



OPEN

Prediction of creep failure time using machine learning

Soumyajyoti Biswas^{1,2}, David Fernandez Castellanos³ & Michael Zaiser¹✉

A subcritical load on a disordered material can induce creep damage. The creep rate in this case exhibits three temporal regimes viz. an initial decelerating regime followed by a steady-state regime and a stage of accelerating creep that ultimately leads to catastrophic breakdown. Due to the statistical regularities in the creep rate, the time evolution of creep rate has often been used to predict residual lifetime until catastrophic breakdown. However, in disordered samples, these efforts met with limited success. Nevertheless, it is clear that as the failure is approached, the damage become increasingly spatially correlated, and the spatio-temporal patterns of acoustic emission, which serve as a proxy for damage accumulation activity, are likely to mirror such correlations. However, due to the high dimensionality of the data and the complex nature of the correlations it is not straightforward to identify the said correlations and thereby the precursory signals of failure. Here we use supervised machine learning to estimate the remaining time to failure of samples of disordered materials. The machine learning algorithm uses as input the temporal signal provided by a mesoscale elastoplastic model for the evolution of creep damage in disordered solids. Machine learning algorithms are well-suited for assessing the proximity to failure from the time series of the acoustic emissions of sheared samples. We show that materials are relatively more predictable for higher disorder while are relatively less predictable for larger system sizes. We find that machine learning predictions, in the vast majority of cases, perform substantially better than other prediction approaches proposed in the literature.

All materials break under sufficiently high stress. However, even when the system can support a load at the instance of its application, it may still break at a later time by creep rupture¹. Local damage may accumulate even at a sub-critical loads. Accumulation of microstructural damage may be associated with the thermally activated crossing of energy barriers: examples include the accumulation of free volume as result of the thermal activation of shear transformations in amorphous materials^{2,3}, or the thermally assisted removal of dislocation barriers in irradiated metals leading to microstructural slip localization and irradiation embrittlement⁴. Local damage accumulation reduces the energy barriers for future damage activation, thus promoting a tendency to localization. Overall, creep deformation is generally known to have three temporal regimes. First, we observe a decelerating strain rate regime associated with (statistical) hardening or aging effects as the weakest elements of the microstructure deform first and become consequentially inactivated by internal back stresses³. The decelerating regime is followed by an intermediate regime of constant strain rate and a final accelerating strain rate regime, associated with damage accumulation and strain localization and leading to catastrophic breakdown².

For obvious reasons, understanding the creep failure dynamics is an important issue for stability analysis of structures across scales. Especially, predicting the residual lifetime of a given sample until its failure under a subcritical load is a question that is actively investigated by both physicists and engineers⁵. Reliable lifetime predictions might not only avoid catastrophic in-service failure of components and systems, but also yield substantial economic benefits in view of the possibility of extending replacement cycles. Sample specific information on the damage accumulation process can, on the one hand, be obtained from the macroscopic sample response, i.e., the time dependent creep strain or strain rate. More detailed information can be drawn from analysis of the spatio-temporal pattern of energy releases as local creep damage accumulates in a material subject to subcritical load. The idea is here that the introduction of local damage is accompanied by a release of elastic energy which can be

¹WW8-Materials Simulation, Department of Materials Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Dr.-Mack-Str. 77, 90762 Fürth, Germany. ²Department of Physics, SRM University - AP, Guntur, Andhra Pradesh 522502, India. ³PMMH, CNRS-UMR 7636, ESPCI Paris, PSL University, Sorbonne Université, Université de Paris, 75005 Paris, France. ✉email: michael.zaiser@fau.de

recorded by monitoring the acoustic emission (AE) of the sample, thus providing a means of non-destructively monitoring the damage accumulation process.

Among the empirical attempts to predict sample specific failure times from macroscopic creep strain rates, one possible approach is to correlate the time t_m of minimum strain rate with the catastrophic failure time t_f , in the simplest case by assuming a linear relationship between both^{6,7}. However, there are multiple issues in using that observation for failure time forecasting: (i) in analyzing time series for an individual sample, it is often difficult to identify a unique minimum for the strain rate. This problem is particularly pronounced when the creep strain rate is itself a stochastic, highly intermittent process; (ii) while empirical observation indicates, on average, a linear relation between t_m and t_f , the scatter is high especially for highly disordered samples; (iii) the prediction for t_f necessarily requires waiting until t_m can be reliably identified. Given that experimentally observed t_m already amount to 60% of t_f and that larger times are needed to reliably identify a minimum, the resulting prediction might be too late to be useful⁸.

A different prediction approach focuses on temporal statistics of the damage accumulation process as monitored by AE. In this case, one looks at the magnitudes, times, and possibly locations of acoustic emission events and tries to identify statistical correlations that allow to interpolate the time of failure. For instance, one may exploit the observation made both in simulations² and experiments⁹ that the AE event rate ν_{AE} may accelerate towards failure according to a reverse Omori law, $\nu_{AE} \propto (t - t_f)^{-p}$ with $p \approx 1$. Such a reverse Omori behavior was also reported to be a generic feature of mean-field models of thermally activated rupture processes¹⁰. In such situations, one can obtain the failure time by fitting the Omori law to the AE record until time t , with the advantage that (unlike predictions based on the strain rate minimum) the ensuing predictions continually improve with increasing record length, i.e. decreasing time-to-failure. At the same time, the approach to failure may be accompanied with other characteristic changes in the AE burst statistics, such as an increase in the AE event magnitude or characteristic changes in the Gutenberg-Richter exponent of the power law type statistics of burst energies^{2,11}, which may also be used for monitoring and prediction purposes.

Even further information can be harnessed by simultaneously monitoring the spatial pattern of damage accumulation, as failure is associated with localization of damage^{2,9,11}. Spatio-temporal correlations in energy release signals, therefore, may hold important information regarding distance to the catastrophic breakdown of the sample. However, given the high dimensionality of the data sets involved and the possible complexity in the correlation measures, it may not be possible to extract the necessary information regarding failure time in terms of simple empirical laws. Indeed, the task of extracting non-trivial correlations from high dimensional data is precisely what machine learning algorithms can do best. In recent times, machine learning found widespread applications in predicting deformation, failure, and flow processes in disordered systems based on complex data. Predictions of irreversible deformation and failure processes were based on data describing local atomic structure in amorphous solids^{12,13}, mesoscale microstructures¹⁴ (for an overview see e.g.¹⁵), as well as monitoring data obtained in macroscopic tests^{16–18}. Here we use Random Forest regression¹⁹ for extracting information regarding sample specific failure times from spatio-temporal records of energy release signals prior to failure. To avoid problems resulting from scarcity of data, we obtain our training and testing data from ensembles of creep rupture simulations performed using the model introduced in Ref.². The trained algorithm is tested over a set of samples previously unseen by the algorithm using various accuracy measures. We investigate the variations in prediction accuracy as a function of loading shear stress, the degree of microstructural disorder, and sample size.

Results

As a model for creep rupture, we use a mesoscale elastoplastic model^{2,3} that considers plastic activity accompanied by damage accumulation in a simulated sample which is driven by temporally constant, subcritical shear loads (see “Methods”). The sample volume is divided into mesoscopic volume elements. Local energy barriers control deformation and damage accumulation within the individual elements. The statistical distribution of these barriers characterizes the microstructural disorder of the material. The barrier height is reduced by stress, hence, if local stresses are high enough, barriers may be crossed and local plastic activity takes place. At the same time internal stresses, which arise from local deformation, couple the deformation response of the individual elements. Plastic deformation generates local damage which reduces, on average, the local barrier height. The coupling between deformation, internal stresses and damage accumulation ultimately leads to damage localization in the form of a macroscopic shear band. Such damage localization induces a divergence of the strain rate which indicates catastrophic failure. The model has been successful in reproducing the temporal regimes of creep, the statistics of damage accumulation in the form of avalanches, and the observation of progressive strain localization in the approach to failure^{2,3}. A detailed model description and default model parameters are provided in the “Methods” section.

The model produces, as raw data output, information that can be interpreted as a simulated Acoustic Emission time series: Deformation activity is characterized by the timings, locations, and amplitudes of deformation avalanches, as well as by the resulting spatio-temporal strain patterns. As shown in^{2,3}, the corresponding time series exhibit correlations which evolve with time. Examples are the variation of the statistics of avalanches or the progressive localization of spatial activity (see Fig. 1). These variations depend on the proximity to failure and can thus be envisaged as precursors with the potential for prediction. As explained in the “Methods” section, from the space-time series of deformation avalanches as produced by a mesoscale creep simulation we extract several features that are used for ML prediction of the failure time.

At each time t measured since the beginning of the creep process, the machine learning algorithm makes a prediction for the remaining time to failure, t_p . To each time t at which a prediction is made, we can *post mortem* assign an actual remaining time to failure t_a . Therefore, we define the (mean) fractional error of the machine

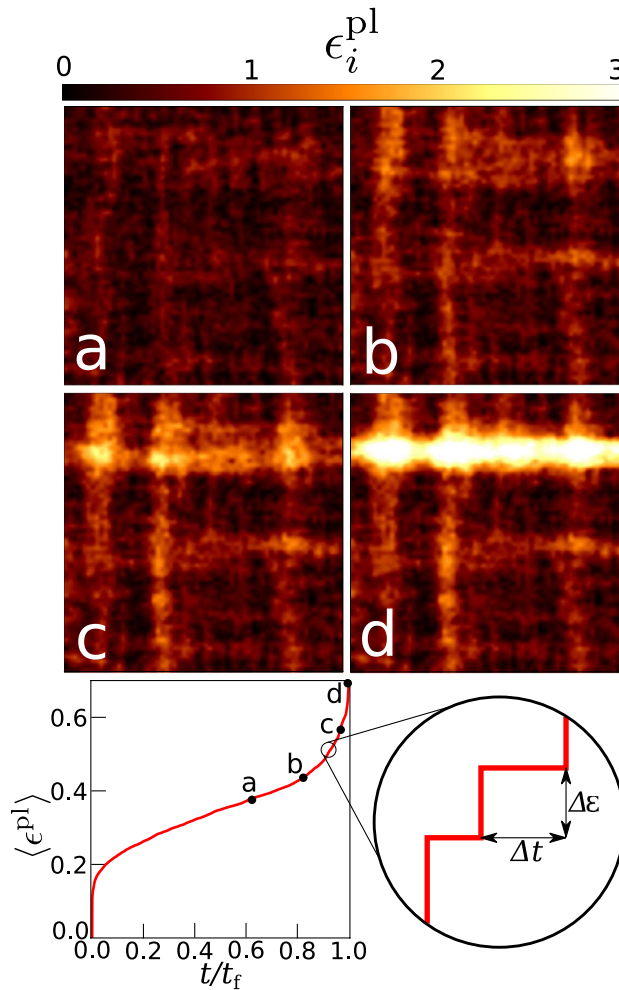


Figure 1. The top panel shows the spatial evolution of damage (cumulative number of local AE events) as time evolves ($t(a) < t(b) < t(c) < t(d)$). At later time, the damage becomes localized. The bottom panel shows the growth of the global AE event number with time. The qualitative signature of localization appears roughly above $0.7t_f$, which is already close to breakdown. The aim of this study is to make predictions substantially ahead of the manifestation of damage localization.

learning prediction as the statistical average over the test set $e_{ML} = \langle \frac{|t_p - t_a|}{t_f} \rangle$; the complementary quantity $1 - e_{ML}$ is denoted as the prediction performance.

In order to assess prediction capabilities, it is appropriate to quantify the performance of the algorithm not in absolute terms but relative to a baseline that can be obtained without involving any monitoring data or ML algorithms. We use as baseline for residual lifetime prediction at time t the mean residual lifetime of samples in the reduced training set S_t , consisting of all training samples with lifetime larger than t . The mean error (averaged over the test set) made by this baseline prediction for individual samples is denoted by e_{woML} . We then define the improvement achieved by machine learning over the baseline prediction as $\epsilon = e_{ML}/e_{woML}$. The complementary quantity

$$S = 1 - \frac{e_{ML}}{e_{woML}}. \tag{1}$$

is denoted as the prediction score. A prediction score of 1 means that there is no error, whereas a prediction score of zero indicates that the prediction is only as good as the baseline prediction that the individual sample lifetime equals the average lifetime of the samples in the training set. Note that, since both lifetimes and predictions are statistically distributed variables, even with a high prediction score there may exist individual samples for which the ML prediction is worse than the baseline.

We use the creep time series of 1000 samples as our training set and 200 different samples as the test set to evaluate the predictions of remaining time to failure. With the trained algorithm, we systematically investigate how the above prediction accuracy measures depend on externally applied stress, disorder and sample size, and we investigate how the prediction accuracy is influenced by the inclusion of spatial features into the monitoring data. Afterwards, we benchmark the machine learning predictions against other methods of prediction that use

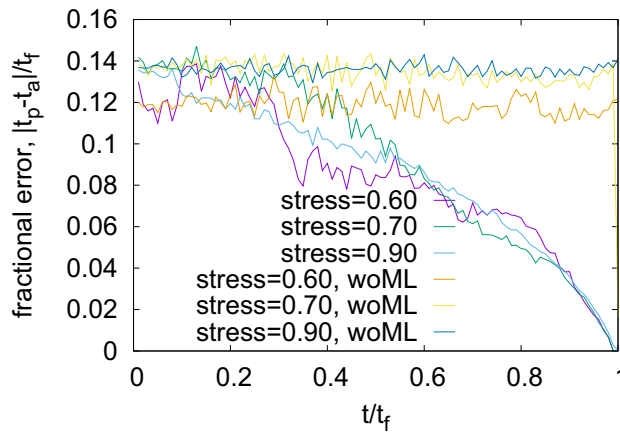


Figure 2. Fractional error in ML predictions at various stress levels and corresponding baseline prediction based upon average failure time, both predictions start from similar levels and the ML predictions improve gradually, while the baseline remains unchanged; there is no systematic dependency of fractional error on stress level even though the absolute lifetime changes dramatically; for parameters see “Methods” section.

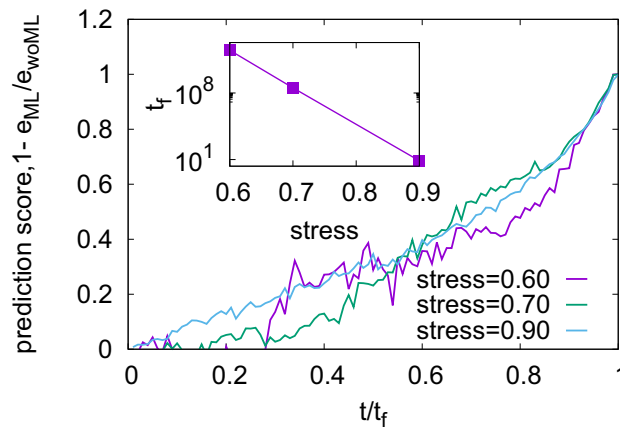


Figure 3. Prediction score as function of time to failure, for different applied stress levels; inset: mean lifetime as a function of stress level; see “Methods” for other model parameters.

empirical laws. Specifically, we use the time minimum of the strain rate t_m to predict the failure time t_f assuming a linear relationship between both, and the distinct Omori-type acceleration of the AE event rate in the approach to failure that has been reported for the model at hand².

Dependence of prediction performance on applied stress level. In order to reproduce creep conditions, the system is loaded with a constant external stress Σ^{ext} which is below the short-term critical stress Σ^c at which the system fails instantaneously. The ensuing failure time t_f depends strongly on the ratio $\Sigma^{\text{ext}}/\Sigma^c$. It is therefore natural to expect a variation in predictability as $\Sigma^{\text{ext}}/\Sigma^c \rightarrow 1$. We have used three values of applied stress—60%, 70% and 90% of the critical stress, respectively, keeping the other parameters of the model fixed.

Figure 2 shows the fractional errors e_{ML} and e_{woML} for different values of stress. Even though the absolute value of the sample lifetime changes by many orders of magnitude for this considerable range of variation in the applied load (see inset in Fig. 3), we find no significant or systematic dependency of the fractional error on stress. Note that for small values of t , the prediction from machine learning is just equal to the average of the training set. This is expected since, at the beginning of the creep dynamics for a particular sample, the algorithm has not yet received any sample specific information. As time progresses, the algorithm utilizes its training and makes, based on the precursor activity up to that time, predictions that improve with increasing length of the precursor record. On the other hand, in the absence of any ‘training’, the naive prediction from the average of training data does not improve with time and remains roughly constant.

Figure 3 shows the prediction score achieved by machine learning for different stress levels, as a function of time-to-failure. Note that the extreme increase of the damage rate just before failure ensures that failure is always correctly identified as it happens, with the consequence that for $t \rightarrow t_f$, $S \rightarrow 1$. The question is, however, whether the machine learning algorithm can achieve good prediction scores at earlier times.

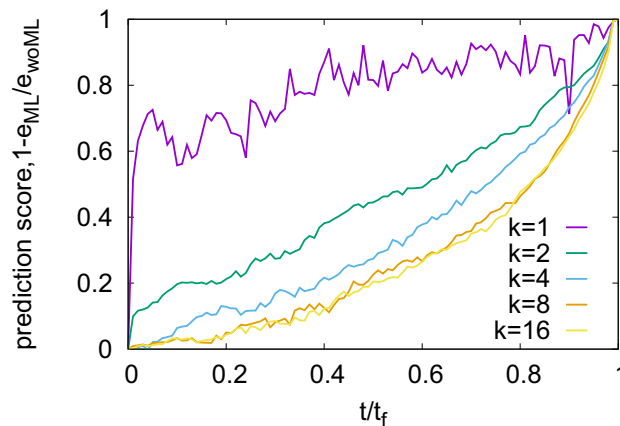


Figure 4. ML prediction scores for different degrees of disorder (Weibull exponents k); see “Methods” for model parameters.

Dependence on material disorder. In the model we use, the local barrier heights which control damage accumulation are statistically distributed to represent a material with a disordered microstructure. If one assumes the weakest-link hypothesis, then the local strength of a mesoscale region is essentially the strength of the weakest microscopic subregion. In this case, the mesoscale distribution function of the local strength is expected to follow a Weibull distribution³. We statistically distribute the local barriers according to a Weibull distribution with shape parameter k , which determines the width of the distribution and hence can be used to quantify the microstructural disorder. Specifically, a small value of k indicates a wide distribution and therefore a high degree of microstructural disorder. This translates into a comparatively large statistical scatter of sample lifetimes. Conversely, very large values of k imply nearly deterministic behavior, i.e., the creep curves of different samples and the corresponding sample lifetimes are almost identical.

Figure 4, left, shows time dependent prediction scores achieved for a range of values of k . High disorder substantially improves predictability of the remaining time to failure. The reason for the higher prediction scores lies in the more complex precursor activity and larger variations in local properties, resulting in spatio-temporal correlations which anticipate strain localization already at early creep stages well before catastrophic failure (see Ref. 20). In particular, significant local strain differences emerge already at an early stage in strongly disordered samples. Thus, spatial signatures can contribute to prediction at a much earlier stage, leading to better-trained algorithms for the same number of training samples. On the other hand, with decreasing disorder prediction scores become independent of the disorder. This can be understood by noticing that when the disorder distribution is very narrow, the stochastic behavior of the time series becomes dominated by thermal noise, which is kept constant. Hence the predictability becomes independent of the disorder of local strengths.

Dependence on sample size. The mesoscale model considers a square lattice composed of $L \times L$ mesoscale regions. Here we vary the linear system size L to study the impact of sample size on the accuracy of the machine learning predictions.

As can be seen from the plots in Fig. 5, the prediction scores decrease with system size. This can be understood as a consequence of strain localization. Sample failure is controlled by processes taking place in a localized shear band which emerges before failure. The width of this band does not depend on system size, hence it occupies a smaller fraction of the sample when the sample is larger. If one assumes that precursory signals that can be used for prediction mainly emanate from the shear band region (see Ref. 9 for a discussion of this phenomenon on a real sample), whereas other regions mainly produce confounding ‘noise’, then it is clear that smaller samples exhibit a better ratio of precursor signal to stochastic noise, and are therefore more predictable.

This observation might have far-reaching consequences in terms of real-world predictions. For example, a catastrophic shear band in a laboratory-scale fracture test of a rock sample occupies a far larger fraction of the overall sample volume than the slip localization zone in the context of an earthquake. Thus, sample size may be an important factor determining predictability, with unfortunate implications for the predictability of geo-scale fracture processes.

Comparison of ML with alternative prediction methods. It is useful to compare the machine learning predictions obtained here with results obtained by other methods. We consider two approaches that have been proposed and used in the literature.

First, we envisage an empirical method which uses the time t_m at which the global minimum strain rate occurs to predict the failure time t_f as a linear function of t_m ^{6,7}. The main drawback of this method lies in the fact that the average minimum is quite flat, whereas the instantaneous creep strain rate is subject to strong fluctuations. To apply this method, a smoothed signal must first be constructed from the discrete sequence of events (see “Methods”). The question is whether the best possible prediction from the empirical method based on the global minimum is better than the one from machine learning. Figure 6, left, shows such a comparison. As expected, the

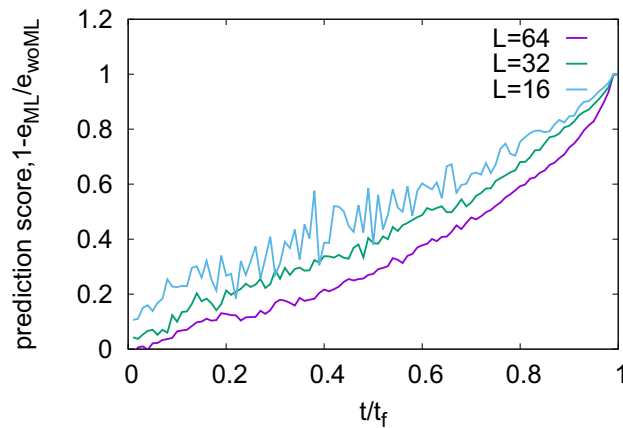


Figure 5. Dependency of ML prediction scores on system size; see “Methods” for model parameters.

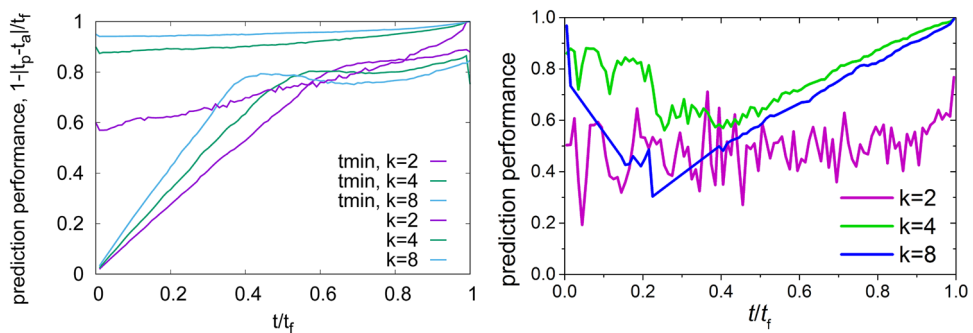


Figure 6. Left: comparison between prediction performance from machine learning and failure time prediction based upon minimum creep rate; right: performance of failure time prediction based upon Omori-type event rate acceleration; see “Methods” for model parameters and other details.

initial predictions from the strain rate minimum are, during the initial stages of creep before the actual minimum has been reached, far worse than the machine learning ones. The same is true during the late stages of creep close to the failure time, since predictions based on the strain rate minimum cease to improve once the minimum is passed, whereas ML predictions continuously improve. For moderate to low disorder ($k = 4 - 8$), the strain rate minimum method performs consistently worse than the machine learning method. Only for high disorder ($k = 2$) and creep times close to t_m it achieves prediction scores that are comparable to ML.

A method that is complementary to the strain rate minimum approach, in the sense that it yields predictions that continuously improve in the approach to failure, might be based on a reverse Omori-type acceleration of the event rate in the approach to failure. Such an acceleration is observed in experimental AE records and has been discussed as a means of prediction⁹, and it is also a generic feature of the creep model considered here². According to the reverse Omori law, the event rate increases in the approach to failure as $\dot{n} \propto (t_f - t)^{-1}$, and one can fit this relationship to the data to obtain the failure time as a fit parameter. Alternatively, one can integrate this relationship to express the event occurrence time as a function of the event number, $t(n) = t_f[1 - \exp(-an)]$ and again fit this relationship to the data. We choose the latter method since it avoids the need of averaging the strongly intermittent and fluctuating event rate. Results are compiled in Fig. 6, right. It is seen that Omori-type predictions indeed improve in the approach to failure, however, they consistently perform well below the machine prediction. Except in the immediate vicinity of failure, predictions based on reverse Omori fits actually perform below the baseline defined by the ensemble averaged lifetime. The reason for this finding is that reliable fits of an Omori type acceleration are a priori impossible until one is well beyond the strain rate minimum. Also, while the reverse Omori law is very clearly borne out when one considers ensemble-averaged statistics with a large ensemble of samples, it is much less evident in single samples, in particular when the samples are small or strongly disordered.

It is interesting to compare the results obtained from an Omori-based prediction strategy with those obtained from a ‘temporal’ ML strategy where we only use avalanche times and magnitudes but not spatial information. The performance of this ‘temporal’ ML is in fact very similar to that of the Omori type prediction scheme: over most of the sample lifetime, the performance is below the baseline and only for ordered samples and in the very last stage of deformation, prediction performance raises above the baseline, while always remaining well below that of the standard ML scheme. We may thus infer that a purely temporal ML scheme essentially probes the

reverse Omori law (and performs equally badly), whereas the performance of our actual ML scheme, and notably its capability of predicting strongly disordered samples, is contingent on the use of use of spatial information—the high early prediction scores in disordered samples are due to an early emergence of spatial localization features.

Discussion and outlook

Prediction of failure time for creep rupture is a crucial problem with wide-ranging potential applications in science and engineering. Empirical prediction methods are often not accurate enough, especially when the disorder is strong. Moreover, sample to sample variations are often high which makes it difficult to extrapolate knowledge gathered from a tested sample to another one. For this reason, we have trained a machine learning algorithm. For machine learning, fluctuations become a source of knowledge that can help in training the algorithm to better recognize precursor patterns to failure and exploit complex correlations which can be used to predict incipient failure.

We have performed a systematic study of the variations of predictability with applied stress, presence of disorder and sample size, using synthetic data generated from a well-established model of creep deformation and failure. We find no systematic variation of predictability with stress. However, predictability increases with increasing disorder while it decreases with increasing sample size. Our benchmark against alternative methods confirms a superiority of machine learning over other approaches suggested in the literature, which can be regarded as a promising method with the potential to improve existing hazard assessment techniques.

The proposed method is based on the use of a comparatively small set of features characterizing the deformation process. These features are chosen in such a manner that they have direct equivalents in laboratory tests, such as the timings and intensities of acoustic emissions or simple signatures of progressive spatial localization. Advances in acoustic emission tests have made possible in recent years to measure AE magnitudes and locations with high resolution, which makes the proposed machine learning method an ideal candidate for analyzing and extracting useful information from experimental data. The method works even if comparatively small numbers of sample records (here: several hundreds) are available for testing. Such ensemble tests have been conducted in the published literature and it will be an interesting task for future investigations to apply the present method to actual monitoring data on quasi-brittle disordered materials. In this context quasi-two-dimensional samples such as paper^{7,8} or semi-brittle polymer sheets, which approximately match the geometry considered in our simulations, are of particular appeal. Since experimental ensembles are in general much smaller than simulated ones, an important question arises in this context: Can one transfer training results from simulation to experimental data? Whether or not this can be effected by appropriate re-scaling depends on whether the simulations correctly represent the essential features of the processes occurring in the experimental samples. This implies, on the one hand, that it is wise to ‘keep it simple’ and focus on basic features of the monitoring records. Also, parameters that strongly influence the prediction scores should be the same in simulations and experiments. This concerns in particular the ‘disorder strength’ as expressed by the k -exponent of the Weibull distribution, and indeed the question whether Weibull distributed thresholds adequately represent the threshold disorder of real samples. In this context, one might use an independent machine-learning method proposed in Ref. 21 to infer local threshold distributions for matching simulations from non universal features of the experimental avalanche statistics.

Methods

Creep model. Synthetic space-time series of creep deformation accompanied by damage accumulation are produced by a mesoscale model of plastic deformation of disordered materials introduced in Refs. 2,3. The model considers a 2D $L \times L$ lattice of mesoscale elements denoted by an index $i \in [1 \dots L^2]$. Each mesoscale element has a volume V which coarse-grains microscopic details of a disordered material. Similar to Ref. 22, we describe the state of each mesoscale element by continuum mechanics variables, namely a tensorial irreversible strain ϵ_i^{pl} and a stress tensor Σ_i which is connected to the reversible (elastic) strain tensor via the tensor of elastic constants, which we assume to represent an isotropic material. The internal microstructure of each element is characterized by a spectrum of stress dependent energy barriers of which we assume the lowest barrier, $\Delta E_{\text{min},i}(\Sigma_i)$ to control activation of irreversible deformation. To make the connection with traditional concepts of mechanics of materials, we introduce element specific, stress dependent yield functions $\Phi_i(\Sigma_i)$ which fulfil the condition

$$\Delta E_{\text{min},i}(\Sigma_i) = 0 \quad \text{if} \quad \Phi_i(\Sigma_i) = 0 \quad (2)$$

We assume that deformation is controlled by deviatoric (shear) stress only and take Φ_i to be of the form

$$\Phi_i = \Sigma_i^{\text{eq}} - \hat{\Sigma}_i \quad (3)$$

where $\Sigma_i^{\text{eq}} = \sqrt{(3/2)\text{dev}(\Sigma) : \text{dev}(\Sigma)}$ is the von Mises equivalent stress and $\text{dev}(\Sigma)$ denotes the deviatoric stress tensor. $\hat{\Sigma}_i$ defines the equivalent stress at which $\Phi_i = 0$, i.e., the stress at which the energy barrier to initiate a plastic strain increment vanishes and, hence, the volume element i becomes mechanically unstable. In the language of plasticity theory, this corresponds to the local flow stress in the limit of zero temperature.

In the regime of negative Φ_i , plastic deformation is controlled by thermal activation over non vanishing barriers. For simplicity, we assume the rate controlling energy barrier to be linearly proportional to Φ_i , i.e.,

$$\Delta E_{\text{min},i}(\Phi_i) = \begin{cases} -V_a \Phi_i, & \Phi_i < 0, \\ 0, & \Phi_i \geq 0. \end{cases} \quad (4)$$

where the proportionality constant V_a is an activation volume. Barrier crossing leads to a discrete plastic event which introduces a finite tensorial plastic strain increment $\Delta \epsilon^{\text{pl}}$. The barrier crossing rate in element

i is $v_i = v_{el} \exp(-\Delta E_{min,i}/k_B T) = v_{el} \exp(\Phi_i/\Sigma_T)$ where the parameter v_{el} defines the local yielding attempt frequency at the mesoscale and v_{el}^{-1} defines the natural timescale of the model. $\Sigma_T = k_B T/V_a$ characterizes the influence of temperature in terms of a scalar, stress-like variable. In mechanically unstable elements ($\Phi \geq 0$) events are assumed to occur instantaneously. Upon activation of an event in an element, the plastic strain field in that element is updated by adding to the local irreversible strain the tensorial increment $\Delta \epsilon_i^{pl} = \Delta \epsilon_i^{eq} \cdot \hat{\epsilon}_i$, where $\Delta \epsilon_i^{eq}$ denotes the scalar magnitude of the strain increment and the tensor $\hat{\epsilon}_i$ defines the direction. In line with J2 plasticity, this direction is given by a maximum energy dissipation criterion, thus $\hat{\epsilon} = \nabla_{\Sigma} \Phi = (3/2) \text{dev}(\Sigma)/\Sigma^{eq}$. On the other hand, the magnitude is given by $\Delta \epsilon_i^{eq} = \chi \Sigma^{eq}/3G$, where χ is a factor between 0 and 1. This choice ensures that the local deviatoric stress (the thermodynamic driving force) cannot change sign upon introduction of an event. The location i of a thermally activated deformation event and the associated time increment $\Delta t > 0$ elapsed since the last thermal activation are in our simulations determined by the Kinetic Monte Carlo method. Upon introduction of an event in element i , we increase the plastic strain tensor in that element, $\epsilon_i^{pl} \rightarrow \epsilon_i^{pl} + \Delta \epsilon_i^{pl}$. Alongside with the plastic strain tensor, we also update the cumulative equivalent strain, $\epsilon_i^{eq} \rightarrow \epsilon_i^{eq} + \Delta \epsilon_i^{eq}$. Using the updated plastic strain field, stresses everywhere in the simulated sample are re-computed. The ensuing stress changes may lead to destabilization of other elements and thus to secondary events. In that case, stresses are again updated considering all such plastic events simultaneously, then checking for further unstable events, and continuing this cycle until the system is mechanically stable and the 'avalanche' terminates. The simulation then returns the following primary data: (i) the time of the thermally activated event, (ii) the overall strain increment, (iii) the change in the spatial strain pattern.

To account for microstructural randomness, the element strength $\hat{\Sigma}_i$ is statistically distributed according to a Weibull distribution of exponent k with cumulative distribution function

$$P(\hat{\Sigma}_i) = 1 - \exp\left(-\left[\frac{\hat{\Sigma}_i}{\bar{\Sigma}_i}\right]^k\right) \quad (5)$$

and mean value $\langle \hat{\Sigma}_i \rangle = \bar{\Sigma}_i \Gamma(1 + 1/k)$ where Γ denotes the Gamma function. Whenever an element undergoes plastic deformation, its strength is renewed from the distribution (5). Deformation-induced damage in element i is described by a variable $\delta_i = f \epsilon_i^{eq}$ which is proportional to the local equivalent plastic strain. The average of the distribution from which local strength values are drawn decreases with local damage as $\bar{\Sigma}_i = \bar{\Sigma}_0 \exp(-\delta_i)$, implementing strain softening.

We load the system under pure shear conditions, with principal axes oriented along $\pm\pi/4$ to the x axis of our Cartesian coordinate system. This gives rise to a spatially homogeneous 'external' stress tensor which represents a pure shear stress state, $\Sigma = \Sigma^{ext}(\mathbf{e}_x \otimes \mathbf{e}_y + \mathbf{e}_y \otimes \mathbf{e}_x)$. In the course of creep deformation, the emerging inhomogeneous plastic strain pattern leads to inhomogeneous and multi-axial internal stresses which add to this external stress. The calculation of these stresses is done by the Finite Element Method with a regular square grid, linear shape functions, assuming linear elasticity and with homogeneous elastic properties. Each mesoscale element is matched with a finite element. The stress at each mesoscale element is computed as the average of the stress field within the associated FEM element. The reader is referred to²³ for further details on the numerical implementation of the model.

To perform simulations under creep loading conditions, the value of Σ^{ext} is kept fixed in time. To establish the specific value, we look first for the critical value $\Sigma^{ext} = \Sigma^c$ beyond which the system is mechanically unstable and fails instantaneously even at zero temperature. This parameter defines our stress scale. We measure stresses in units of Σ^c , strains in units of Σ^c/E where E is the Young's modulus of the material, and time in units of v_{el}^{-1} . Externally controllable parameters are Σ^{ext} and Σ_T . Default parameters in our simulations are, unless otherwise stated, $L = 64$, $\Sigma_T = 0.015$, $\chi = 0.1$, $f = 0.1$, $k = 4$ and $\Sigma^{ext} = 0.7 \Sigma^c$.

The model produces, as output, raw data in the form of the times, locations, and magnitudes (strain increments) of all local deformation events between initial loading and failure. We can envisage this output as a simulated acoustic emission record which monitors the deformation activity within the sample throughout the creep process.

Machine learning method. For predicting the failure time, we use a supervised learning algorithm—Random Forest regression¹⁹—as implemented in the Scikit-learn Python library²⁴. The algorithm is trained over a training set where we use various features of the creep simulation data. Predictions are made at the times where an avalanche is ended. For making a prediction at a given point of time (necessarily at the end of an avalanche) we use the data features fed to the algorithm at that point of time, which include both temporally and spatially aggregated information up to that point. Specifically, the features used for prediction include (1) the elapsed time since the beginning of the creep process up until the point where a prediction is being made (again, necessarily at the end of an avalanche), (2) the size of the last avalanche at the end of which the prediction is being made. We also add spatially aggregated information such as the (3) maximum and (4) minimum damage magnitudes (accumulated local AE events as shown in Fig. 1a–d). Finally, the last column of the training data set is the target variable, i.e. time remaining before macroscopic failure, which we want to predict (in the test data set).

Attributes (3) and (4) are calculated by taking the sum over the individual rows and columns of the matrices shown in Fig. 1a–d, and determining the *magnitude* of the maximum and minimum accumulated damage along a row or column, i.e. if d_{ij} represents the accumulated AE damage matrix, and we let $D_{max}^{row} = \max(\sum_j d_{ij}, j = 1, 2, \dots, L)$ and $D_{max}^{col} = \max(\sum_i d_{ij}, i = 1, 2, \dots, L)$, then the third attribute is simply $\max(D_{max}^{row}, D_{max}^{col})$. The fourth feature is just the corresponding minimum. At this point a comment is needed: From continuum mechanical stability analysis we know that, given the nature of the applied loading (the direction of the applied surface tractions), an incipient shear band will be oriented parallel to either the x or y axis of our

coordinate system, and not diagonal or along any other angle. Hence the rows and columns of the accumulated damage matrix (d_{ij}) represent potential shear band 'candidates'. The absolute maximum of the summed-up damage along a row or column can be thought of as representing the damage in the center of an incipient shear band, whereas the difference between the maximum and the corresponding absolute minimum of the summed-up damage gives an indication of the degree of deformation localization.

For every prediction set, a typical training data set consists of the above mentioned features of 1000 samples and the test data set is typically comprised of data of 200 samples (over which the prediction accuracy is measured). The training set data file is the combination of all 1000 samples time series features appended together. Due to the slow time variation and the size of file, one in 50 time steps are considered for the training set to have a significant change in the feature values. Further details about the data processing and parameter sets for the regression model are given below.

General features of the random forest regression model. Random Forest (RF) regression is an ensemble algorithm that makes predictions based on the average prediction of an ensemble of decision trees. A decision tree is a flow-chart like structure, where starting from a root node, the samples are split depending on their feature values or attributes. For example, a particular attribute A could be used to split the samples into two parts, those having values less than A_0 and those having values greater than A_0 . Each of these parts can be further split depending on other attribute values and so on. The splitting values of the attributes at each stage are optimized by the algorithm used until all samples at a given node have the same value of the target variable (in this case, time to failure), a prescribed maximum depth of the tree (number of splittings) is reached, or further splitting does not improve predictions. To assess the latter point we use a variance reduction criterion which works as follows: If N is the number of samples in a node, then the node is split (provided it is not restricted by maximum depth of the tree) for a threshold value, say A_0 of a feature A , provided A and A_0 maximize Δ , which is given by $\Delta = var(N) - \frac{n_l}{N} var(n_l) - \frac{n_r}{N} var(n_r)$, where n_l and n_r are the numbers of samples in the two nodes if the split is accepted and $var(\dots)$ refers to the variance of the target variable i.e. the prediction time of the set of samples. This is checked for all features and all threshold values for all splitting decisions. The end nodes are called leafs and they hold the predictions for the given set. For each of the trees, the training data are subsampled using a Bootstrapping algorithm (see below). Consequently, each tree is fed with a random subset of the training data (hence Random Forest). Following the training, the test data, which is unseen by the model until this point, are passed through each tree and they end up in the leaf nodes which are then the predictions for each of the test data points. In case of regression, where the target variable is continuous such as here, the prediction of the RF for a given test data point is the average value of the predictions of each of the trees for that data point.

Data processing for regression. From the training data set, a number of subsampled data sets are generated by randomly selecting data points from the training set (selecting, say, N rows randomly and uniformly from N rows, but with replacement, i.e. bootstrapping). Due to bootstrapping, some of the data points will be repeated, which acts as mitigation towards outliers in the training set. The number of subsampled, randomly generated sets is equal to the number of decision trees used in the RF (see below). Each of the trees are then fed with a different training set (randomly sampled) and in the case of regression, as in our case, the average prediction of all the trees is the prediction of the forest.

Parameters of the model. As mentioned above, the RF algorithm uses a set of decision trees for making predictions. The number of decision trees used here is 1000, however, a parametric study showed no appreciable effect on prediction scores when this number was varied between 200 and 1200. The algorithm uses a specified maximum depth (number of splittings) for each tree, set at 10. Higher depth was found to lead to an increase in error due to overfitting. The minimum samples required to split an internal node is 2 and the minimum number of samples to be in a leaf node is 1. Remaining parameters are default for the Scikit-learn 0.19.1 version.

Analysis based on strain rate minimum. In order to estimate the failure time based on the time of minimum strain rate, t_m , which separates the decelerating and accelerating creep regimes, we need to reconstruct a smooth signal from the discrete sequence of events. First, we note that the strain rate minimum occurs during the linear creep regime and that during this regime plastic activity is almost exclusively thermally activated, with subsequent mechanically activated events being rare. In this case, avalanches have a typical size of a single plastic event and the strain increment measured over a certain observation interval is proportional to the number of avalanches occurring in that interval. Consequently, we can estimate the minimum of strain rate by looking for the maximum of a smoothed time increment signal. To obtain such smoothed signal, we substitute the n th value of the discrete time increments Δt_n by the average of the increments whose numbers lie in the window $[n - h, n + h]$ of width $2h$ centered at n . Averaging over a window of fixed width defined in terms of event number can be interpreted as averaging with an adaptive time window (i.e., a narrow time window in stages of small characteristic time increments and vice versa). The value h of the window width must be set arbitrarily. We check the stability of the results upon variations of h in order to decide its specific value. To this end, we compute the probability distribution of t_m , where each t_m corresponds to a different realization of the creep process. We find that for a wide range of values $h \in [200, 2000]$ the results are nearly independent of h for all the different values of the simulation parameters considered in this work, and we set $h = 500$.

Analysis based on Omori-type event rate acceleration. For the present creep model, the ensemble-averaged event rate shows an Omori-type acceleration to failure, $\dot{n} = a(t_f - t)^{-1}$. Since the event rate is a strongly fluctuating quantity, we integrate this relationship to obtain

$$n(t) = a \ln \left(\frac{t_f}{t_f - t} \right), \quad t(n) = t_f (1 - \exp(-n/a)). \quad (6)$$

We may fit this relationship to the actual event times in order to obtain t_f as a fit parameter. However, a straightforward fit may not work since any data before the strain-rate minimum cannot be represented by Eq. (6) and the resulting errors may compromise the fit result. We resolve this problem by observing that, for the data in our simulations, also the *decelerating* part of the creep curve shows a (reverse) Omori law³, hence $\dot{n} = a[(t_f - t)^{-1} + t^{-1}]$ which gives

$$n(t) = a \left[\ln \left(\frac{t_f - t_0}{t_f - t} \right) + \ln \left(\frac{t}{t_0} \right) \right], \quad t(n) = t_f (1 + \exp(-n/a)(t_f/t_0 - 1))^{-1}. \quad (7)$$

We fit this relation to the event time list and use the t_f fit parameter as our estimate for the sample specific lifetime.

Received: 5 May 2020; Accepted: 27 July 2020

Published online: 09 October 2020

References

- Andrade, E. N. D. C. On the viscous flow in metals, and allied phenomena. *Phil. Trans. R. Soc. A* **84**, 1 (1910).
- Castellanos, D. F. & Zaiser, M. Avalanche behavior in creep failure of disordered materials. *Phys. Rev. Lett.* **121**, 125501 (2018).
- Castellanos, D. F. & Zaiser, M. Statistical dynamics of early creep stages in disordered materials. *Eur. Phys. J. B* **92**, 139 (2019).
- Zaiser, M., Moretti, P., & Chu, H. J. Stochastic crystal plasticity models with internal variables: application to slip channel formation in irradiated metals. *Adv. Eng. Mater.* <https://doi.org/10.1002/adem.201901208>.
- Biswas, S., Ray, P. & Chakrabarti, B. K. *Statistical Physics of Fracture, Breakdown, and Earthquakes: Effects of Disorder and Heterogeneity* (Wiley, Hoboken, 2015).
- Hao, S.-W., Zhang, B.-J., Tian, J.-F. & Elsworth, D. Predicting time-to-failure in rock extrapolated from secondary creep. *J. Geophys. Res.* **119**, 1942 (2014).
- Koivisto, J., Ovaska, M., Miksic, A., Laurson, L. & Alava, M. J. Predicting sample lifetimes in creep fracture of heterogeneous materials. *Phys. Rev. E* **94**, 023002 (2016).
- Viitanen, L., Ovaska, M., Ram, S. K., Alava, M. J. & Karppinen, P. Predicting creep failure from cracks in a heterogeneous material using acoustic emission and speckle imaging. *Phys. Rev. Appl.* **11**, 024014 (2019).
- Lennartz-Sassinek, S., Main, I. G., Zaiser, M. & Graham, C. C. Acceleration and localization of subcritical crack growth in a natural composite material. *Phys. Rev. E* **90**, 052401 (2014).
- Saichev, A. & Sornette, D. Andrade, Omori, and time-to-failure laws from thermal noise in material rupture. *Phys. Rev. E* **71**, 016608 (2005).
- Kun, F., Varga, I., Lennartz-Sassinek, S. & Main, I. G. Approach to failure in porous granular materials under compression. *Phys. Rev. E* **88**, 062207 (2013).
- Cubuk, E. D. *et al.* Identifying structural flow defects in disordered solids using machine-learning methods. *Phys. Rev. Lett.* **114**, 108001 (2015).
- Bapst, V. *et al.* Unveiling the predictive power of static structure in glassy systems. *Nat. Phys.* **16**, 448–454 (2020).
- Papanikolaou, S. Microstructural inelastic fingerprints and data-rich predictions of plasticity and damage in solids. *Comp. Mech.* **1–14** (2020).
- Liu, W. K., Karniadakis, G., Tang, S. & Yvonnet, J. A computational mechanics special issue on: data-driven modeling and simulation-theory, methods, and applications. *Comp. Mech.* **64**, 275–277 (2019).
- Rouet-Leduc, B. *et al.* Machine learning predicts laboratory earthquakes. *Geophys. Res. Lett.* **44**, 9276 (2017).
- Salmenjoki, H., Alava, M. J. & Laurson, L. Machine learning plastic deformation of crystals. *Nat. Commun.* **9**, 5307 (2018).
- van der Baan, M. & Jutten, C. Neural networks in geophysical applications. *Geophysics* **65**, 1032 (2000).
- Hastie, T., Tibshirani, R. & Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference and Predictions* (Springer, New York, 2001).
- Vasseur, J. *et al.* Heterogeneity: the key to failure forecasting. *Sci. Rep.* **5**, 13259 (2015).
- Papanikolaou, S. Learning local, quenched disorder in plasticity and other crackling noise phenomena. *NPJ Comput. Mater.* **4**, 27 (2018).
- Budrikis, Z., Castellanos, D. F., Sandfeld, S., Zaiser, M. & Zapperi, S. Universal features of amorphous plasticity. *Nat. Commun.* **8**, 15928 (2017).
- Castellanos, D. F. *Stochastic modeling of plastic flow and failure in disordered materials*, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), (2019).
- Pedregosa, F. *et al.* Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).

Acknowledgements

The authors acknowledge support by DFG under Grant No. 1Za 171/8-1 and under 377472739/GRK 2423/1-2019 FRASCAL. Support by the European commission under H2020-MSCA-RISE Project No. 734485 FRAMED is also gratefully acknowledged.

Author contributions

D.F.C. wrote the creep code and performed creep simulations, S.B. analyzed the obtained data using random forest regression and strain rate minimum based forecasting, M.Z. designed the overall research and provided failure forecasts based on Omori acceleration. All authors contributed to the manuscript and reviewed the final version.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.Z.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020