forecasting>

Keywords:

deep learning

laboratory earthquakes

auto-regressive forecasting

LSTM

TCN

Transformer

ABSTRACT

Earthquake forecasting and prediction have long and in some cases sordid histories but recent work has rekindled interest based on advances in early warning, hazard assessment for induced seismicity and successful prediction of laboratory earthquakes. In the lab, frictional stick-slip events provide an analog for earthquakes and the seismic cycle. Labquakes are also ideal targets for machine learning (ML) because they can be produced in long sequences under controlled conditions. Indeed, recent works show that ML can predict several aspects of labquakes using fault zone acoustic emissions (AE). Here, we extend these works with: 1) deep learning (DL) methods for labquake prediction, 2) by introducing an autoregressive (AR) forecasting DL method to predict fault zone shear stress, and 3) by expanding the range of lab fault zones studied. The AR methods allow forecasting stress at future times via iterative predictions using previous measurements. Our DL methods outperform existing ML models and can predict based on limited training. We also explore forecasts beyond a single seismic cycle for aperiodic failure. We describe significant improvements to existing methods of labquake prediction and demonstrate: 1) that DL models based on Long-Short Term Memory and Convolution Neural Networks predict labquakes under conditions including pre-seismic creep, aperiodic events and alternating slow/fast events and 2) that fault zone stress can be predicted with fidelity, confirming that acoustic energy is a fingerprint of fault zone stress. Our DL methods predict time to start of failure (TTsF) and time to the end of Failure (TTeF) for labquakes. Interestingly, TTeF is successfully predicted in all seismic cycles, while the TTsF prediction varies with the amount of pre-seismic fault creep. We report AR methods to forecast the evolution of fault stress using three sequence modelling frameworks: LSTM, Temporal Convolution Network and Transformer Network. AR forecasting is distinct from existing predictive models, which predict only a target variable at a specific time. The results for forecasting beyond a single seismic cycle are limited but encouraging. Our ML/DL models outperform the state-of-the-art and our autoregressive model represents a novel framework that could enhance current methods of earthquake forecasting.

© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Earthquake forecasting and prediction have long been of interest because of the obvious practical and societal implications. While research has waxed and waned with many failed directions, recent work on early warning systems, hazard assessment, and precursors has provided renewed interest (Allen and Stogaitis, 2022; Ben-Zion and Lyakhovskiy, 2001; Beroza et al., 2021; Wang et al., 2021; Denolle et al., 2014; Kohler et al., 2020; Kong et al., 2018; Pritchard et al., 2020). Laboratory work has fuelled this interest via: 1) the discovery that machine-learning can predict several as-

pects of lab earthquakes (Rouet-Leduc et al., 2017, 2018; Johnson et al., 2021) and 2) recent work on the mechanisms of precursors to labquakes that adds to earlier studies (Shreedharan et al., 2021; Bolton et al., 2021; Dieterich, 1978; Scholz, 1968; Dresen et al., 2020; Scuderi et al., 2016; Hedayat et al., 2014; Acosta et al., 2019; Johnson et al., 2013; Main et al., 1989, 1992; Thompson et al., 2005; Passelègue et al., 2017).

Recent works show that acoustic emissions can be used to predict labquake failure time, fault zone shear stress, and labquake stress drop (Rouet-Leduc et al., 2017, 2018; Lubbers et al., 2018; Hulbert et al., 2018; Bolton et al., 2020). Existing works also show that seismic radiation from lab faults scales with time to failure and that its recent history can be used to predict the current fault zone stress state (Bolton et al., 2019; Corbi et al., 2019; Jasper-

* Corresponding author.

E-mail address: laura.laurenti@uniroma1.it (L. Laurenti).

son et al., 2021) with both active and passive seismic monitoring of the fault zone (Shreedharan et al., 2021; Shokouhi et al., 2021). These studies show that both passive signals coming from the fault zone and active acoustic signals passing through the fault can be used to predict failure. The active source signals record changes in fault properties prior to failure and thus offer the possibility of incorporating physics-based models (i.e., of asperity contact mechanics) in the ML/DL algorithm. Other studies show that ML can be used to connect lab AE events with fault zone microstructure (Chaipornkaew et al., 2022; Trugman et al., 2020; Ma et al., 2021) and that ML methods can be augmented by directly incorporating physics into the prediction models (Raissi et al., 2019). Here, we extend this approach to investigate the use of DL methods and also to introduce a new approach based on autoregressive forecasting. The AR methods are distinct because they predict future time horizons using present and past values rather than the current time based on previous values.

Recent lab studies have also identified reliable precursors to failure in connection with labquake prediction. These include strong relations between acoustic transmissivity, elastic properties and fault strength (Schubnel et al., 2006; Nagata et al., 2008; Scuderi et al., 2016; Tinti et al., 2016; Shreedharan et al., 2020, 2021; Bolton et al., 2021; Rivière et al., 2018). Other previous studies provided insight on labquake nucleation processes and the evolution of frequency-magnitude (b-value) statistics during the seismic-cycle prior to failure (Dresen et al., 2020; Rivière et al., 2018; Scholz, 2015; Goebel et al., 2017; Kwiatak et al., 2014; Latour et al., 2013; Shreedharan et al., 2020; Johnson et al., 2021; McBeck et al., 2020; McLaskey and Lockner, 2014). These studies provide a framework for understanding rupture nucleation and therefore how machine learning can predict labquakes using statistics of the continuous AE signal emanating from lab faults (Hulbert et al., 2018). Previous works have used ML to predict geodetic signals of episodic tremor and slip and other data from tectonic faults (Hulbert et al., 2020; Johnson et al., 2020; Wang et al., 2021; McBrearty et al., 2019).

The majority of existing ML studies of labquakes use decision-tree-based algorithms and a gradient boosting framework (e.g., XGBoost model). In such studies models are built and errors are minimized by gradient descent methods using multiple learning algorithms that iteratively improve predictive performance beyond that of a single classical learning algorithm. These studies show that it is possible to successfully predict labquakes with reasonable accuracy. The success of the original approaches was based on continuous records of AE events broadcast from the lab fault zones. These AE represent a form of lab microearthquakes. AE are also detected during stable frictional slip. In this case, their origin is less clear, as they could represent microfracture events or micro-instabilities that do not impact the macroscopic strength (and therefore they do not appear as a stress drop).

Despite the dramatic recent advances in lab earthquake prediction our understanding of how and why these methods work is far from complete. Simple questions remain such as how the AE signal scales with labquake timing and magnitude and how AE signal characteristics encode fault zone shear stress even during aperiodic failure. Here, we address these questions by exploring a wider range of ML/DL methods. We also address the open discussion regarding the predictability of earthquakes and if faults slip deterministically or stochastically.

We explore two different problems and three families of Neural Network (NN) architectures. The two problems refer to two different tasks: prediction and forecasting. The former has been addressed in previous works (Hulbert et al., 2018; Rouet-Leduc et al., 2017), and the main goal is how to predict the present value of target variables (e.g., shear stress) given past information and some memory of sequence evolution. With the second problem we

introduce a novel autoregressive (AR) forecasting approach to predict not only the present value of the target variable, but also its step-by-step future evolution.

The three families of NN architectures adopted in this paper are Long short-term memory (LSTM), Temporal Convolutional Network (TCN), and Transformer Network (TF). LSTM is a Recurrent Neural Network (RNN) designed to capture long-range data dependencies in time series and sequential data (Hochreiter and Schmidhuber, 1997). RNNs are Deep Neural Networks (DNNs) that recursively update an internal status (this may represent the current status of the fault) from which predictions are made. This approach provides the capability of modelling sequences of events. LSTM is designed to solve gradient vanishing problems of RNN's for long-term prediction. LSTM has shown great potential for modelling seismic time series (Johnson et al., 2021). TCN, one of the simplest NN architectures for sequence modelling, is a Convolutional Neural Network (CNN) with adjustments that allow longer sequences by addressing the relation between the depth and the length of the considered input sequence. Because it involves convolution, TCN may be applied to sequences of variable length. The TF architecture is the most recent and it has shown promising results for earthquake detection (Mousavi et al., 2020). It's based on mechanisms of self-attention and is effective in capturing long-term dependencies in a variety of sequence modelling tasks (Vaswani et al., 2017). One notable weakness of TF is its complexity and that it requires significantly larger datasets to be trained than other architectures.

2. Laboratory earthquake experiments

We use data from experiments conducted in the double-direct shear (DDS) configuration with a biaxial deformation machine (see Fig. 1). The DDS experiments consist of two identical fault zones sheared simultaneously between three loading blocks. Two stationary side blocks are supported from below and a central block is driven between them to impose shear.

Our experiments were performed following a standard procedure that was developed to obtain highly reproducible results (Karner and Marone, 1998; Bolton et al., 2020). First, fault normal stress was imposed with a servo-controlled hydraulic system and maintained at a controlled value throughout shear. Then, fault zone shear was imposed by driving the central block of the DDS configuration at a constant displacement rate, thus imposing a constant shear strain rate within the fault zone. Fault zones were composed of granular materials that began as 3-mm thick layers of nominal contact area $10 \times 10 \text{ cm}^2$. We use data from two types of lab fault zones: 1) glass beads (experiment p4581) and 2) quartz powder, (experiments p4679 and p5198). The glass beads range in size from 100-150 μm and for these experiments the normal stress is held below 10 MPa so that grain fracturing is minimal (e.g., Mair et al., 2002). The quartz powder has a mean particle size of 10 μm (Bolton et al., 2020) and a power-law particle size distribution that does not change significantly with shear for the stresses of our experiments. The fault normal stress was 7 MPa for experiment p4679 and it ranged from 6 to 11 MPa in experiment p5198 and from 2 to 8 MPa in p4581. These experiments exhibit a range of stick-slip behaviours from fast to slow (Leeman et al., 2016) and have seismic cycles that range from highly periodic to complex and aperiodic.

We study a range of frictional sliding behaviours from stable sliding to highly unstable slip. In the framework of rate-and-state friction this range of behaviours is understood in terms of the ratio of loading stiffness to a critical rate of fault weakening with slip, which scales with normal stress (Leeman et al., 2016). Each experiment includes hundreds of lab earthquakes (Fig. S1). By adjusting the loading stiffness and/or the fault normal stress, the same fault gouge can host slow-slip events, fast lab earthquakes,

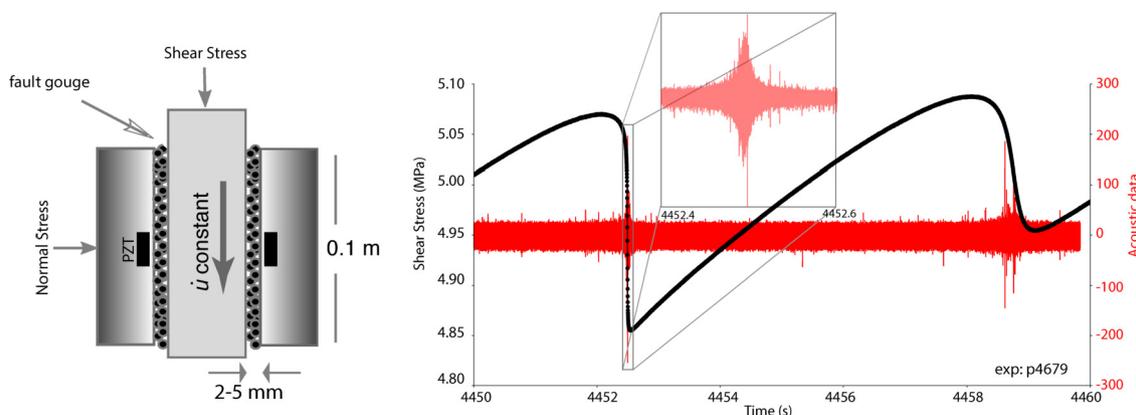


Fig. 1. Left: schematic of the double direct shear configuration used for the experiments. Granular layers simulate an earthquake fault and are sheared between rough surfaces. Acoustic signals are recorded with Piezo-Electric Transducers (PZTs) embedded in the loading blocks. Faults are loaded with a normal stress that is maintained constant via servo control. The central block is driven downward at constant displacement rate to produce frictional shear. Acoustic emission (AE) data are recorded continuously at a sampling rate of 4 MHz. Right: typical data for shear stress and AE coming from the fault zone. Zoomed window above shows data from a lab earthquake. Note that the acoustic signal differs for the large labquake compared to the smaller event that follows (near 4459 s).

Table 1

Results for DL prediction of using the continuous lab seismic signal. For each target we show goodness of fit (GOF) in terms of the coefficient of determination R^2 and the root mean square error RMSE. Shear stress is reasonably well predicted for each experiment. Of the three targets, TTsF is the hardest to predict. Experiment p4581 is the hardest one to predict.

Target	GOF	p4581 Glass beads	p5198 Quartz powder	p4679 Quartz powder
Shear Stress	R^2	0.9254	0.9884	0.9574
	RMSE	0.0670	0.0305	0.0519
Time To Start Of Failures	R^2	0.6317	0.9313	0.8229
	RMSE	0.15844	0.0728	0.1087
Time To End Of Failures	R^2	0.8721	0.9697	0.9200
	RMSE	0.0972	0.0476	0.0723

or complex, quasi-periodically events (Leeman et al., 2016; Scuderi et al., 2017; Mele Veedu et al., 2020). Mechanical data of stress and strain are recorded at 1 kHz and acoustic signals are recorded at 4 MHz using lead-zircon-titanate piezoceramic sensors embedded in the DDS (Fig. 1) (Rivière et al., 2018; Bolton et al., 2020). A high-precision signal is used to synchronize the mechanical and acoustic data.

In the lab three stages of the seismic cycle can be identified: an initial –interseismic– stage of stress increase, a stage with pre-seismic slip and then a last stage with a co-seismic stress drop (Figs. 1 and 2). The interseismic period of the lab seismic cycle is identified from the initial, linear-elastic portion of the curve, where load increases in proportion to stiffness. Pre-seismic slip is marked by a deviation from elastic loading. This typically corresponds to a relatively short time interval preceding failure and seismic energy release. Pre-seismic slip is part of the labquake nucleation process. We study data from three experiments chosen to represent a range of lab earthquake behaviours. In some cases the events are quasi-periodic, as for experiment p5198 (Fig. 2 a), and in others, as for experiment p4679 (Fig. 2 d) they are aperiodic with alternating slower and faster events. Experiments like p4679 often have complex stress evolution throughout the lab seismic cycle. Some of this complexity can be seen at the scale of the whole experiment (Fig. S1) and other aspects of it are visible only at a larger scale (Fig. 2). In the case of p4581, the pre-seismic phase of the seismic cycle is characterized by almost no pre-slip prior to failure. The acoustic emission record for this type of experiment is distinct because the signal variance is quite low until just prior to the labquake (Fig. S2). Nonetheless there is still AE activity before the

main labquake. In contrast, experiment p5198 shows significant pre-seismic slip before the peak stress and failure (see Fig. 2). Both of these experiments show somewhat regular, quasi-periodic seismic cycles (Figs. 2 and S2). The third experiment (p4679) shows aperiodic seismic cycles that are much more challenging to predict. For each experiment we focus on a data segment of 30 to 50 events shown in the red boxes in Fig. S1.

3. Prediction and forecasting models

We use DNNs by adopting three of the most well-known sequence modelling approaches: LSTM (Long short-term memory) (Hochreiter and Schmidhuber, 1997), TCN (Temporal Convolutional Network) (Bai et al., 2018) and TF (Transformer Network) (Vaswani et al., 2017). Figs. 3 and 4 and Appendix 9.1 provide a summary of our models.

3.1. Input, output and model performance

We use measurements of fault zone shear stress and radiated elastic energy from AE as ML features. Model input consists of statistical measures of the continuous seismic data generated from AE (for the prediction task, see Section 3.2) or the lab measurements of shear stress (for forecasting, see Section 3.3). The model outputs (different for prediction and forecasting) are compared to the input measurements of ground truth. Because we are dealing with time series of physical processes it is inappropriate to shuffle data in time. This limitation holds for both our prediction and forecasting tasks.

As in Rouet-Leduc et al. (2017), we consider the variance of the continuous acoustic signal as the most important ML feature for the prediction task. An example of raw data for the acoustic signal and shear stress is shown in Fig. 1 while the variance is shown in Fig. 2.

For the ML prediction task, our model predicts as output, often called target, the Shear Stress and/or the Time To Failure (TTF), defined as the time remaining before the next labquake, derived from the shear stress time series. We predict both the time of start of failure (TTsF) and end of failure (TTeF). The former derives from the time of maximum shear stress preceding labquakes (Fig. 2, orange line). That is, TTsF is 0 at the maximum shear stress and it increases backward in time within the lab seismic cycle. TTeF is derived from the minimum shear stress after an event (Fig. 2, red line). TTeF is 0 at the stress minimum and it increases backward in time within the lab seismic cycle. We show both TTsF

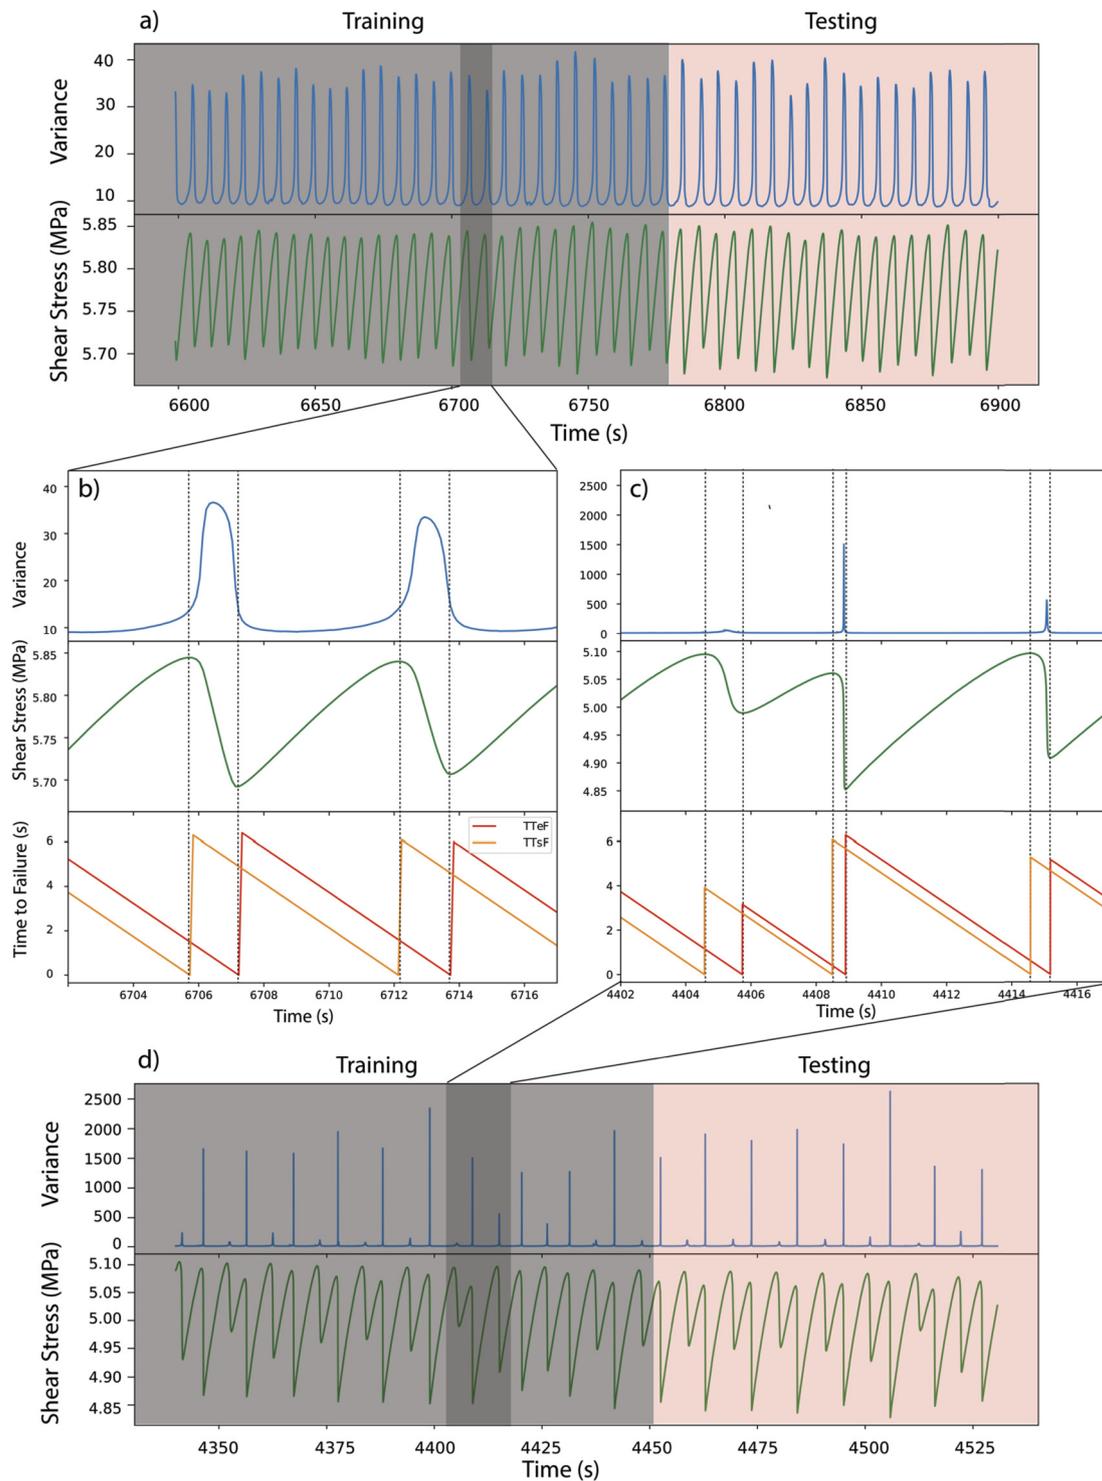


Fig. 2. Data for a series of labquakes showing the training and testing phases. (a) Quasi-periodic events during experiment p5198. Panels (b) and (c) are zooms of (a) and (d), respectively, but also show time to the start of failure TTsF (i.e. time of maximum shear stress preceding labquakes) and time to the end of failure TTeF (i.e. time of minimum shear stress after events). (d) Labquakes from experiment p4679 showing complex behaviour including period doubling and both fast and slow events. For Panels (a) and (d), variance of the continuous seismic data is plotted above shear stress. Note that we just plot the variance of Channel 1's AE; the variance of Channel 2 is similar. Gray and pink shading show examples of training and testing data split (note that we use a different split for prediction vs. forecasting). (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

and TTeF along with shear stress and the variance of the continuous AE record for several lab seismic cycles in Fig. 2. Note here the stepwise nature of TTsF and TTeF as linear functions that our models are designed to predict. Note also the differences in the seismic cycles for our experiments: p5198 shows a somewhat periodic sequence of labquakes with lower acoustic energy compared

to p4679, which has complex seismic cycles and greater acoustic energy release (Fig. 2 and S2).

To calculate AE statistics and determine features for our ML algorithms, we use a moving window on the continuous data which were acquired at 4 MHz (Table S1). We adjust this window for each experiment based on the seismic cycle duration (Fig. 2), so as

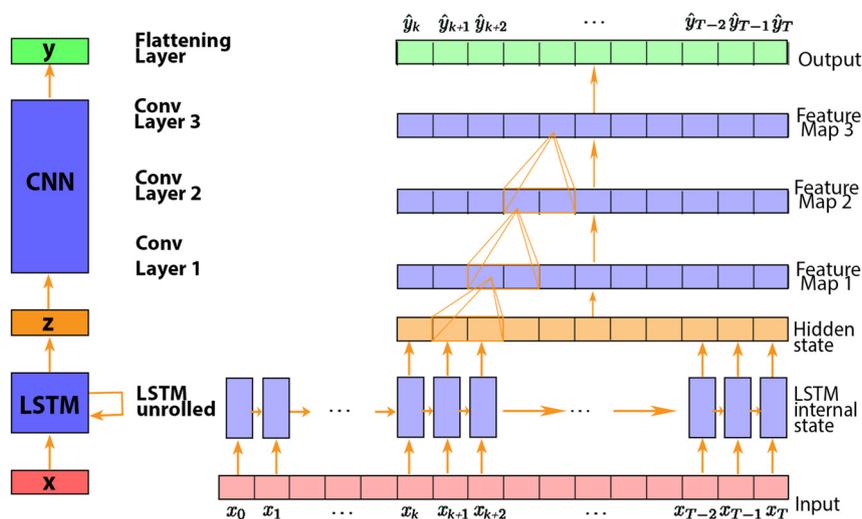


Fig. 3. Schematic model of our combined LSTM and CNN architectures. LSTM scans the input to produce an embedding (lowest level above input). The LSTM layer is followed by three successive CNN layers (see far left) that make the predictions. The input for each layer is the output of the previous layer. The LSTM representation (centre and right side) is unrolled in time, which means that each \hat{z}_t (hidden state z at time t) is predicted by considering information coming from the whole sequence from 0 to t . Thus, the LSTM output starts from k , which represents the LSTM memory length. Further information, as well as discussion about best k value selection, are in Section 4.1.1. The convolutional layers are used to extract features from the input signal (z in this case). The mathematical operation of convolution is performed between the input signal and a convolutional sliding filter of dimension 2. Here, the output of each layer is called a feature map, which gives information about the signal features. Red and pink denote input; green denotes output. Blue denotes hidden states and orange denotes the connections and components inside hidden states. Further details are provided in the Supplementary Section 9.1.

to preserve all of the shear stress evolution while not losing small differences between various cycles. In particular, we adopt window lengths similar to those used previously (e.g., Rouet-Leduc et al., 2018; Hulbert et al., 2018) so that we can make reliable comparisons of model performance.

In previous work they use a so called “subwindows” procedure to address non-single valued functions. This means that one computes the ML features from the continuous seismic signal in two, slightly offset time windows. This is done because the targets (both shear stress and TTF) are symmetric in time within the lab seismic cycle. By using two subwindows the algorithm is able to differentiate between the loading and the slipping part of the seismic cycle. However, our DL models provide a more robust solution to this problem and one that has much more transportability. We do not need subwindows because the LSTM layer carries information about where we are within the seismic cycle. Thus we just use one window for each AE recording.

For our forecasting work we use shear stress measurements, with past data as input and future times as target output. For forecasting, we do not use TTF data but rather forecast shear stress and therefore failure events directly from the shear stress. We smooth the signal in this case too, using a running average window. Then we reduce the resolution to make the length of the sequences suitable for our ML models (details in Table S1). In particular, we sub-sample the signal by a factor of 100 (upon smoothing to limit aliasing).

A typical ML protocol requires at least two subsets of data, one for training, to learn the model parameters, and one for testing, to measure the generalizability of the results to unseen data (Figs. 2 and S2). We also use a validation dataset to evaluate and optimize the model fit during training and to avoid overfitting. Validation is done with a small portion of data (generally 10%) not used for training nor for testing. The split among the three segments preserves the statistical properties of the dataset.

A misfit (loss) function is required to build the ML algorithm and optimize the model parameter weights. We use a root mean squared error (RMSE) loss function to compare output prediction and ground truth data. From this comparison, we iteratively adjust model parameters (weights and biases) to minimize the loss.

Table 2

Results for autoregressive forecasting of fault zone shear stress showing a comparison for each experiment and three models. The goodness of fit (GOF) values are averages computed for all the segments in the testing data. The GOF values vary among the segments as described in the text.

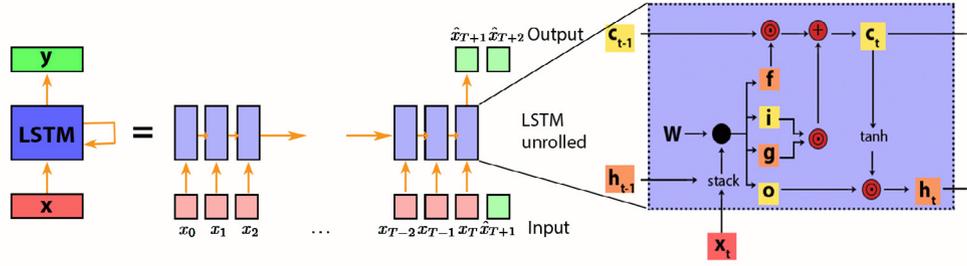
Model	GOF	p4581 Glass beads	p5198 Quartz powder	p4679 Quartz powder
TCN	R^2	0.3935	0.9419	0.8273
	RMSE	0.1245	0.0549	0.0732
LSTM	R^2	-4.5193	0.8021	-0.2704
	RMSE	0.1521	0.0904	0.1634
TF	R^2	0.1172	0.8914	0.7940
	RMSE	0.1460	0.0707	0.0738

Features and target are generally normalized because their ranges typically vary widely and because normalization helps DL algorithms converge with better results. The final performance of our ML/DL models is evaluated with respect to ground truth in the form of our lab data (Table 1 and 2). For evaluation metrics we use both the coefficient of determination R^2 and the RMSE. Summarizing, RMSE is used as loss, following common practise in ML, but also as a test metric representing the discrepancy between model output and ground truth.

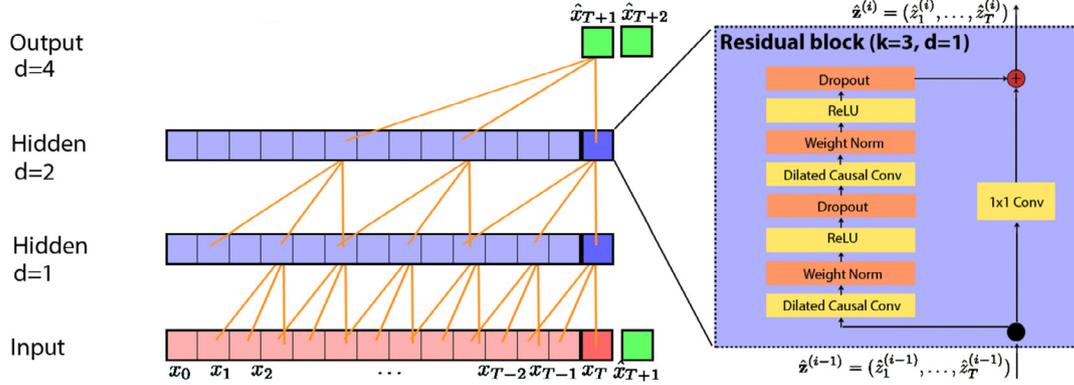
3.2. Prediction

3.2.1. Problem definition

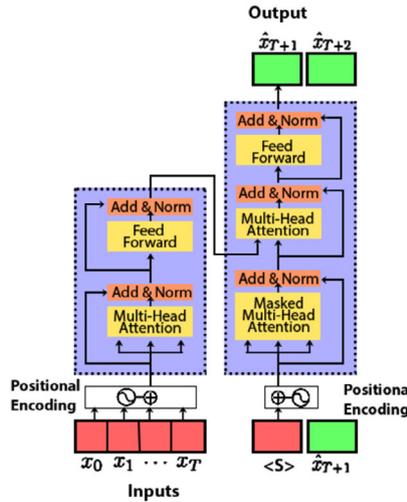
Our goal is to predict the present value of shear stress or TTF, as target variables (e.g., y_t where t is the present time) given current information about AE, as an input variable, (e.g., x_t) and its recent history (e.g. from x_{t-k} to x_{t-1}). We note that the temporal evolution of the AE signal during the lab seismic cycle differs for our range of experiments, in particular during the creep phase preceding labquakes (Figs. 2 and S2). Previous works have observed that AE statistics (in particular signal amplitude variance and higher-order moments) are highly effective at predicting laboratory earthquakes (Rouet-Leduc et al., 2017). Thus we begin by following this approach. However, whereas previous works focused



(a) Time-unrolled LSTM (Hochreiter and Schmidhuber, 1997). Zoom at right shows details of the LSTM cell representation. The state-vector (h_t) is recursively updated at each t , including the new information coming from the current input and also the cell state (c_t).



(b) Architectural elements in a TCN Bai et al. (2018). A dilated causal convolution with dilation factors $d = 1, 2, 4$ and filter size $k = 3$. The receptive field is able to cover all values from the input sequence. Zoom on the right shows the TCN residual block. Here a 1×1 convolution is added when residual input and output have different dimensions.



(c) Transformer architecture Vaswani et al. (2017).

Fig. 4. Schematic showing how TCN, LSTM and TF are used for our forecasting models. In each case, red denotes input, green denotes output, blue denotes hidden states, and orange denotes the connections and components within hidden states. These models are autoregressive (AR). Thus, the inputs are prediction \hat{x}_{T+1} , coming from the previous iteration. Further details are provided in the Supplementary Section 9.1.

on data from only one acoustic sensor, we use data from two sensors, one on either side of the DDS assembly (Fig. 1). This was done by simply using each data stream as a separate input for the continuous seismic signal. Using multiple sensors will clearly be valuable when these techniques are applied to tectonic faults and/or to event location. We did limited testing of the lab data and did not see dramatic improvement over what one would expect by having additional data to constrain the model loss during training.

We consider for the ML process two channels of AE variance x as input, from time $t - k$ to time t , and we predict the tar-

get y (shear stress or TTF) at time t using the input variable and its recent history. We want to learn the ideal function f that maps our input x into the desired output y . With ML we can approximate that function with \hat{f} , leveraging input data, and then approximating the variable of interest \hat{y} . In particular, the estimator $\hat{f}(x_{t-k}, \dots, x_t)$ makes predictions for sample y_t . The quantity k determines the length of the input sequence ($k + 1$ in this case) and represents the LSTM's memory in the internal state. We chose the hyperparameter k using an ad hoc iterative approach (Section 4.1.1).

3.2.2. Deep learning model

We use a DL model based on the implementation of Levinson et al. (<https://www.kaggle.com/c/LANL-Earthquake-Prediction/overview> <https://www.linkedin.com/pulse/my-team-won-20000-1st-place-kaggles-earthquake-corey-levinson/>). This model, from Team “The Zoo,” won the Kaggle competition for lab earthquake prediction (Johnson et al., 2021). We optimize the model for our purposes, including varying hyperparameters and other minor changes, i.e. we adapt the model to allow it to predict one or more targets, as desired. The modified hyperparameters are: 1) using misfit (loss) of ‘mean squared error’ instead of ‘mean absolute error’, 2) in `model.compile()`: we disregard the parameter `loss_weights`, that is set to “none”, because our model doesn’t use a multitarget approach, and 3) in callbacks: `monitor='val_loss'`, `mode='min'` instead of `monitor='val_regressor_mean_absolute_error'`

As shown in Fig. 3, our architecture is a combination of an LSTM layer, which scans the input sequence and produces an embedding, and three stacked convolution (CNN) layers (Fig. 3). This LSTM + CNN architecture extracts the patterns of the sequence and predicts the target. We predict the target for each time t separately, not as a sequence. Therefore, it is sufficient to specify the number of targets and thus we predict N targets for each time step. The model has a total of 421277 parameters. We implement the model in Keras (<https://keras.io/>), an open-source software library that acts as an interface for the TensorFlow machine learning library (<https://www.tensorflow.org/about>). We used Keras to be consistent with the work of Levinson et al.

3.3. Forecasting

3.3.1. Problem definition

Autoregressive forecasting methods are distinct from existing predictive models. They predict not only the current state of the target (e.g. y_t , where t is the present time), but also future steps (e.g. from y_{t+1} to y_{t+N} , where N is the length of the sequence to predict) given past information (e.g. from x_{t-k} to x_{t-1}). In particular, in the forecasting problem the input and output variables need to be of the same nature; that is, we can input shear stress data from time $t-k$ to $t-1$ and predict shear stress in the future from time t to $t+N$. This type of DL model uses an autoregressive technique (AR) in which regression-based training occurs on the input variable itself. Autoregressive models predict their own input features one step at a time and keep predicting longer-term future horizons by using the previous output as input for the next step (Fig. 4). The estimator $\hat{f}(x_{t-k}, \dots, x_{t-1})$ makes predictions for samples \hat{x}_t from measurements of the target x_{t-k}, \dots, x_{t-1} and continues predicting autoregressively $\hat{x}_{t+1}, \hat{x}_{t+2}, \dots, \hat{x}_{t+N}$, using the previous outputs $\hat{x}_t, \hat{x}_{t+1}, \dots, \hat{x}_{t+N-1}$.

This proceeds iteratively as:

$$\text{STEP 0 : } \hat{f}(x_{t-k}, \dots, x_{t-1}) = \hat{x}_t$$

$$\text{STEP 1 : } \hat{f}(x_{t-k+1}, \dots, x_{t-1}, \hat{x}_t) = \hat{x}_{t+1}$$

...

$$\text{STEP } i : \hat{f}(x_{t-k+i}, \dots, x_{t-1}, \hat{x}_t, \dots, \hat{x}_{t-1+i}) = \hat{x}_{t+i}$$

where $i = 0, \dots, N$. During AR training we apply the Teacher forcing technique. Teacher forcing is known to introduce a so-called exposure bias (He et al., 2019), since the training procedure focuses on predicting the next step, while at inference the model predicts from history and from its own (auto-regressive) predictions. However, it has been shown that the exposure bias may effectively be neglected in most cases (He et al., 2019).

With our use of teacher forcing, the prediction at time $t+1$ uses ground truth at t rather than the prediction of t , and proceeds iteratively as:

$$\text{STEP 0 : } \hat{f}(x_{t-k}, \dots, x_{t-1}) = \hat{x}_t$$

$$\text{STEP 1 : } \hat{f}(x_{t-k+1}, \dots, x_{t-1}, x_t) = \hat{x}_{t+1}$$

...

$$\text{STEP } i : \hat{f}(x_{t-k+i}, \dots, x_{t-1}, x_t, \dots, x_{t-1+i}) = \hat{x}_{t+i}$$

Thus, we introduce and accept differences between the train and test protocol (i.e., applying the “teacher-forcing” during training) but we still leverage the full data set. This training is more efficient because data are more informative than the predicted values. Model validation and testing are done without teaching-forcing, because we lack ground truth in those cases. The AR approach has several advantages. First, it is self-supervised and does not require labelling. This also extends to TTF: in fact, TTF is computed from the actual stress time series. So it is computed from the actual sequence of stress values for the (self-supervised) training and it is computed from the forecast stress values at test. Since TTF is a quantity that is manually built, this may be prone to ambiguity due to pre-processing. Also, AR has the potential to predict beyond the TTF estimates, by predicting multiple cycles into the future.

For AR training we set the time interval for input history k (from x_{t-k} to x_{t-1}) as the “steps in” variable. Then to establish the prediction time (from y_t to y_{t+N}) we use length N as the “steps out” variable. The values of k and N for training are also used for validation and testing. While it is possible to increase N and extend the forecast horizon, we expect a deterioration in performance for values greater than the N used for training, and we explore such work below. We implemented the AR work in Pytorch, an open source machine learning library based on the Torch library (<https://pytorch.org/>). Additional details of the procedure are provided in the Supplement.

3.3.2. Forecasting with LSTM

A key component of an LSTM (Fig. 4a) is the memory cell that regulates information flow using three gates (i.e., Forget gate, Input gate, Output gate). The Forget gate deletes useless information by not releasing it to the next stage. The Input gate regulates new information into the network. The Output gate decides which part of the memory to output from the cell. During the training process, inputs and model weights pass to all gates, which in turn are connected to the self-recurrent memory cell (Fig. 4a). Our model has 3 stacked layers with size 300, for a total of 1808701 parameters (see Supplement 9.1.2 for details).

3.3.3. Forecasting with Temporal Convolutional Networks (TCN)

Temporal Convolutional Networks (TCN) (Fig. 4b) are Convolutional Neural Networks (CNN) that have been adopted for sequence modelling because of their performance (Dessi and Baroni, 2019). TCN consists of causal and dilated 1D convolutional layers with the same input and output lengths (Bai et al., 2018). Here, the term causal refers to the fact that convolution is applied only with the present and past elements but not the future. The term dilation in the context of a convolutional layer refers to the distance between the elements of the input sequence that are used to compute one entry of the output sequence: this increases the receptive field in each layer, making it possible to model long temporal patterns. In sequence modelling, TCN can be viewed as the process of sliding a 1D-window over the sequence to predict the part of the sequence with the length of the receptive field, using the chosen kernel. Such predictions are passed to the subsequent layers and the procedure continues until the receptive field has the size of

the input sequence. The convolutions in TCN can be parallelized at training, because the filter that extracts the features, used in each layer, can be computed jointly. RNNs instead need to be unrolled during backpropagation and can not be parallelized.

We adopt a model composed of three convolutional 1D layers and scan the sequence in a causal fashion, with dilation=1. Our models have a hidden size of 64 in the first layer and 256 in the second layer. The last layer has the dimension of the output, which is 1 in this case (because the model forecasts the next step). The model has a total of 29761 parameters, which is 2 orders of magnitude less than what we use for LSTM.

3.3.4. Forecasting with Transformer Network

Our Transformer Network (TF) consists of a modular architecture with an encoder and a decoder that contains three blocks: attention modules (self-attention and encoder-decoder attention), a feed-forward fully-connected module, and residual connections (Fig. 4c). The network's capability to capture sequence nonlinearities lies mainly in the attention module. TF maintains the encoding output (memory) separate from the decoded sequence, which means that training is parallelizable. TF is parallelizable also because of the absence of the internal status (as in LSTMs), so interactions between input and output are direct and do not need recursion for loops of backpropagation.

The model we adopt is based on the work of Giulieri et al. (2020). This model has $d_{model}=128$, 2 layers and 4 attention heads, for a total of 663938 parameters. We use RMSE and train the network via backpropagation with the NoamOpt optimizer; dropout value of 0.1 (Supplementary Material 9.1.3).

Although TF is often favoured for its high performance, a weakness is the need for huge training datasets. This explains the poor performance of TF for our experiments. Thus to improve TF we pre-trained using a sine function as input. This allows the model to learn the oscillatory behaviour of the experiments. We used a pre-training dataset of $1e7$ rows to describe a long set of sine waves with the same resolution of our lab data. The sine function roughly matches our lab data with amplitude equal to one frequency of 0.1 Hz and sample rate of 1000 Hz. This approach shows that 1 or 2 epochs are enough (an epoch is one complete pass of the training dataset through the algorithm) to pre-train (Section 9.1.4). After pre-training, we fine-tune the model using training data for each experiment.

4. Results

4.1. Prediction

Shear stress and time to failure are reasonably well predicted by the LSTM+CNN architecture (Fig. 3). The predictions match ground truth with an accuracy $> 93\%$ (Table 1). Fig. 5 shows model predictions and indicates that long term memory length is a key hyperparameter.

4.1.1. Best length for the past memory of LSTM

We investigate LSTM long term memory duration k using both RMSE and R^2 to assess performance. The optimum length for k is about one seismic cycle (Fig. S3). Note that the red lines show the optimum k . For experiments p4581 and p5198 we adopt a lower resolution because of event periodicity and because we compute variance every 0.1 seconds (see Section 3.1 and Table S1). Thus, one seismic cycle corresponds to about 100 data points and the best length for the observation is $k = 70$ or 7 seconds. For p4679 we adopt a higher resolution, 0.003 s, to describe both fast and slow events. Here, one seismic cycle corresponds to about 1700 data points, and the optimum value of $k = 2000$ (~ 6 seconds) for

predicting shear stress (Fig. S3). For predicting TTF, the optimum value of $k = 1100$ (~ 3 s).

We found high R^2 values for all three experiments using an observation length of about one seismic cycle. This suggests that there is a saturation of performance at one seismic cycle, so this may be sufficient to understand the signal. After one cycle there is a decrease in performance, due to experiment aperiodicity and LSTM memory limitations.

LSTMs struggle to learn long-term trends. The presence of slow and fast events in p4679 requires higher resolution to calculate AE variance. Compared to sequences of quasiperiodic events, p4679 requires about a factor of 10 more data. Long seismic cycles with rapid stress drops represent challenging conditions for LSTM. The problem arises because of observation lengths and the fact that Forget gates tend to remove too much information.

4.1.2. Prediction dataset split

We train and validate with 70% of the data and test with 30% (Fig. 2). Validation data were chosen randomly from the first 70% of the data and this value (10%) was removed from the training set. Thus the final division was: 63% for training, 7% for validation and 30% for testing. Details in Table S2. We normalized our data using: $X_{norm} = \frac{X - \min(X_{train})}{\max(X_{train}) - \min(X_{train})}$.

4.1.3. Prediction results

Fig. 5 and Table 1 summarize prediction results. Black lines show measurements and the coloured lines (green, yellow and red) are predictions for shear stress, TTsF and TTeF, respectively. Note that the predictions are quite accurate. The model is able to accurately predict shear stress, with $R^2 > 0.9$. Also TTF predictions are accurate (R^2 up to 90%), even if noisy in a few cases. We observe a general trend of better performance for TTeF than TTsF. Also the performance is better for experiments p5198 and p4679 than for p4581. Fig. 2 shows that TTeF is maximum where shear stress is minimum and variance is maximum. The peak variance scales directly with stress drop amplitude, with smaller values of peak variance preceding labquakes with smaller stress drop. As a consequence of smaller stress drop the time to reach a critical failure stress is smaller for constant loading rate. Indeed, the maximum shear stress and the slope of the restrengthening phase are quite similar in all the seismic cycles. In contrast, the values of TTsF derived from maximum shear stress show greater variability –possibly because of creep prior to the stress drop.

We used only AE variance as model input to predict shear stress and TTF and found the same result for data from each AE sensor. Our model performance for each experiment exceeds the state-of-the-art (Table S3) as defined by existing works (Rouet-Leduc et al., 2017, 2018). While this suggests that we could further improve model performance, by using more complex statistical features, rather than just variance. However, we did not pursue this direction and instead focused on new DL models.

4.2. Experiments on forecasting

Autoregressive forecasting models provide a method to predict shear stress variations based on past observations. We test AR models using our three DNN architectures: LSTM, TCN and TF. We analyzed all experiments but focus here on p4679 because of its complex behaviour and aperiodic seismic cycles. For this task, we decimated shear stress data to $dt = 0.1$ s and we apply an average running mean (Table S1). Although this reduces the original data resolution, downsampling is necessary for computational load and to limit the size of the step-in hyperparameter, which is particularly important for representing multiple seismic cycles. To establish the AR models we use $k = 200$ (20 seconds) which is a

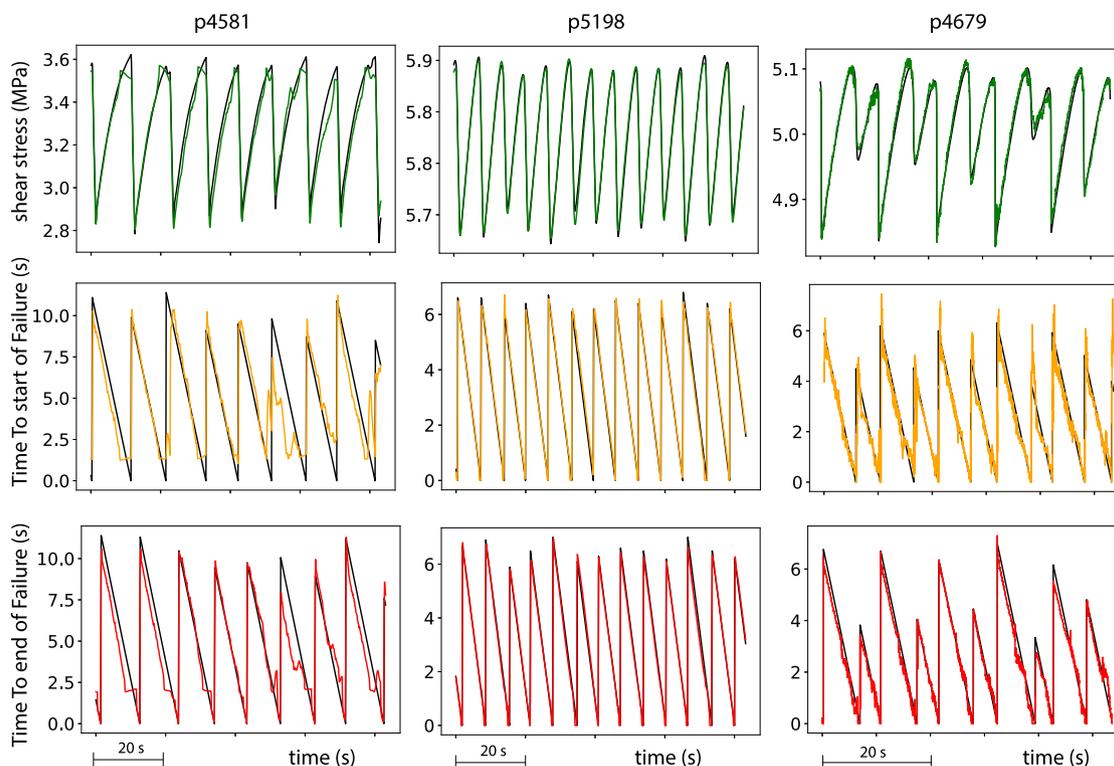


Fig. 5. Results of the model predictions for three experiments. Black lines show lab measurements (ground truth) and coloured lines are ML predictions. Note that shear stress (green lines) and time to end of failure (red lines) are well predicted in all the experiments (see also Table 1). Predictions of time to start of failure (orange lines) are also quite good in general, excepting a few sections of p4581. For p4679 we can see that, even if the prediction seems a bit noisy, the behaviour of the function is always well predicted.

value that both describes seismic cycles and fits GPU memory limitations.

Our data include many seismic cycles but this number is actually quite small for training DL models. Thus, we forecast using overlapping windows. To compare performance of our DL architectures (LSTM, TCN and TF) for each experiment, we set up the model to predict 100 data points (10 seconds), which corresponds to 1-2 seismic cycles depending on the experiment. See Fig. 6 for a sketch of these overlapped sequences. Shear stress forecasting is possible, although prediction accuracy varied between the DL models.

4.2.1. Forecasting dataset split

Our AR work used the same normalization as for prediction model, described above, while the division into train / validation / test subsets is different. The training set consists of the first 70% of the data and the testing set consists of the last 20%. Model validation was done with the remaining 10% of the data. Details in Table S2. Model performance was evaluated for future predictions in each window (Fig. 6). We also do an average among windows to measure overall performance (Table 2). Different data segments result in different performances, thus in Fig. 6 we show predictions for a range of different starting locations within the data stream. Note that AR models perform well for a range of starting times and data segments.

4.2.2. LSTM forecasting results

LSTM models generally produced the poorest predictions of future stress states (Fig. 7). Our AR forecast results are reasonable for only p5198, which has the least complex seismic cycles. For the other experiments LSTM did not provide accurate predictions. While LSTM is generally good for time series forecasting, because of its flexibility and optimization, it has trouble with data from our

experiments because of the data density, even if significantly decimated ($dt = 0.1$ s), and seismic cycle length (of order 10 s). LSTM works well with fewer data points. For example, in the work of Giuliani et al. (2020) they have 8 data points in input and 12 data points in output, whereas we have 200 in input and 100 in output. In essence, for our experiments the LSTM network “forgets” the past too quickly and does not predict accurately the signal in the future.

4.2.3. Temporal Convolution Network forecasting results

While TCN is the simplest of our models, it produced some of the best AR results. The TCN models provide the best compromise between training complexity and the number of parameters required. TCN requires only one tenth of the parameters needed for LSTM and TF. From Table 2 and Fig. 6 we can appreciate the TCN capability for forecasting. Results for experiment p5198 are quite good as are those for the more complex case of p4679. For p4581, TCN is the best of the tested models based on average forecast accuracy. However this experiment turned out to be challenging for each of the AR models as none of the forecasts were very good (Fig. 6), possibly due to the lack of appreciable preseismic creep in this experiment.

4.2.4. Transformer Network forecasting results

TF is the second best model, after TCN. It is not the best on average in any experiment, however with the flexible attention focusing it captures higher order variations of the data. For the same reason it does not forecast well the average behaviour of the signal, as a simpler model like TCN does. TF is the most complex among the tested models for optimization and training, and it is also the most flexible, because of its complex connections. It requires a lot of data to start working, and for this reason we designed a pretraining stage in which we fed the model sine waves with a frequency similar to our data, as described in 3.3.4.

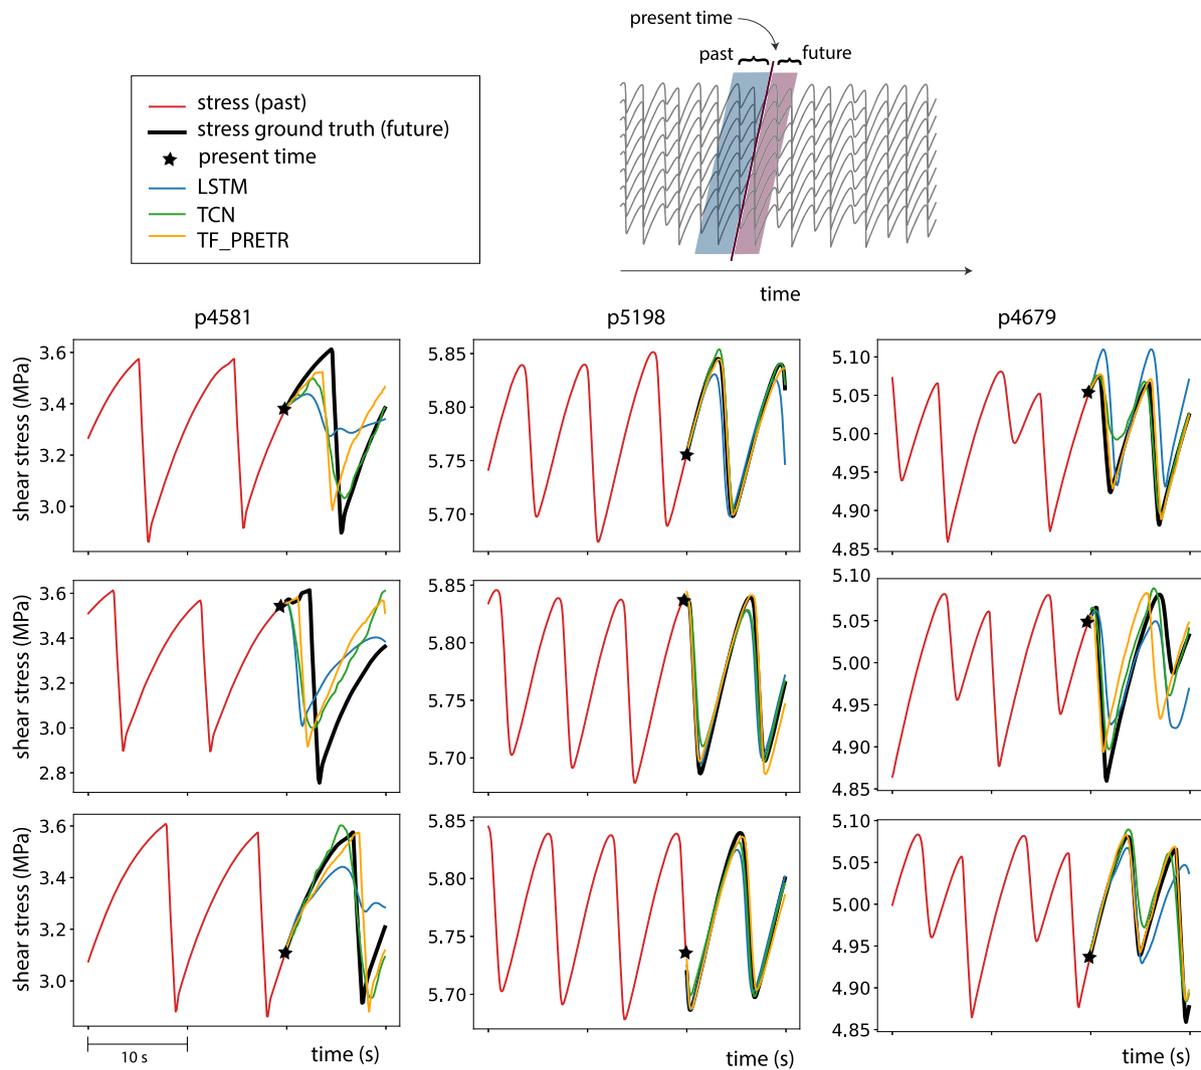


Fig. 6. Results of the autoregressive (AR) forecasting models for three experiments. Forecasts vary throughout the experiment, so for each experiment we show three time windows for the forecast. In each case red lines show stress measurements used as model input, black lines show ground truth future stress and coloured lines show model output forecasts. Forecasts are poor for experiment p4581 and the accuracy varies significantly from one window to another. For p4581, TF does better than LSTM and TCN at times. The forecasts are best for p5198, where all models predict the target quite well. Experiment p4679 shows a complex set of large and small events and is the most challenging. Forecasts here are quite variable depending on the time interval. However the models are able to predict reasonably well, especially TCN and TF. In each case, LSTM provides the poorest fits.

In several segments of p4581 TF is capable of detecting irregularities in the seismic cycle (Fig. 6). However, in other places TF tries to predict irregularities that are noise rather than signal. One example of this occurs close to the local maxima in Fig. 6 for p4581. Noise features like this caused trouble for TF. On the other hand, TF did well to forecast the aperiodicity of p4679 (Fig. 6). For p5198, TF did well, possibly because the behaviour is quite periodic. This is an indication that data overfitting during training may be a problem. We note that the improvement based on simple pre-training using a periodic signal provides some interesting future direction.

4.2.5. Forecasting results: model comparison and analysis

A key comparison between our AR models is how they perform with respect to the longer term temporal variations of the seismic cycle. Thus we evaluate results for several different data segments within the test set (Fig. 7). Note that the performances in Table 2 are averaged for the entire test set. In Fig. 7 we plot results for windows starting at several positions within the data. We can see that all the models have high variability depending on the tested window, hence the capability of the model in predicting the signal

depends on the shape and level of irregularity of the input-output window. As expected from our summary of general results, LSTM produced the worst AR forecasts. The other three models are in general able to forecast the signal with quite low variability. We note also that for p4581 the performance decreases with time in some areas simply because the signal becomes more challenging, for example where there are irregularities before peak stresses as seen on the left in Fig. 7.

An important question for the AR forecasting is that of how far into the future these models can predict and in particular if they can predict beyond a single seismic cycle. Fig. 8 summarizes such testing. Here we plot the average of all the performance values for each forecast time in the future. Our models are trained to forecast 100 data points in the future (white section), however we extended these to forecast 200 data points in the future (grey section). LSTM simply gets worse as time is extended. TCN decreases in performance slowly, whereas TF has a peculiar behaviour. For forecasting times from 0 to 100, TF is similar to TCN, while afterwards it becomes suddenly worse. This is because in the training phase TF forecasts from 0 to 100, but the behaviour is somewhat different in the next seismic cycle – from [101, 200]. This is per-

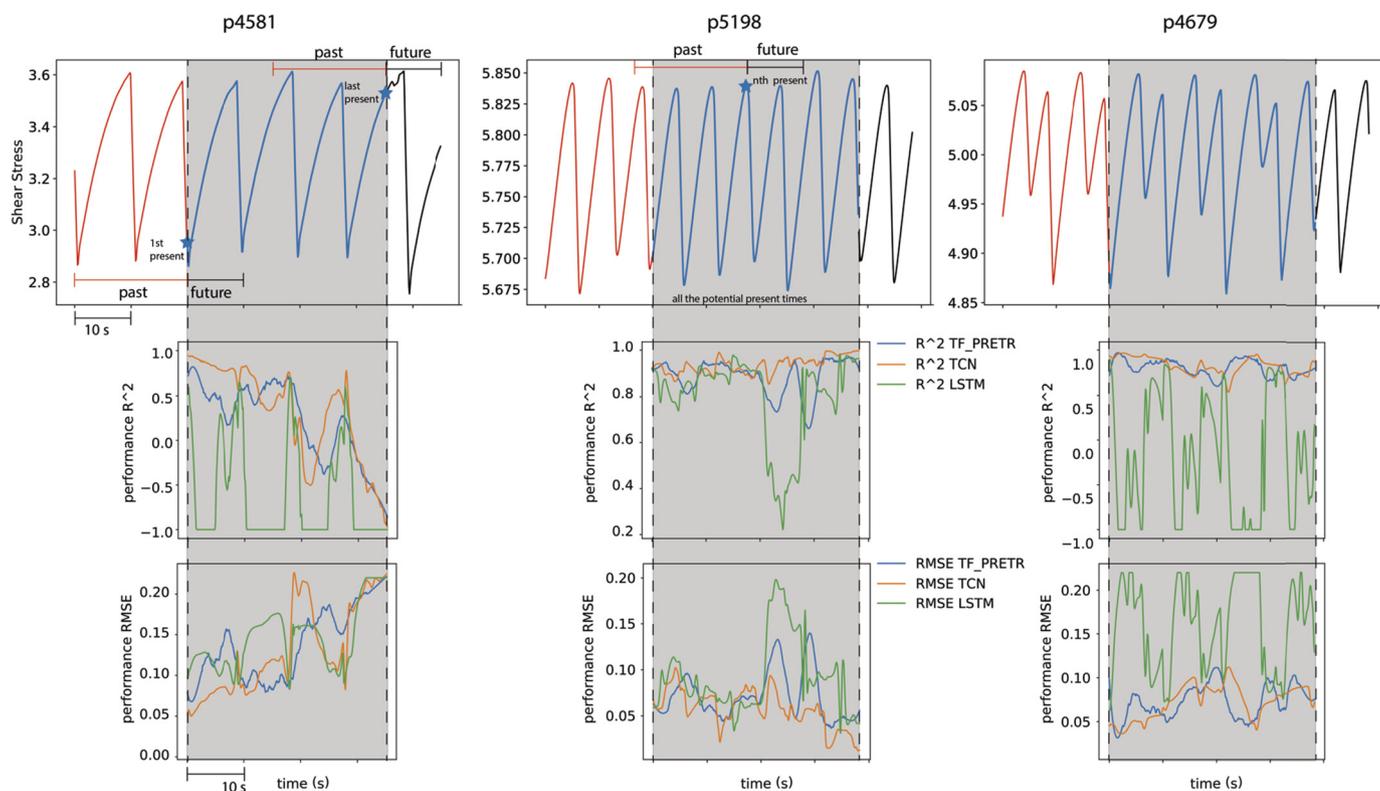


Fig. 7. Results of the AR forecasting shown as performance variation with respect to present time. Red and black sections represent input for the firsts windows and ground truth for the lasts windows (respectively). Blue sections are times when model predictions are compared to data, with performance shown below (stars in Fig. 6). Note that R^2 metric is not shown below -1 and RMSE is not shown above 0.22 . For TF we show a smoothed version of the original, noisier data because TF is the most complex model with the largest variation in performance.

haps surprising given that TF is renowned for generalization. Here, it is possibly the result of data complexity. This is clearly an important question for future work.

5. Discussion

The first part of our work was devoted to evaluating DL models based on a combination of LSTM and CNN. Compared to the state-of-the-art based on existing works we find that DL models perform well. As suggested by previous works, we used the variance of the continuous acoustic emissions emanating from the laboratory faults. The DL models work reasonably well for all tested experiments, however results are better when creep occurs before the mainshock (i.e. as in p5198 and p4679). This is perhaps not surprising given that creep represents fault slip which could result in micro fracture and breakage of frictional asperity contacts. Moreover, the shear stress curves reflect creep and AE via the gradual reduction in the rate of increase prior to reaching a peak at the onset of a labquake (Fig. 2). In the absence of creep, such as in p4581 (Fig. S2), the acoustic variance is initially very low and increases suddenly at the onset of failure. As a result the DL models struggle to identify the exact time of the drop in shear stress. This appears to be part of the reason why Time To Start of Failure (TTsF) is always more challenging to predict than Time To End of Failure (TTeF), because the former is zero when the variance is beginning to grow, while the latter is zero right after the variance has first increased and then decreased.

The limitation of TTF is that it represents the time remaining for just one event. Moreover this quantity is not recorded directly during the experiment in the laboratory. Therefore we manually labelled a suitable dataset for it, before the model training phase. The autoregressive models solve this problem nicely because they are designed to forecast.

Our autoregressive forecasting procedure allows one to predict future values of shear stress during the lab seismic cycle. To the best of our knowledge, this is the first time such an AR approach has been used for lab earthquake data. Here, the innovation is that we can indefinitely forecast into the future without the need for labelling because we use the same feature (shear stress) for both model input and output.

We tested several DL networks for the AR forecasting. All of the models work at some level. The LSTM network produced the poorest results. In time series forecasting, LSTM is one of the most common models, because of its flexibility, but in our application it has memory problems due to the considerable length of the sequences. TCN is the simplest model in terms of number of parameters and structure, and it performed better, perhaps because the target (shear stress) during the lab seismic cycle is somewhat periodic. TF is also flexible but it is the most complex for optimization and requires a lot of data to initiate a reasonable model because of the huge number of parameters to be trained. We dealt with that problem by pretraining the TF models with sine waves, however that represents another step in the processing.

Of the experiments we evaluated, p5198 was fit best by the DL models. Experiment p4679 was challenging because of its aperiodicity and the presence of slow and fast events. However our TCN and TF models were able to forecast reasonably well. The forecasts for p4581 were more challenging, perhaps because of the lack of appreciable fault creep and/or because of noise in the data. Further work is needed to resolve these issues.

Overall, DL is capable of understanding the behaviour of the signal and this opens new perspectives in seismology research with the use of AR methods. First of all, this confirms that there are some patterns in the signal that the model is able to recognize in most of the cases. Moreover, the model after training, just needs a limited knowledge about the past of a specific window to fore-

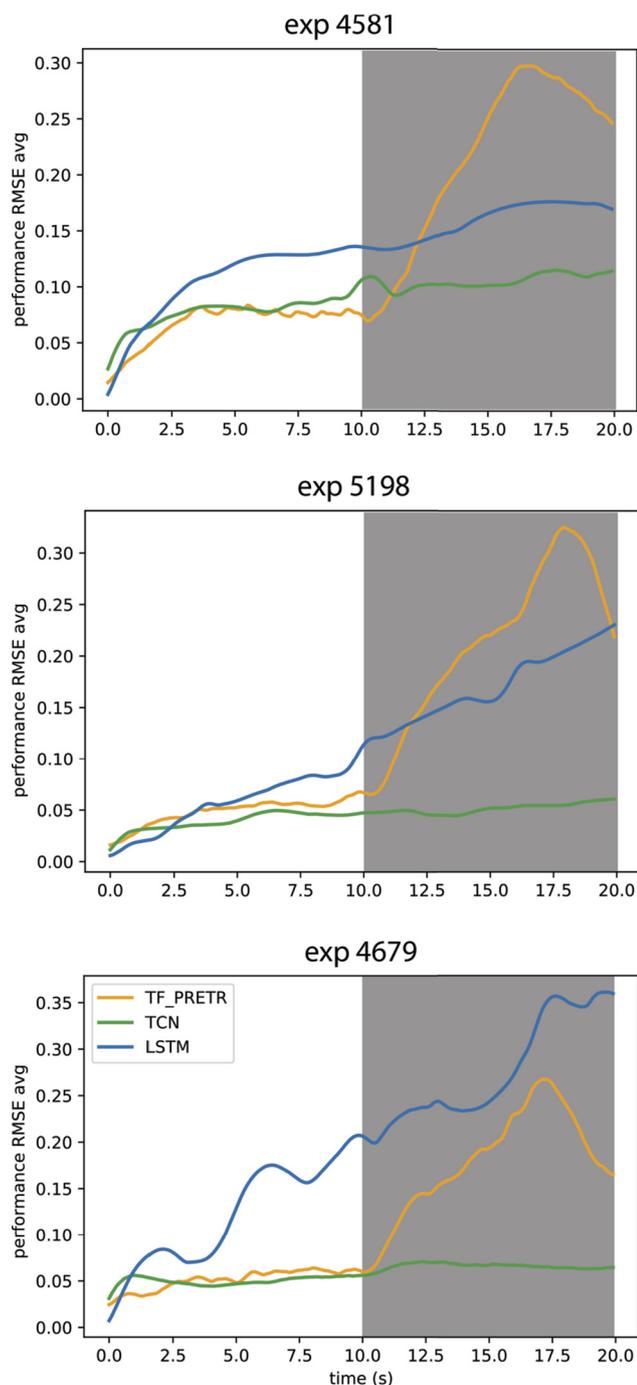


Fig. 8. Data showing how AR model performance evolves using a 20 second window for forecasting. Note that the lab seismic cycle is variable but typically < 10 sec. (Fig. 6). Performance values are the RMSE of the average of all the i -th at each time predictions compared with the i -th time ground truth value, in all the windows in the test set. We use only RMSE because it is not possible to compute R^2 point by point. White section is for times from present to present+10 s (100 steps in the future). Gray section is for times from present+10 s to present+20 s (from the 100th to the 200th step in the future). The model has been trained to predict 100 steps in the future, so this is just a test to investigate the generalizability of the procedure.

cast reasonably the future. This provides insight on the physical processes and their deterministic nature.

One good way to advance our work would be to apply the method on real fault data. This would require measurements of shear stress, which are not available. The idea could be to use seismograms directly prior to or during an earthquake and infer the shear stress using the prediction method, such as we did with AE

in the lab. This would allow forecasts of the shear stress for future time steps, using the forecasting method. This procedure has the limitation that labquakes are simplistic compared to earthquakes, and this is one of the reasons why it is important to improve our understanding of ML methods for prediction and forecasting. Another approach would be to simply predict the temporal evolution of the seismogram based on the initial part of the signal. A key question here is then, are lab AE signals uniquely related to fault shear zone stress and if so, would that transfer to tectonic faults? Yet another approach would be transfer learning from lab earthquakes to real fault data.

6. Conclusions

We used Deep Neural Networks to predict and forecast laboratory earthquakes and lab measurements of fault zone shear stress based on seismic signals emanating from lab faults. Previous work showed that using the variance of lab seismic signals (from fault zone acoustic emissions) it is possible to predict fault shear stress and the time to failure. We systematically tested a range of DL models with a variety of lab faults and found that our models significantly outperform the state-of-the-art. Moreover we proved that it is possible to forecast future fault zone shear stress based on the previous history of stress values. This result has significant potential because shear stress is representative of the state of the fault and forecasting it in time means implicitly to forecast the timing of the future failure states. To our knowledge, this is the first application of a forecasting procedure with the goal of inferring autoregressively the future shear stress knowing the stress itself in the past.

CRediT authorship contribution statement

Laura Laurenti: Conceptualization, Data curation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Elisa Tinti:** Conceptualization, Investigation, Writing – original draft, Writing – review & editing. **Fabio Galasso:** Methods, Data curation, Validation, Supervision, Writing – review & editing. **Luca Franco:** Methods, Validation, Writing – review & editing. **Chris Marone:** Conceptualization, Data collection, Investigation, Methods, Supervision, Resources, Validation, Software, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data and resources

All simulation input files and the Jupyter notebooks are accessible at GitHub: <https://github.com/lauralaurenti/DNN-earthquake-prediction-forecasting>.

Acknowledgements

We would like to thank M. Denolle and two anonymous reviewers for helpful comments that improved the manuscript. We would also like to thank M. Scuderi, C. Colletini, P. Johnson, and G. Paoletti for helpful discussions. L.L., E.T., and C.M. were supported by European Research Council Advance grant 835012 (TECTONIC). C.M. also acknowledges support from US Department of Energy grants DE-SC0020512 and DE-EE0008763.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.epsl.2022.117825>.

References

- Acosta, M., Passelègue, F.X., Schubnel, A., Madariaga, R., Violay, M., 2019. Can precursory moment release scale with earthquake magnitude? A view from the laboratory. *Geophys. Res. Lett.* 46, 12927–12937. <https://doi.org/10.1029/2019GL084744>.
- Allen, R.M., Stogaitis, M., 2022. Global growth of earthquake early warning. *Science* 375, 717–718. <https://doi.org/10.1126/science.abc5435>.
- Bai, S., Koller, J.Z., Koltun, V., 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *CoRR*. arXiv:1803.01271.
- Ben-Zion, Y., Lyakhovskiy, V., 2001. Accelerated seismic release and related aspects of seismicity patterns on earthquake faults. *Pure Appl. Geophys.* 159, 2385–2412.
- Beroza, G., Segou, M., Mousavi, S., 2021. Machine learning and earthquake forecasting—next steps. *Nat. Commun.* 12. <https://doi.org/10.1038/s41467-021-24952-6>.
- Bolton, D.C., Shokouhi, P., Rouet-Leduc, B., Hulbert, C., Rivière, J., Marone, C., Johnson, P.A., 2019. Characterizing acoustic signals and searching for precursors during the laboratory seismic cycle using unsupervised machine learning. *Seismol. Res. Lett.* 90, 1088–1098. <https://doi.org/10.1785/0220180367>.
- Bolton, D.C., Shreedharan, S., Rivière, J., Marone, C., 2020. Acoustic energy release during the laboratory seismic cycle: insights on laboratory earthquake precursors and prediction. *J. Geophys. Res., Solid Earth* 125, e2019JB018975. <https://doi.org/10.1029/2019JB018975>.
- Bolton, D.C., Shreedharan, S., Rivière, J., Marone, C., 2021. Frequency-magnitude statistics of laboratory foreshocks vary with shear velocity, fault slip rate, and shear stress. *J. Geophys. Res., Solid Earth* 126, e2021JB022175. <https://doi.org/10.1029/2021JB022175>.
- Chaiornkaew, L., Elston, H., Cooke, M., Mukerji, T., Graham, S.A., 2022. Predicting off-fault deformation from experimental strike-slip fault images using convolutional neural networks. *Geophys. Res. Lett.* 49, e2021GL096854. <https://doi.org/10.1029/2021GL096854>.
- Corbi, F., Sandri, L., Bedford, J., Funicello, F., Brizzi, S., Rosenau, M., Lallemand, S., 2019. Machine learning can predict the timing and size of analog earthquakes. *Geophys. Res. Lett.* 46, 1303–1311. <https://doi.org/10.1029/2018GL081251>.
- Denolle, M., Dunham, E., Prieto, G., Beroza, G., 2014. Strong ground motion prediction using virtual earthquakes. *Science* 343, 399–403. <https://doi.org/10.1126/science.1245678>.
- Dessi, R., Baroni, M., 2019. CNNs found to jump around more skillfully than RNNs: compositional generalization in seq2seq convolutional networks. arXiv:1905.08527.
- Dieterich, J.H., 1978. Preseismic fault slip and earthquake prediction. *J. Geophys. Res., Solid Earth* 83, 3940–3948. <https://doi.org/10.1029/JB083iB08p3940>.
- Dresen, G., Kwiatek, G., Goebel, T.H.W., Ben-Zion, Y., 2020. Seismic and aseismic preparatory processes before large stick-slip failure. *Pure Appl. Geophys.* 177, 5741–5760.
- Giuliari, F., Hasan, I., Cristani, M., Galasso, F., 2020. Transformer networks for trajectory forecasting. arXiv:2003.08111.
- Goebel, T.H.W., Kwiatek, G., Becker, T.W., Brodsky, E.E., Dresen, G., 2017. What allows seismic events to grow big?: insights from b-value and fault roughness analysis in laboratory stick-slip experiments. *Geology* 45, 815–818.
- He, T., Zhang, J., Zhou, Z., Glass, J.R., 2019. Quantifying exposure bias for neural language generation. ArXiv arXiv:1905.10617 [abs].
- Hedayat, A., Pyrak-Nolte, L., Bobet, A., 2014. Precursors to the shear failure of rock discontinuities. *Geophys. Res. Lett.* 41. <https://doi.org/10.1002/2014GL060848>.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hulbert, C., Rouet-Leduc, B., Jolivet, R., Johnson, P., 2020. An exponential build-up in seismic energy suggests a months-long nucleation of slow slip in Cascadia. *Nat. Commun.* 11, 4139. <https://doi.org/10.1038/s41467-020-17754-9>.
- Hulbert, C., Rouet-Leduc, B.P.G., Johnson, P.A., Ren, C., Rivière, J., Bolton, D.C., Marone, C., 2018. Similarity of fast and slow earthquakes illuminated by machine learning. *Nat. Geosci.* 12, 69–74.
- Jaspersen, H., Bolton, D.C., Johnson, P., Guyer, R., Marone, C., de Hoop, M.V., 2021. Attention network forecasts time-to-failure in laboratory shear experiments. arXiv:1912.06087.
- Johnson, C., Hulbert, C., Rouet-Leduc, B., Johnson, P., 2020. Learning the low frequency earthquake daily intensity on the central San Andreas Fault. <https://doi.org/10.1002/essoar.10504699.1>.
- Johnson, P., Rouet-Leduc, B., Pyrak-Nolte, L., Beroza, G., Marone, C., Hulbert, C., Howard, A., Singer, P., Gordeev, D., Karaflos, D., Levinson, C., Pfeiffer, P., Puk, K.M., Reade, W., 2021. Laboratory earthquake forecasting: a machine learning competition. *Proc. Natl. Acad. Sci.* 118, e2011362118. <https://doi.org/10.1073/pnas.2011362118>.
- Johnson, P., Ferdowsi, B., Kaproth-Gerecht, B., Scuderi, M., Griffa, M., Carmeliet, J., Guyer, R., Le Bas, P.Y., Trugman, D., Marone, C., 2013. Acoustic emission and microslip precursors to stick-slip failure in sheared granular material. *Geophys. Res. Lett.* 40. <https://doi.org/10.1002/2013GL057848>.
- Karner, S., Marone, C., 1998. The effect of shear load on frictional healing in simulated fault gouge. *Geophys. Res. Lett.* 25, 4561–4564. <https://doi.org/10.1029/1998GL900182>.
- Kohler, M., Smith, D., Andrews, J., Chung, A., Hartog, R., Henson, I., Given, D., Groot, R., Guiwits, S., 2020. Earthquake early warning shakealert 2.0: public rollout. *Seismol. Res. Lett.* 91. <https://doi.org/10.1785/0220190245>.
- Kong, Q., Trugman, D.T., Ross, Z.E., Bianco, M.J., Meade, B.J., Gerstoft, P., 2018. Machine learning in seismology: turning data into insights. *Seismol. Res. Lett.* 90, 3–14. <https://doi.org/10.1785/0220180259>.
- Kwiatek, G., Goebel, T.H.W., Dresen, G., 2014. Seismic moment tensor and b value variations over successive seismic cycles in laboratory stick-slip experiments. *Geophys. Res. Lett.* 41, 5838–5846.
- Latour, S., Schubnel, A., Nielsen, S., Madariaga, R., Vinciguerra, S., 2013. Characterization of nucleation during laboratory earthquakes. *Geophys. Res. Lett.* 40, 5064–5069. <https://doi.org/10.1002/grl.50974>.
- Leeman, J., Saffer, D., Scuderi, M., Marone, C., 2016. Laboratory observations of slow earthquakes and the spectrum of tectonic fault slip modes. *Nat. Commun.* 7, 11104. <https://doi.org/10.1038/ncomms11104>.
- Lubbers, N., Bolton, D., Mohd-Yusof, J., Marone, C., Barros, K., Johnson, P., 2018. Earthquake catalog-based machine learning identification of laboratory fault states and the effects of magnitude of completeness. *Geophys. Res. Lett.* 45. <https://doi.org/10.1029/2018GL079712>.
- Ma, G., Mei, J., Gao, K., Zhao, J., Zhou, W., Wang, D., 2021. Machine learning bridges microslips and slip avalanches of sheared granular gouge. *Earth and Space Sci. Open Arch.* 13. <https://doi.org/10.1002/essoar.10506338.1>.
- Main, I., Meredith, P., Sammonds, P., 1992. Temporal variations in seismic event rate and b-values from stress-corrosion constitutive laws. *Tectonophysics* 211, 233–246.
- Main, I.G., Meredith, P.G., Jones, C.M., 1989. A reinterpretation of the precursory seismic b-value anomaly from fracture mechanics. *Geophys. J. Int.* 96, 131–138.
- Mair, K., Frye, K.M., Marone, C., 2002. Influence of grain characteristics on the friction of granular shear zones. *J. Geophys. Res.* 107, 2219.
- McBeck, J., Ben-Zion, Y., Renard, F., 2020. The mixology of precursory strain partitioning approaching brittle failure in rocks. *Geophys. J. Int.* 221, 1856–1872. <https://doi.org/10.1093/gji/ggaa121>.
- McBrearty, I., Delorey, A., Johnson, P., 2019. Pairwise association of seismic arrivals with convolutional neural networks. *Seismol. Res. Lett.* 90. <https://doi.org/10.1785/0220180326>.
- McLuskey, G.C., Lockner, D.A., 2014. Preslip and cascade processes initiating laboratory stick slip. *J. Geophys. Res., Solid Earth* 119, 6323–6336. <https://doi.org/10.1002/2014JB011220>.
- Mele Veedu, D., Giorgetti, C., Scuderi, M., Barbot, S., Marone, C., Collettini, C., 2020. Bifurcations at the stability transition of earthquake faulting. *Geophys. Res. Lett.* 47, e2020GL087985. <https://doi.org/10.1029/2020GL087985>.
- Mousavi, S., Ellsworth, W., Weiqiang, Z., Chuang, L., Beroza, G., 2020. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nat. Commun.* 11, 3952. <https://doi.org/10.1038/s41467-020-17591-w>.
- Nagata, K., Nakatani, M., Yoshida, S., 2008. Monitoring frictional strength with acoustic wave transmission. *Geophys. Res. Lett.* 35. <https://doi.org/10.1029/2007GL033146>.
- Passelègue, F.X., Latour, S., Schubnel, A., Nielsen, S., Bhat, H.S., Madariaga, R., 2017. Influence of Fault Strength on Precursory Processes During Laboratory Earthquakes. Chapter 12. American Geophysical Union (AGU), pp. 229–242.
- Pritchard, M.E., Allen, R.M., Becker, T.W., Behn, M.D., Brodsky, E.E., Bürgmann, R., Ebinger, C., Freymueller, J.T., Gerstenberger, M., Haines, B., Kaneko, Y., Jacobsen, S.D., Lindsey, N., McGuire, J.J., Page, M., Ruiz, S., Tolstoy, M., Wallace, L., Walter, W.R., Wilcock, W., Vincent, H., 2020. New opportunities to study earthquake precursors. *Seismol. Res. Lett.* 91, 2444–2447. <https://doi.org/10.1785/0220200089>.
- Raissi, M., Perdikaris, P., Karniadakis, G., 2019. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Rivière, J., Lv, Z., Johnson, P., Marone, C., 2018. Evolution of b-value during the seismic cycle: insights from laboratory experiments on simulated faults. *Earth Planet. Sci. Lett.* 482, 407–413. <https://doi.org/10.1016/j.epsl.2017.11.036>.
- Rouet-Leduc, B., Hulbert, C., Bolton, D.C., Ren, C.X., Riviere, J., Marone, C., Guyer, R.A., Johnson, P.A., 2018. Estimating fault friction from seismic signals in the laboratory. *Geophys. Res. Lett.* 45, 1321–1329. <https://doi.org/10.1002/2017GL076708>.
- Rouet-Leduc, B., Hulbert, C., Lubbers, N., Barros, K., Humphreys, C.J., Johnson, P.A., 2017. Machine learning predicts laboratory earthquakes. *Geophys. Res. Lett.* 44, 9276–9282. <https://doi.org/10.1002/2017GL074677>.
- Scholz, C., 1968. The frequency-magnitude relation of microfracturing in rock and its relation to earthquakes 58. <https://doi.org/10.1785/BSSA0580010399>.
- Scholz, C., 2015. On the stress dependence of the earthquake b-value. *Geophys. Res. Lett.* 42. <https://doi.org/10.1002/2014GL062863>.

- Schubnel, A., Benson, P., Thompson, B., Hazzard, J., Young, R.P., 2006. Quantifying damage, saturation and anisotropy in cracked rocks by inverting elastic wave velocities. *Pure Appl. Geophys.* 163, 947–973. <https://doi.org/10.1007/s00024-006-0061-y>.
- Scuderi, M., Collettini, C., Marone, C., 2017. Frictional stability and earthquake triggering during fluid pressure stimulation of an experimental fault. *Earth Planet. Sci. Lett.* 477, 84–96. <https://doi.org/10.1016/j.epsl.2017.08.009>.
- Scuderi, M., Marone, C., Tinti, E., Di Stefano, G., Collettini, C., 2016. Precursory changes in seismic velocity for the spectrum of earthquake failure modes. *Nat. Geosci.* 9. <https://doi.org/10.1038/ngeo2775>.
- Shokouhi, P., Girkar, V., Rivière, J., Shreedharan, S., Marone, C., Giles, C.L., Kifer, D., 2021. Deep learning can predict laboratory quakes from active source seismic data. *Geophys. Res. Lett.* 48, e2021GL093187. <https://doi.org/10.1029/2021GL093187>.
- Shreedharan, S., Bolton, D., Rivière, J., Marone, C., 2021. Competition between preslip and deviatoric stress modulates precursors for laboratory earthquakes. *Earth Planet. Sci. Lett.* 553. <https://doi.org/10.1016/j.epsl.2020.116623>.
- Shreedharan, S., Bolton, D.C., Rivière, J., Marone, C., 2020. Machine learning predicts the timing and shear stress evolution of lab earthquakes using active seismic monitoring of fault zone processes. *Earth and Space Sci. Open Arch.* 48. <https://doi.org/10.1002/essoar.10505562.1>.
- Thompson, B., Young, R.P., Lockner, D., 2005. Observations of premonitory acoustic emission and slip nucleation during a stick slip experiment in smooth faulted westerly granite. *Geophys. Res. Lett.* 32, 10304. <https://doi.org/10.1029/2005GL022750>.
- Tinti, E., Scognamiglio, L., Michelini, A., Cocco, M., 2016. Slip heterogeneity and directivity of the M_L 6.0, 2016, Amatrice earthquake estimated with rapid finite-fault inversion: rupture process of 2016 amatrice event. *Geophys. Res. Lett.* 43. <https://doi.org/10.1002/2016GL071263>.
- Trugman, D., McBrearty, I., Bolton, D., Guyer, R., Marone, C., Johnson, P., 2020. The spatiotemporal evolution of granular microslip precursors to laboratory earthquakes. *Geophys. Res. Lett.* 47. <https://doi.org/10.1029/2020GL088404>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. *arXiv:1706.03762*.
- Wang, K., Johnson, C.W., Bennett, K.C., Johnson, P.A., 2021. Predicting fault slip via transfer learning. *Nat. Commun.* 12. <https://doi.org/10.1038/s41467-021-27553-5>.