



Predicting the Timing of Catastrophic Failure During Triaxial Compression: Insights from Discrete Element Method Simulations

JESSICA McBECK¹

Abstract—Examining the deformation of rocks during triaxial compression may provide insights into the precursory phase that leads to large earthquakes by revealing the components of the deformation field that evolve with predictable, systematic behavior preceding catastrophic failure. Here, we build three-dimensional discrete element method simulations of the triaxial compression of rock cores that include a variety of fault geometries in order to identify the components of the deformation field, including the velocity and strain components, that enable machine learning models to predict the timing of macroscopic failure. The results suggest that the velocity field provides more valuable information about the timing of macroscopic failure than the strain field, and in particular, the velocity component parallel to the maximum compression direction. The models also strongly depend on the second invariant of the strain deviator tensor, J_2 , indicative of the shear strain. The importance of J_2 on the model predictions increases with confining stress, consistent with laboratory observations that show a transition from tensile- to shear-dominated deformation with increasing confining stress. In contrast to expectations from previous machine learning analyses, none of the models strongly depend on components of the strain tensor indicative of dilation, such as the first invariant of the strain tensor. This difference may arise because the simulations host less dilative strain than the experimental rocks.

Keywords: Triaxial compression, macroscopic failure, fracture development, machine learning, discrete element method.

1. Introduction

Identifying the factors that control macroscopic failure and how they evolve toward failure are critical goals in geoscience and engineering. Determining which strain components, for example, develop with a systematic evolution towards macroscopic failure

may enable analysts to recognize the precursory signals that indicate approaching macroscopic failure in the crust and engineered structures. Precursory activity has been detected preceding some large earthquakes (e.g., Whitcomb et al., 1973). However, for many large earthquakes, including earthquakes that occurred on faults that were heavily instrumented, analyses detected scant evidence of precursors (e.g., Bakun et al., 2005). Variations in seismic wave properties, including the velocity and attenuation, are one of the identified precursory changes in geophysical properties documented in laboratory experiments and some natural earthquakes (e.g., Main et al., 2012). Previous work attributes this variation to the dilation of the crust via fracture propagation and/or opening, preslip on existing fault surfaces, increasing normal stress in the wall rock surrounding the fault, and the movement of fluids in the crust (Main et al., 2012; Scholz et al., 1973). Using the evolution of seismic wave properties, previous machine learning analyses have successfully predicted the timing of slip within pre-cut blocks in the laboratory (Shreedharan et al., 2021a).

The observed reduction in the P-wave velocity suggests that dilation occurs preceding fault slip in the laboratory and the crust, and thus that the dilation of the crust may indicate the proximity of macroscopic failure. However, the evolutions of these seismic properties cannot reveal whether the shear strain accumulates with a recognizable evolution that indicates approaching macroscopic failure. Thus, tracking the seismic properties cannot provide insight into which strain component provides more systematic information, and thus potentially more predictive ability.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s00024-022-03151-4>.

¹ Njord Centre, Departments of Geosciences and Physics, University of Oslo, Oslo, Norway. E-mail: j.a.mcbeck@geo.uio.no

To directly compare the predictive ability of the dilation and shear strain, in a previous analysis we used the local incremental strain fields calculated from in situ X-ray tomograms acquired during triaxial compression to predict the timing of macroscopic failure, when the rock experiences a sudden, catastrophic stress drop after achieving its maximum differential stress (McBeck et al., 2020a). This work found that the machine learning models depend primarily on the local dilation of the rock to predict the timing of failure, rather than a quantity that measures the local shear strain components, the curl of the displacement field. This result is consistent with observations from the laboratory and the crust of temporary variations in the P-wave properties prior to fault slip and large earthquakes (e.g., Aggarwal et al., 1973; Chiarabba et al., 2020; Niu et al., 2008; Pio Lucente et al., 2010; Scuderi et al., 2016; Whitcomb et al., 1973). However, we performed this machine learning analysis on experiments with intact rocks that lack a preexisting fault, unlike the laboratory and crustal observations. Because faults can heal and thus strengthen in the interseismic period (e.g., Tenthorey & Cox, 2006), the nucleation process leading to large earthquakes may include fracture propagation and opening within relatively intact rocks, similar to the triaxial compression experiments of intact rocks. However, recent analyses highlight that the variation in seismic properties observed in the crust arises from a combination of processes within the host rock and on the preexisting fault surface that are integrated into the overall signal (Shreedharan et al., 2021b). In laboratory experiments with two sliding blocks, fault preslip tends to reduce the P-wave velocity, while increasing normal stress in the wall rock can increase the P-wave velocity, thereby reducing the strength of any observed reduction of the P-wave velocity (Shreedharan et al., 2021b). Thus, systems that include a preexisting healed fault surface may more closely represent the systems in which earthquakes nucleate in the crust than intact rocks. Moreover, the relative importance of the strain components in predicting the timing of macroscopic failure may differ in systems that include and exclude a preexisting fault. The inclusion of a preexisting healed fault with an orientation that maximizes the Coulomb shear

stress will likely promote shear deformation rather than dilation.

To constrain the characteristics of the deformation field that evolve in a systematic evolution toward failure, and thus may be used to recognize the precursory phase of approaching catastrophic failure, we build discrete element method simulations of the triaxial compression of nominally intact rocks, and rock cores that include a preexisting, healed fault. We develop machine learning models to predict the timing of macroscopic failure using characteristics of the deformation field, including the components of the three-dimensional velocity and strain fields. This analysis thus enables comparing the predictive power of each velocity and strain component to each other in systems that more closely resemble natural crustal fault systems than the previous machine learning analysis that used intact rocks (McBeck et al., 2020a). Performing this analysis on simulations, rather than rocks deformed in the laboratory, enables testing a wider range of confining stresses and fault geometries than feasible in the laboratory. We systematically vary the confining stress from uniaxial compression conditions to 50 MPa, vary the orientation of the preexisting fault from 15° to 60° from the maximum compression direction, and vary the fault length from system-spanning to shorter lengths. We test the influence of confining stress on the predictive power of the deformation components because laboratory observations document how confining stress controls the fracture geometry, and proportion of shear deformation relative to dilative or contractive deformation (e.g., Paterson & Wong, 2005). Similarly, laboratory experiments show that the orientation of precut faults can control how the rock fails, and in particular, whether the deformation is concentrated along the preexisting fault surface, or spread more diffusely in the surrounding rock volume (e.g., Giorgetti et al., 2019). We vary the fault length in the simulations because whether the fault spans the rock core, or only crosses a portion of the system, likely controls the precise geometry of the fracture network (e.g., Bobet & Einstein, 1998).

The results suggest that the components of the velocity fields are more valuable for predicting the timing of macroscopic failure than the components of the strain fields, contrary to expectations. In

particular, the deformation parallel to the maximum compression direction provides the most useful information to the machine learning models. Some of the machine learning models also depend on the second invariant of the strain tensor, J_2 . Increasing the confining stress increases the dependence of the models on J_2 . None of the models strongly depends on the components of the strain tensor indicative of dilation, the first invariant of the strain tensor, I_1 , even when the models only have access to the normal strain tensor components and I_1 , in contrast to results from previous machine learning analyses. This work thus examines the influence of fault geometry and confining stress on the predictive power of each component of the deformation field, including both the velocity and strain components.

2. Methods

2.1. ESyS-Particle Discrete Element Method Simulations

We use the discrete element method tool ESyS-Particle to build the simulations (Abe et al., 2004). ESyS-Particle approximates the behavior of materials with indestructible spherical particles that may be unbonded or connected to other particles with one or several bonds. The bonds have both shear and normal stiffness. During a simulation, a bond may break if the forces acting on it exceed the Coulomb criterion determined from its prescribed microparameters of cohesion and friction. Unbonded particles experience frictional and elastic forces. Thorough calibration of the microparameters of the bonds and particles is required for the simulations to approximate the mechanical response of natural rock. ESyS-Particle requires a damping parameter that decreases the kinetic energy in order to maintain the stability of the simulations. Thus, deformation and fault growth are quasistatic, similar to the laboratory experiments to which we compare the results of the present analysis (McBeck et al., 2020a).

We simulate the uniaxial and triaxial compression of blocks that include one rough fault embedded in a three-dimensional system. The blocks are 40 mm \times 80 mm \times 10 mm along the x , y , and z -axes

respectively (Figs. 1, 2). We chose to restrict the dimension of the block along the z -axis in order to improve processing time. Because the preexisting faults extend from the front to the back of the system, along the z -axis, varying the length of this dimension would not significantly change the local stress field that develops within the system, and thus the resulting fracture network.

We apply a constant normal stress to the four sides of the simulations on the front, back, left and right sides, and then apply constant axial velocity loading. To ensure stability, we apply a linear ramp from zero to the desired normal stress and axial velocity over 40,000 timesteps. We identify the timestep size using the maximum bond stiffness and minimum particle radius, following the Courant condition (e.g., Karimabadi et al., 2005) (e.g., Table S1). The normal stress increases in the first 40,000 timesteps, and then the axial velocity loading increases from 50,000 to 90,000 timesteps. Thus, the simulation achieves the desired confining stress before the onset of the axial velocity loading. To apply the axial velocity loading, we move the top and bottom boundaries of the simulations toward each other at 0.125 mm/s. We apply 0, 10, and 50 MPa of confining stress.

To build the preexisting fault, the simulations have three groups of bonds with different microparameters (cohesion and stiffness) that represent the principal slip zone, the damaged rock surrounding this zone, and the more intact host rock (Fig. 1). We refer to these bonds, and the types of fractures that develop within these materials via bond breakage, as on-fault, damage zone, and off-fault, respectively. The simulations include one fault that may be three different lengths, L : 15 mm long, 20 mm long, or system-spanning, in which the fault extends across the x -axis of the simulation (Fig. 2). These faults have roughness parameters of natural faults, with Hurst exponents parallel and perpendicular to the down-dip direction of 0.6, and an amplitude of one millimeter (Candela et al., 2012). Following the approach of McBeck et al. (2021b), we build the faults by labelling particles with surfaces that lie within one millimeter of a mesh surface that represents the fault as on-fault, and then labelling all the other particles as off-fault. Because the maximum

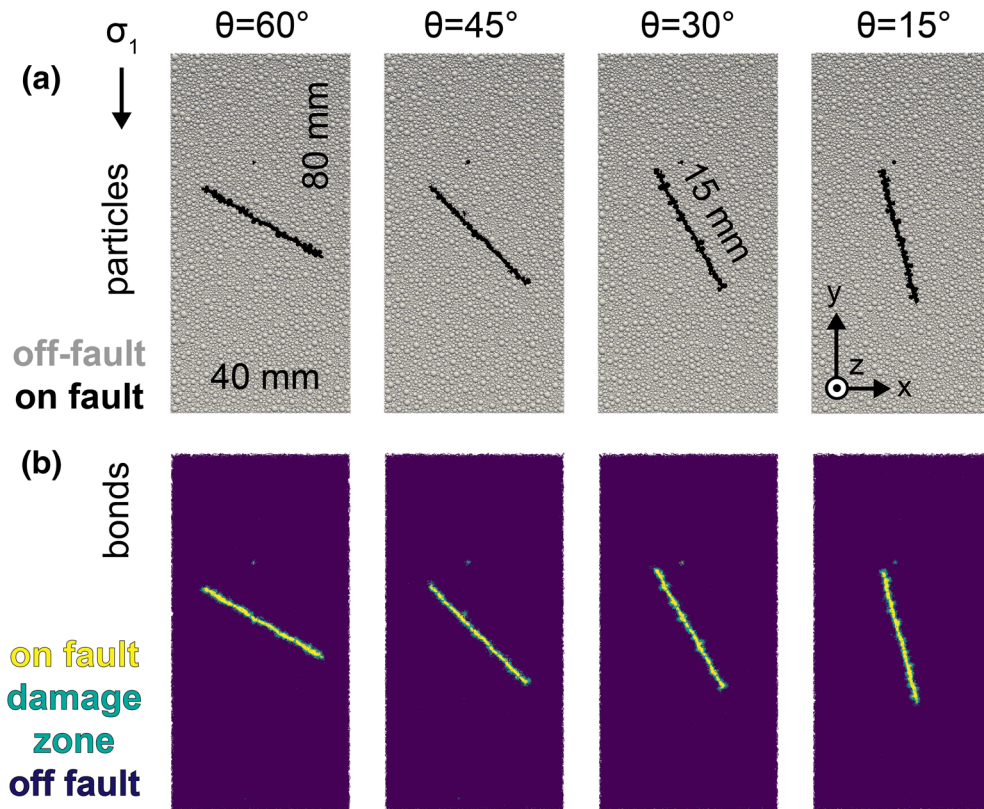


Figure 1

Geometries with 15 mm long faults and the four tested fault orientations: **a** particles, **b** bonds. The simulations are 40 mm \times 80 mm \times 10 mm along the x , y , and z -axes, respectively. **a** The simulations consist of two types of particles that represent the relatively intact host rock (gray), and the weaker preexisting fault material (black). **b** Three types of bonds connect the particles. The bonds represent the on-fault (yellow), damage zone (blue–green), and off-fault material (dark blue). The on-fault bonds connect the weak on-fault particles. The off-fault bonds connect the strong off-fault particles. The damage zone bonds connect pairs of on-fault and off-fault particles. The faults extend from the front to the back of the system along the z -axis. The faults have the roughness parameters (Hurst exponents) of natural faults. To build the faults, we label particles within 1 mm of a surface with the prescribed roughness parameters (Hurst exponents and amplitude) as the weak on-fault particles. The remaining particles in the system are the off-fault particles. We vary the orientation of the fault, θ , between 15 to 60° from the maximum compression direction, σ_1

particle radius is one millimeter, using this distances ensures that the resulting fault surface does not contain gaps. The particle labels then determine the bond labels, which indicate whether the bond represents on-fault, damage zone, or off-fault material.

Fault zones consist of a principal slip zone and surrounding damage zone embedded within a more intact host rock (e.g., Faulkner et al., 2010). Based on the relationship between fracture density and mechanical properties (e.g., Paterson & Wong, 2005), these materials are expected to increase in strength and stiffness with increasing distance from the fault core. This range in material parameters is

also consistent with seismic imaging that reveals low velocity zones around faults, and increasing velocities in the surrounding host rocks (e.g., Allam et al., 2014). We therefore prescribe three different sets of cohesion, c , and normal stiffness, k_n , to the on-fault, damage zone, and off-fault bonds. We vary the values of these microparameters to achieve the desired mechanical properties, including the uniaxial compressive strength, UCS , and Young's modulus, E .

We first find the set of microparameters that produces the UCS and E of intact granite, and then prescribe these microparameters to the off-fault bonds. To find these parameters, we build

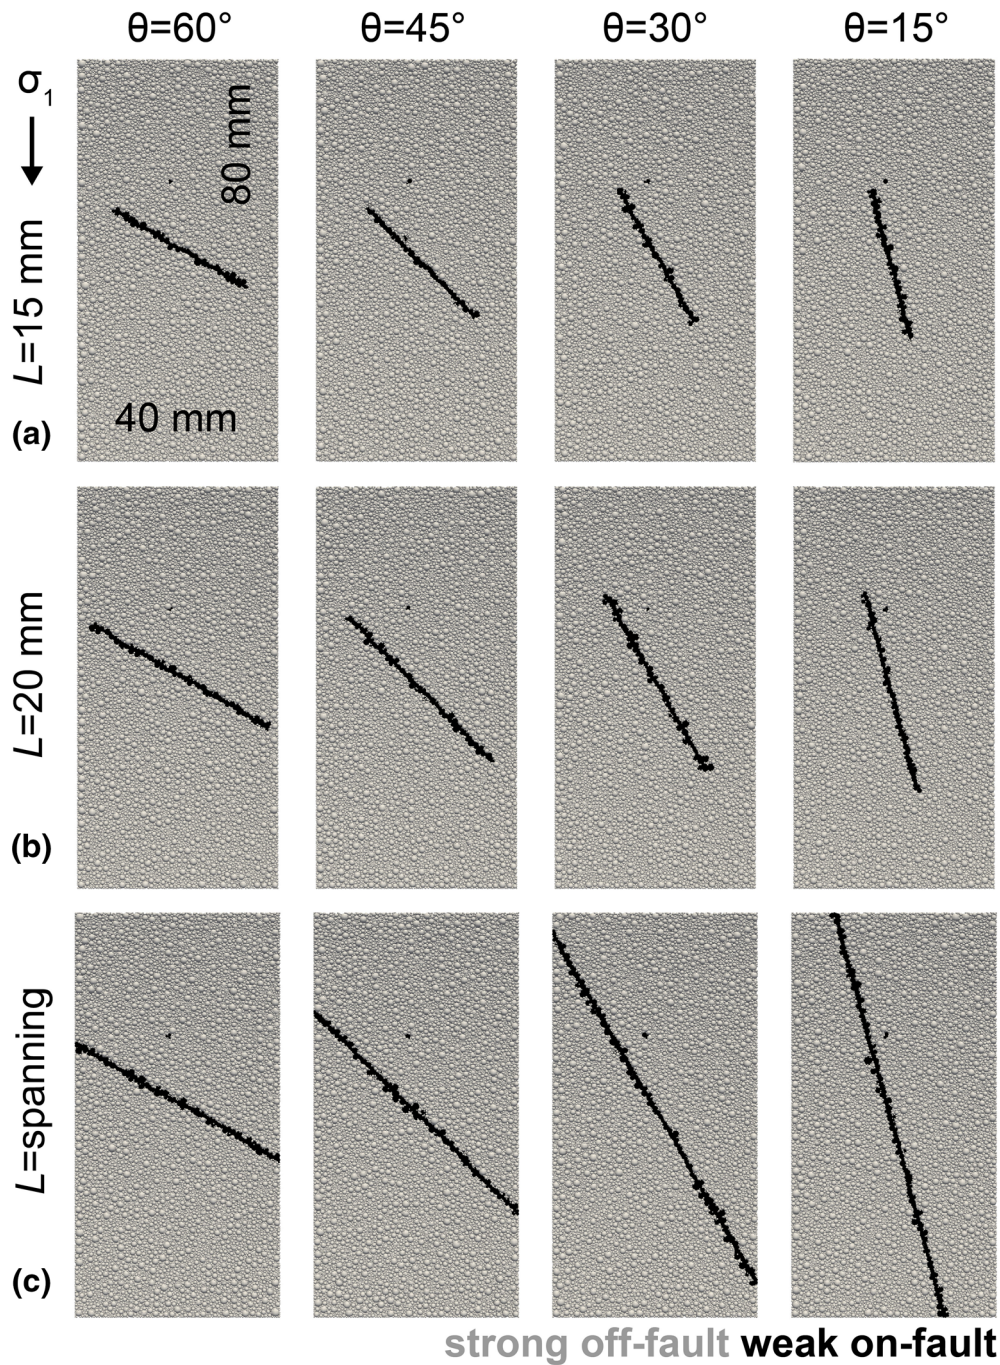


Figure 2

Fault geometries with the four tested orientations and **a** 15 mm long faults, **b** 20 mm long faults, and **c** system-spanning faults

homogeneous simulations in which all the bonds have the same set of microparameters, and then load them in uniaxial compression (Figure S1). Following this

step of the calibration, we identify the set of microparameters that produces UCS and E of 270 MPa and 40 GPa, respectively, within the ranges

of intact granite (e.g. Griffiths et al., 2017). Table S1 lists the full suite of microparameters found through calibration.

The strength and stiffness of the principal slip zone of a fault depend on the rate of healing, rate of crustal loading and damage accumulation, and time elapsed since the previous earthquake (e.g., Lyakhovskiy et al., 2001). We model systems with a relatively weak fault, so we find the set of microparameters for the on-fault bonds that produce UCS and E of 10 MPa and 4.5 GPa, respectively. Then, we prescribe microparameters to the damage zone bonds between the endmembers of the on-fault and off-fault zones. This process produces increasing UCS and E from the on-fault to the damage zone, and then to the off-fault volume. These simulations thus assume that the crustal loading that reactivates the preexisting healed fault occurs quickly compared to the healing of the surrounding host rock.

2.2. Machine Learning Models

Throughout this analysis, we refer to the machine learning models as “models”, and the discrete element method simulations as “simulations” in order to differentiate between these tools. We build the models to predict the timing of macroscopic failure, and thus the peak axial stress, using the particle velocities and corresponding strain fields. We do not include information about the local stress fields because the stress tensor is difficult to measure in the field with high spatial resolution. We provide the three-dimensional velocity and strain fields from timesteps of the simulations to the models, and build the models to predict the normalized distance to macroscopic failure. We define the normalized time to failure as $(t_f - t)/t_f$, where t_f is the timestep when macroscopic failure occurs and t is the present timestep of the simulation. Because the applied axial velocity loading of the simulations is constant, the timestep is proportional to the cumulative axial strain. Thus, the time to failure calculated from the timesteps is equal to the time to failure calculated from the cumulative axial strain. Following previous laboratory work (McBeck et al., 2020a), we consider the time when the system achieves the maximum axial stress as the time of macroscopic failure.

We use XGBoost regression models to perform these predictions (Chen & Guestrin, 2016). We chose XGBoost because of its efficiency and accuracy (e.g., Bühlmann & Yu, 2003; Friedman, 2001). We develop different sets of models to predict each target for each numerical simulation. The combination of four fault orientations, three fault lengths, and three levels of applied confining stress, leads to 36 unique simulations, with one simulation for each combination of parameters.

From the velocity and strain fields, we extract a series of features as for the machine learning models. We only use the timesteps of the simulations that occur preceding macroscopic failure. We calculate the mean of the three velocity components of all the particles within cubes with 1 mm long sides at increments of 1000 timesteps throughout the simulation. Using this grid of velocity vectors, we calculate the nine components of the three-dimensional strain tensor. We then calculate quantities that measure the volumetric strain: the divergence or first invariant of the strain tensor, I_1 , and quantities that measure the shear strain: the second invariant of the strain deviator tensor, J_2 , and the curl. Table S2 lists all the features.

This method of feature extraction produces several hundred measurements of each feature at each simulation timestep, and thus at various times preceding macroscopic failure. Each sample of the dataset provided to the machine learning models is unique in its combination of time and position. We split the training and testing dataset so that none of the data in the training dataset occurs in the testing data. In particular, we split the data so that data from the same timestep of the simulation does not occur in both the training and testing datasets, and the testing dataset includes about 20% of the data. Consequently, we split the training and testing data sets so that the models are trained on different portions of time in the simulation than they are tested. To capture the influence of random variations in how the data are split, for each of the 36 unique simulations we build ten different models with different splits of training and testing data.

We estimate the influence of each feature on the predictions using a widely used metric: Shapely Additive Explanations (SHAP) (Lundberg & Lee,

2017). We compare the mean absolute value of the SHAP value (mean |SHAP|) of each feature, and thus focus on the overall influence of that feature on the model prediction. We perform a grid search over the hyperparameter-space to find the best set of hyperparameters for each model (Lundberg & Lee, 2017).

3. Results

Here, we first describe the fracture networks that develop in all of the simulations, and how the confining stress and fault geometry control the fracture networks that develop. Then, we follow the evolution of the fracture networks throughout the simulations up to macroscopic failure, when the simulation achieves the maximum axial stress. We then compare the ability of the models to predict the timing of failure using the strain and velocity fields of the simulations. We finally identify the characteristics of the strain and velocity fields that control the model predictions, and examine how they evolve toward failure.

3.1. Fracture Network Development

The confining stress, fault orientation and fault length play key roles in the final fracture geometry that develops by the end of the simulation, following the peak axial stress and subsequent stress drop (Fig. 3). Systems with longer faults tend to concentrate deformation within the preexisting healed fault (i.e., on-fault), whereas systems with shorter faults develop more fractures outside of the preexisting fault (i.e., off-fault). For example, when the fault orientation, θ , is 30° , the volume of off-fault fractures is larger when the fault length, L , is 15 mm, than when $L = 20$ mm or is system-spanning, at each level of applied confining stress. This relationship also holds for the simulations with $\theta = 45^\circ$ at each level of confining stress, σ_2 . However, when the faults are less well-oriented for shear failure, with $\theta = 60^\circ$ and $\theta = 15^\circ$, this relationship breaks down. Instead, fractures develop throughout the off-fault region with similar densities in simulations with varying fault length at the same σ_2 . The similarity of the relationship between fault orientation and the magnitude of

off-fault deformation in these simulations and laboratory experiments (Giorgetti et al., 2019) provides evidence that the simulations closely approximate the behavior of natural rocks.

The fault length appears to exert a stronger control on the resulting fracture geometry when the preexisting faults are less well-oriented for shear failure, with $\theta = 60^\circ$ and $\theta = 15^\circ$ (Fig. 3). In simulations with $\theta = 30^\circ$ – 45° , the off-fault fractures propagate approximately coplanar to the preexisting faults when L is not system spanning, thereby producing systems similar to those with system-spanning preexisting faults. However, in simulations with $\theta = 60^\circ$ and $\theta = 15^\circ$, similar behavior only occurs when $\sigma_2 = 0$. Higher confining stress suppresses coplanar propagation from the tips of the preexisting fractures in these simulations.

Confining stress also controls the prevalence of fault-oblique fractures (Fig. 3). In simulations with no confining stress, the off-fault fractures tend to develop at oblique angles to the preexisting fault. In simulations with higher confining stress, the prevalence of these fractures tends to decrease. Simulations with system-spanning faults, for example, host either one or no fault-oblique fractures when $\sigma_2 > 0$. Simulations that do not include a preexisting fault develop diffuse fracture networks prior to failure, similar to some laboratory observations (e.g., Renard et al., 2019).

Next, we examine the evolving fracture networks in four characteristic simulations leading to macroscopic failure. These simulations include fault geometries with the longest and shortest tested faults, and orientations that are less and more well-oriented for shear failure, $\theta = 15^\circ$ and $\theta = 30^\circ$. The simulations include three phases of macroscopic stress accumulation due to the applied axial strain (Figs. 4, 5, Figure S2). First, the axial stress increases quasi-linearly with the axial strain, then the rate of axial stress accumulation slows, and finally the axial stress decreases. Macroscopic failure thus occurs when the simulations achieve the maximum axial stress.

Fracture development includes several key phases preceding macroscopic failure in simulations with shorter faults, $L = 15$ mm, (Fig. 4), that do not occur when the faults are system-spanning (Fig. 5). In simulations with $L = 15$ mm, fractures first develop

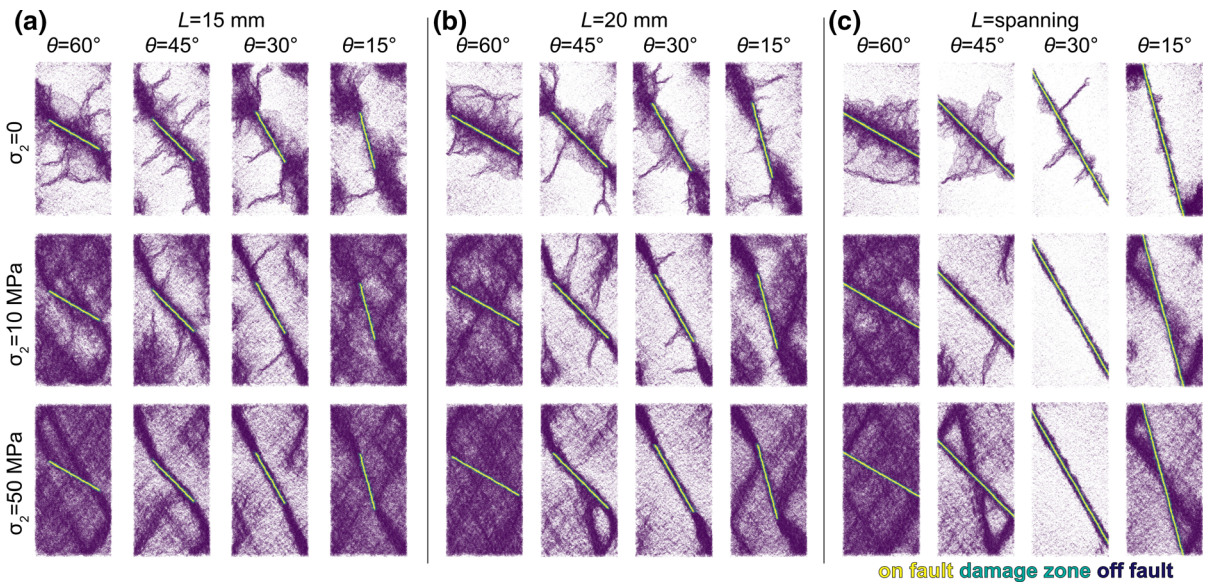


Figure 3

Final fracture geometry for simulations with **a** 15 mm long faults, **b** 20 mm long faults and **c** system-spanning faults at the end of the simulation, following macroscopic failure. The images show the full three-dimensional system along the z -axis. The columns show the simulations with different fault orientations. The rows show the simulations with different confining stresses. The colors correspond to the different types of broken bonds: on-fault (yellow), damage zone (blue-green), and off-fault (dark blue)

along the weaker preexisting fault, and then they propagate from the tips of the preexisting fault (Fig. 4). Throughout this process, fractures also develop off-fault with a diffuse distribution. The confining stress and orientation of the preexisting fault control how the fractures propagate from the tips of the preexisting fault. When the preexisting fault is well-oriented for shear failure, with $\theta = 30^\circ$, increasing confining stress promotes coplanar propagation from the tips of the fault. In particular, when the system experiences no confining stress, the fractures propagate approximately parallel to the σ_1 direction (Fig. 4c). When the system experiences higher confining stress, the off-fault fractures propagate coplanar to the preexisting fault (Fig. 4d). The observed influence of confining stress on the fracture geometry is consistent with laboratory observations and numerical simulations that observed that increasing confining stress promoted coplanar propagation from the tips of preexisting faults rather than out-of-plane propagation (Madden et al., 2017).

When the preexisting fault is less well-oriented for shear failure, with $\theta = 15^\circ$ (Fig. 4a, b), confining stress has a similar influence on the resulting fracture

geometry as when $\theta = 30^\circ$. Increasing the confining stress reduces the density of fault-oblique off-fault fractures. Varying θ also influences the spatial localization of the off-fault fractures that develop at the tips of the preexisting faults. These fractures are more localized when $\theta = 30^\circ$ than when $\theta = 15^\circ$.

When the faults are system spanning (Fig. 5), fracture development concentrates primarily within the preexisting fault (i.e., on-fault), as expected. Similar to the systems with shorter faults, fracture development first initiates along the preexisting fault, and later fractures propagate throughout the off-fault region. Increasing the confining stress and rotating the fault away from the optimal orientation for shear failure increases the density of off-fault fractures.

The varying spatial distribution of off-fault fractures and propagation from the tips of the preexisting faults due to confining stress, fault orientation and fault length suggest that the characteristics of the velocity and strain fields that the machine learning models use to identify the timing of macroscopic failure may differ between the simulations. For example, the orientations relative to the maximum compression directions of the off-fault fractures that

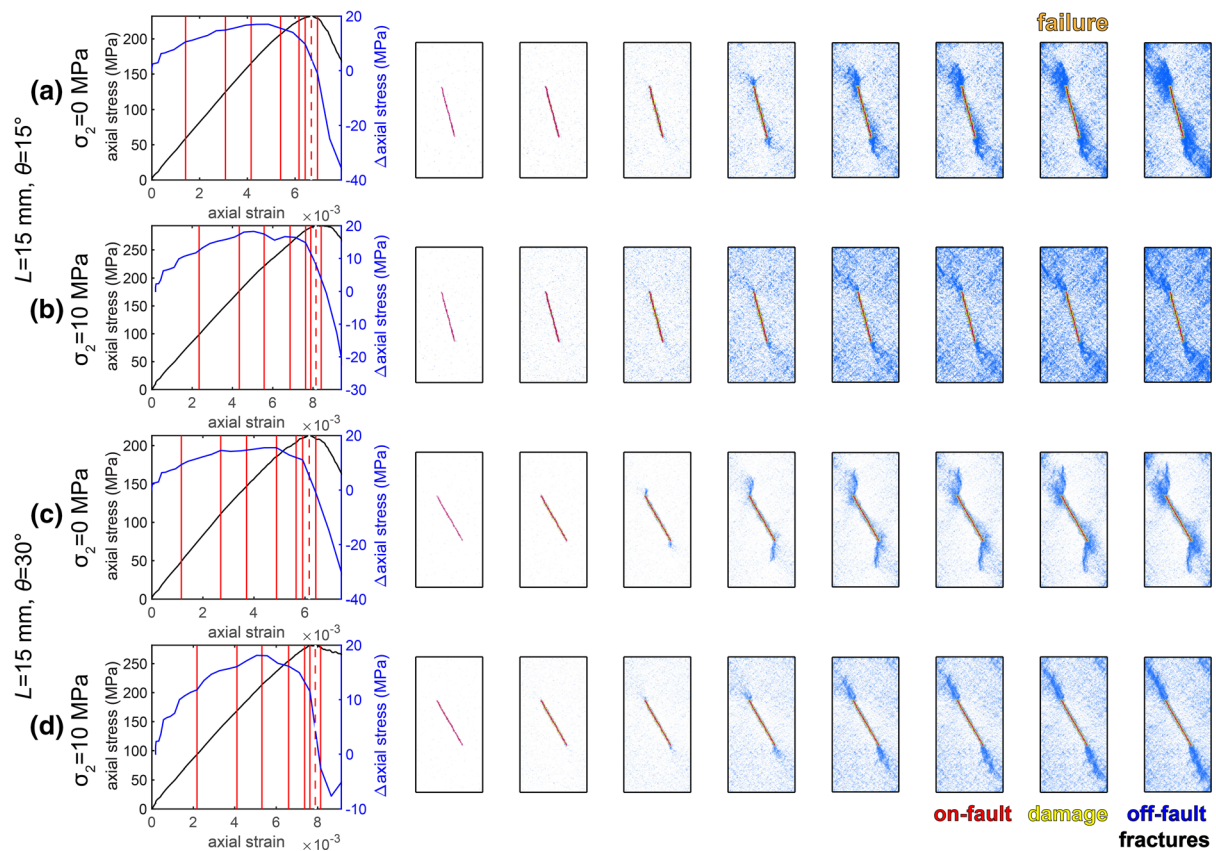


Figure 4

Fracture development in simulations with $L = 15$ mm, $\theta = 15^\circ$, and $\sigma_2 = 0$ (a), $\sigma_2 = 10$ MPa (b), and simulations with $L = 15$ mm, $\theta = 30^\circ$, and $\sigma_2 = 0$ (c), $\sigma_2 = 10$ MPa (d). Plots on the left show the macroscopic axial stress (black), change in axial stress (blue), and axial strain throughout each simulation, measured from the forces acting on and the movement of the top and bottom boundaries of the simulations. The red lines in the plots indicate the macroscopic conditions of the fracture networks shown on the right. The dashed lines highlight when macroscopic failure occurs. The images show the full three-dimensional system along the z -axis. The colors in the fracture networks show the fractures that develop (e.g., the broken bonds) in the on-fault material (red), damage zone (yellow), and off-fault material (blue). The second to last column of each row shows the fracture geometry at macroscopic failure

propagate from the tips of the preexisting faults suggest that they accommodate dilation in simulations with $L = 15$ mm, $\theta = 30^\circ$, and $\sigma_2 = 0$, and accommodate shear strain in simulations with higher confining stress, $\sigma_2 = 10$ MPa (Fig. 3).

Because we develop the models to predict the timing of macroscopic failure, we now compare the peak axial stress achieved in each simulation (Fig. 6). Increasing the fault length decreases the peak axial stress for simulations with faults well-oriented for shear failure, with $\theta=30^\circ$ and 45° , at all the tested confining stresses, as expected. Under no confining stress, increasing fault length also promotes lower peak axial stress when the faults are less well-

oriented for shear failure, $\theta=15^\circ$ and 60° . However, at higher confining stress $\sigma_2 = 10$ MPa and 50 MPa, the peak axial stress is similar for faults of varying lengths when $\theta = 15^\circ$ and 60° . Thus, when the fault is less well-oriented for shear failure, increasing the fault length does not lead to a significant reduction in the macroscopic strength.

The systems fail at the lowest axial stress when $\theta = 30^\circ$ and 45° (Fig. 6), as expected from theoretical expectations (e.g., Jaeger & Cook, 1979) and laboratory experiments (e.g., Donath, 1961; Giorgetti et al., 2019). Because these orientations maximize Coulomb shear stress, the lower axial stresses at these orientations and the corresponding approximately

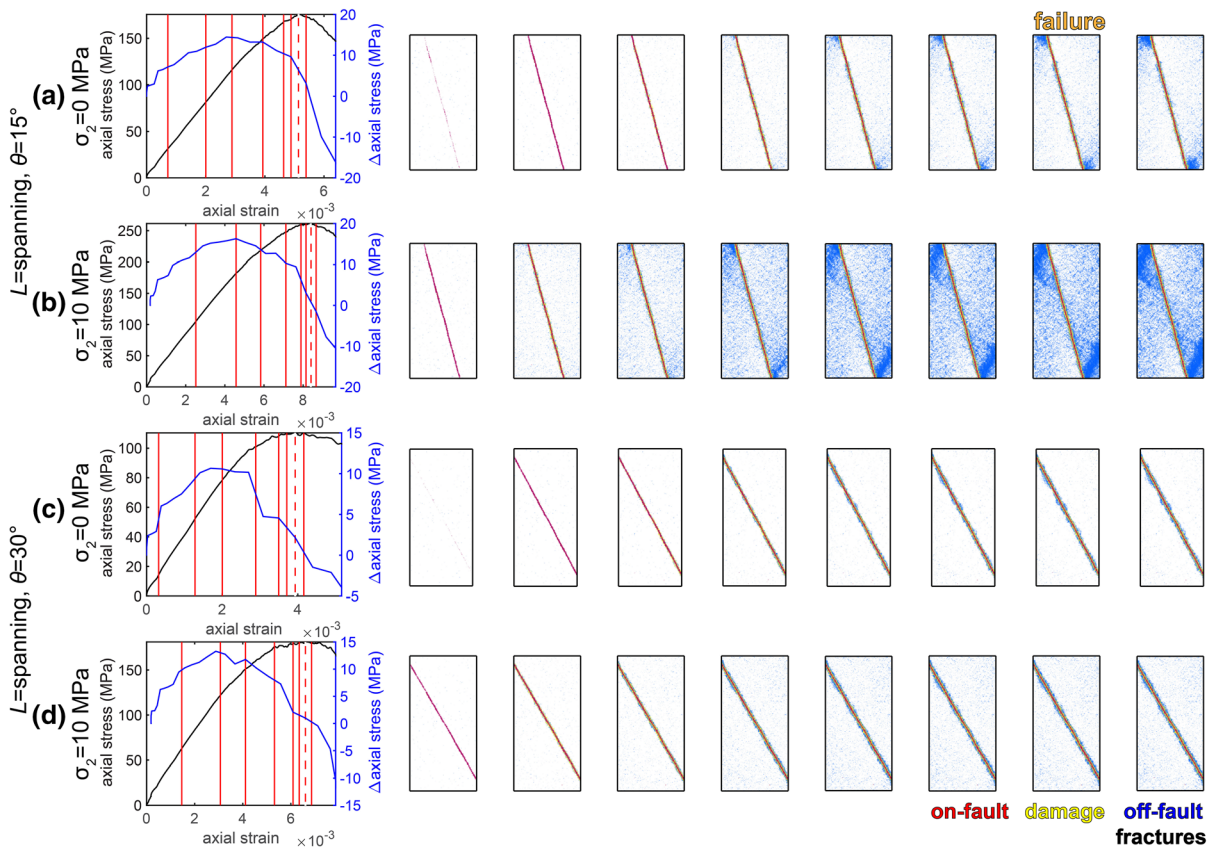


Figure 5

Fracture development in simulations with system-spanning faults, $\theta = 15^\circ$, and $\sigma_2 = 0$ (a), $\sigma_2 = 10$ MPa (b), and simulations with system-spanning faults, $\theta = 30^\circ$, and $\sigma_2 = 0$ (c), $\sigma_2 = 10$ MPa (d). Plots on the left show the macroscopic axial stress (black), change in axial stress (blue), and axial strain throughout each simulation, measured from the forces acting on and the movement of the top and bottom boundaries of the simulations. The red lines in the plots indicate the macroscopic conditions of the fracture networks shown on the right. The dashed lines highlight when macroscopic failure occurs. The images show the full three-dimensional system along the z -axis. The colors in the fracture networks show the fractures that develop (e.g., the broken bonds) in the on-fault material (red), damage zone (yellow), and off-fault material (blue). The second to last column of each row shows the fracture geometry at macroscopic failure

parabolic distribution of the peak axial stresses relative to θ provides additional evidence that the simulations closely approximate the behavior of natural rocks. When the faults are system-spanning, the peak axial stress with $\theta = 30^\circ$ is significantly lower than with $\theta = 45^\circ$, producing a more asymmetric distribution of the peak axial stress relative to θ than when the faults are shorter. Similar to the varying fracture geometries, these varying failure stresses suggest that the machine learning models may use different characteristics of the strain and velocity fields to predict the timing of failure.

3.2. Predicting the Proximity to Macroscopic Failure

To assess the ability of the models to predict the proximity to macroscopic failure, we examine the R^2 score of the models. High positive R^2 scores (> 0.7) indicate strong correlations between the predictions and observations. Scores of one indicate a perfect correlation between the predictions of the models and the observations. The models perform exceptionally well, with mean R^2 scores between 0.93 and 0.98 across all the simulations (Fig. 7). The performance of the models developed for simulations with the

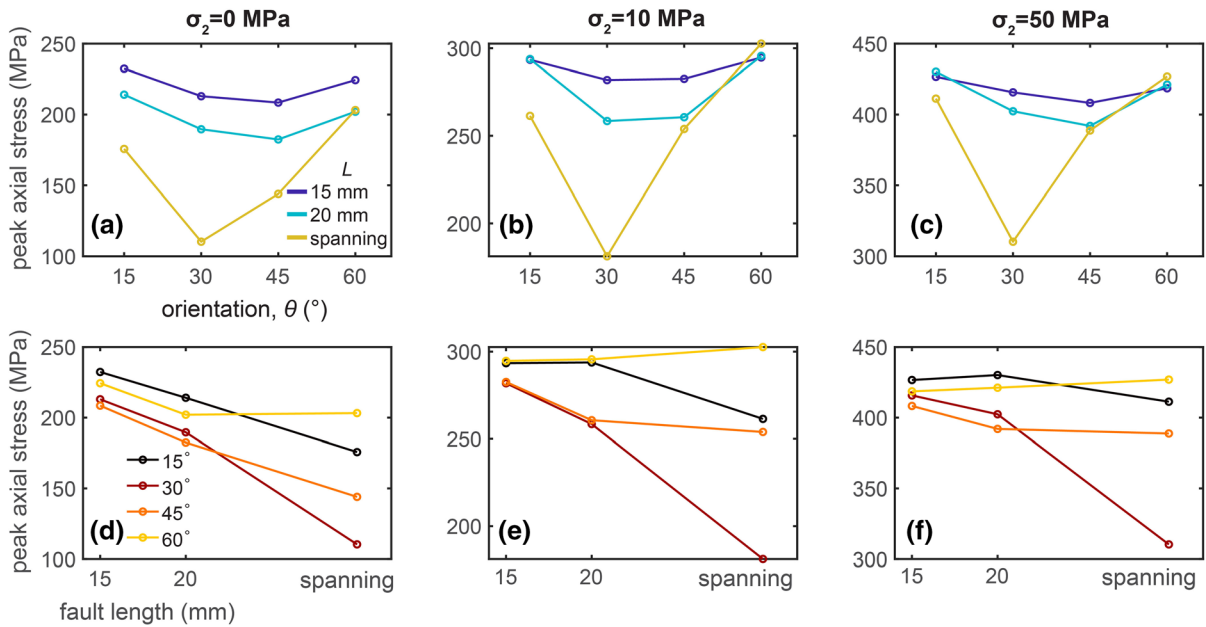


Figure 6

The influence of confining stress, orientation, and fault length on the peak axial stress. **a–c** Peak axial stress relative to orientation at each confining stress: **a** 0 MPa, **b** 10 MPa, and **c** 50 MPa. The colors of the curves indicate the fault length: $L = 15$ mm (dark blue), $L = 20$ mm (light blue), and system-spanning faults (yellow). **d–f** Peak axial stress relative to fault length at each confining stress: **d** 0 MPa, **e** 10 MPa, and **f** 50 MPa. The colors of the curves indicate the fault orientation: $\theta = 15^\circ$ (black), 30° (dark red), 45° (orange), and 60° (yellow)

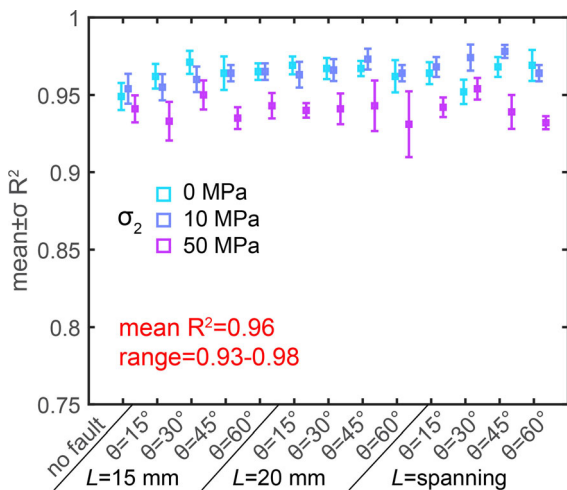


Figure 7

Performance of the models for all of the tested simulations: the mean \pm one standard deviation of the R^2 scores of the groups of ten models developed for each simulation. The ten models only differ in the split of the training and testing datasets. The horizontal axis indicates the geometry of the preexisting fault in the simulations. The color indicates the amount of confining stress: $\sigma_2 = 0$ (light blue), $\sigma_2 = 10$ MPa (blue), $\sigma_2 = 50$ MPa (pink)

highest tested confining stresses is slightly lower than the performance of the other models.

To identify the features that control the model predictions, we calculate a cumulative representation of the SHAP values of each feature. Because we develop ten models for each simulation that only differ in how the training and testing datasets are split, we calculate a normalized importance $s/\max(s)$ from the mean |SHAP| value of each feature, s . We then find the mean of the normalized values weighted by the R^2 score of the model across all the models for a given simulation, $s_s = \sum R^2(s/\max(s))/n$, where n is the number of models. We weight the normalized importance by the R^2 score of the model so that more successful models (with higher R^2) have a greater influence on the results compared to less successful models. The value s_s thus reflects the distribution of the |SHAP| values across all the features of all the machine learning models for a particular simulation (Fig. 8a). We also report the sum of s_s across all the simulations (Fig. 8b). To further constrain the features that most strongly control the model results, we

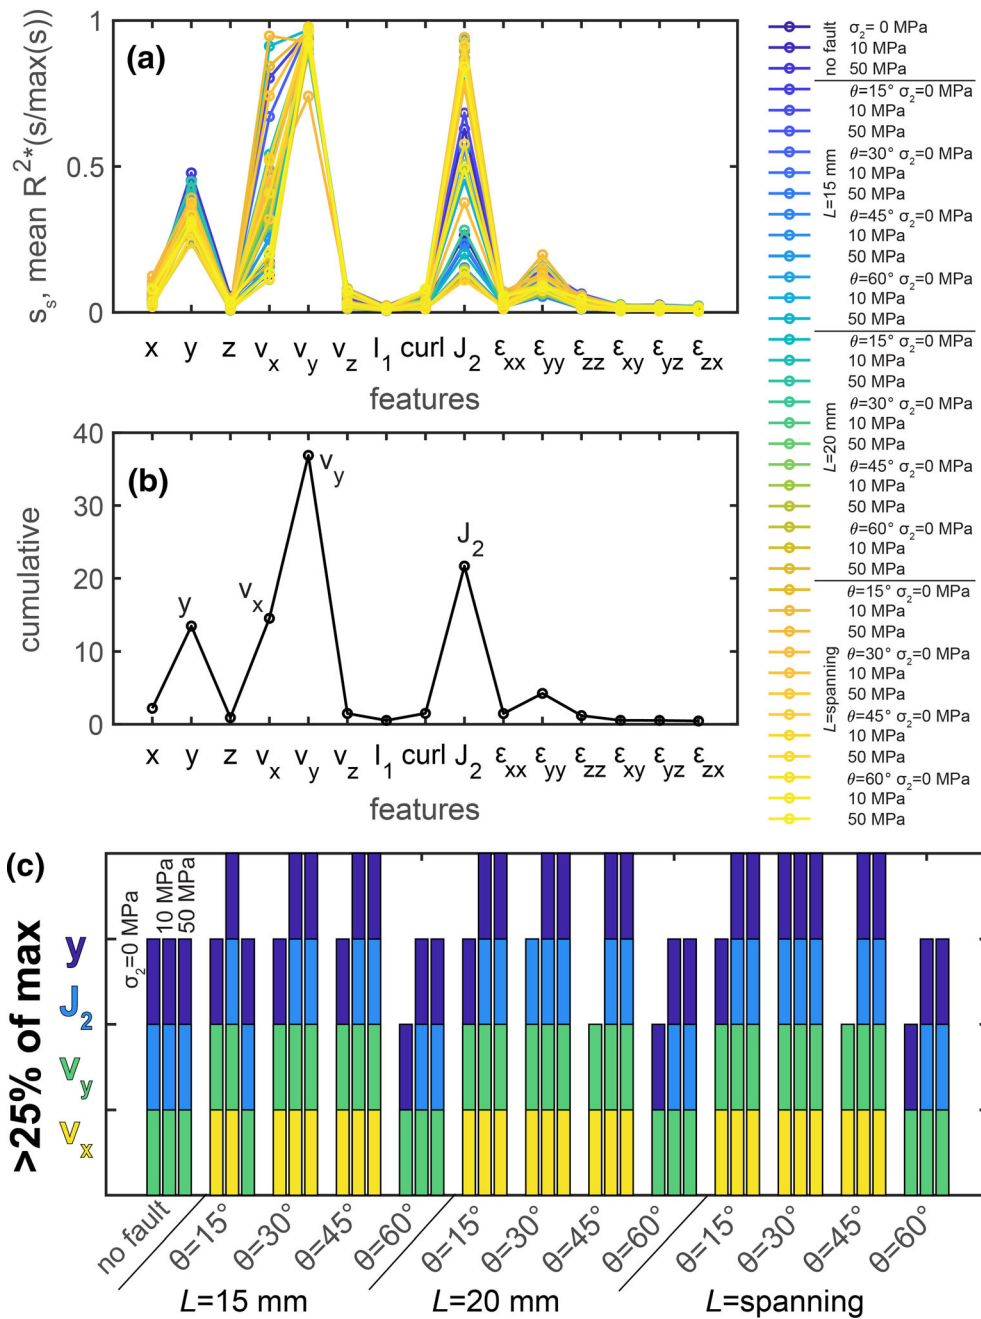


Figure 8

Influence of the velocity and strain characteristics on all the simulations: **a** distribution of the ISHAPL values, s_s , for each simulation, **b** cumulative distribution across all of the simulations, **c** the characteristics with s_s greater than 25% of the maximum value for each simulation. The color of the bars indicates the feature: dark blue (y-position), light blue (J_2), green (v_y), and yellow (v_x). Each group of three bars includes the results for the simulations with $\sigma_2 = 0$ MPa (left bar), 10 MPa (middle bar), and 50 MPa (right bar). The labels below the horizontal axis indicate the preexisting fault configuration. For example, the second group of three bars shows the results for the fault geometry with $L = 15$ mm and $\theta = 15^\circ$

identify the features with s_s greater than 25% of the maximum s_s of each simulation (Fig. 8c).

The cumulative distribution of s_s across all the simulations indicates that four features most strongly control the model results: the y -position in the simulation, the velocity components parallel, v_y , and perpendicular, v_x , to the maximum compression direction, and the second invariant of the strain deviator tensor, J_2 (Fig. 8b). The distribution of s_s for the individual simulations, however, indicates some variations between the simulations (Fig. 8a). Examining the characteristics that have s_s greater than 25% of the maximum s_s for each simulation (Fig. 8c) indicates that all of the models depend strongly on v_y , and nearly all of them depend on y . The models that depend strongly on J_2 tend to be those developed for simulations with higher confining stress. In particular, all of the models developed for the simulations with no applied confining stress do not depend on J_2 , except when $\theta = 30^\circ$ and when the fault is system-spanning or $L = 20$ mm. The models that depend on v_x tend to be those developed for simulations with $\theta < 60^\circ$. When $\theta = 60^\circ$ and when the simulations do not include a preexisting fault, v_x does not strongly control the model predictions. The only exception to this trend is when $L = 15$ mm and $\sigma_2 = 50$ MPa.

To better understand why the models depend on the characteristics identified from the SHAP values (Fig. 8), we now examine the evolution of three characteristics identified as highly important, J_2 , v_y , and v_x , and one that is not identified as important, I_1 (Fig. 9). We track these characteristics at two positions throughout three simulations, in the middle of the simulations and along the preexisting fault ($x = 20$ mm, $y = 40$ mm), and at a position in the lower portion of the system that does not cross the preexisting fault ($x = 20$ mm, $y = 20$ mm). At each of these positions, we report the mean value of the characteristic across all of the z -positions.

The evolution of v_y of the simulations that do not include a fault highlights that these simulations do not experience an extended macroscopic yielding phase (Fig. 9a). Instead, the local deformation is relatively homogeneous until about 10% to failure, when the v_y at the middle and lower positions begin to change, increasing at the middle position, and decreasing at the lower position.

These evolutions help indicate why the models depend on the various characteristics. Each of the example simulations experience a relatively systematic evolution of v_y toward macroscopic failure, while I_1 fluctuates more rapidly, sometimes switching between negative and positive values, and thus contractive and dilative strain (Fig. 9). The more systematic evolution of v_y compared to I_1 allows the models to build reliable relationships between this characteristic and the timing of macroscopic failure that produce strong correlations between the model predictions and the observed timing of failure.

The varying systematic evolution of v_x and J_2 also indicates why some models depend on these characteristics while others do not. The velocity component v_x fluctuates rapidly from positive to negative values, indicative of movement to the right and left in the simulation under no confining stress with no preexisting fault (Fig. 9a), and in the simulation with a system-spanning fault, $\theta = 60^\circ$ and $\sigma_2 = 50$ MPa. (Fig. 9c). In both of these simulations, the models do not depend strongly on v_x .

Similarly, the model developed with the simulation under no confining stress with no preexisting fault does not depend strongly on J_2 , whereas the models developed with the other two example simulations depend on this characteristic (Fig. 8c). Although the simulation without a preexisting fault experiences a relatively systematic evolution of J_2 , similar in its overall evolution to the other two simulations, the maximum value that J_2 achieves in this simulation at the selected sampling points is two orders of magnitude less than the maximum values in the other simulations. Thus, both the value of the characteristic and the consistency of its evolution toward failure appear to control whether the models use a particular characteristic to predict the distance to macroscopic failure.

3.3. Predicting Failure with Subsets of the Deformation Components

To further identify the characteristics that control the model predictions, we now examine the model performance when the models only use subsets of the data. The distribution of the SHAP values indicates that models that use the shear strain component J_2 ,

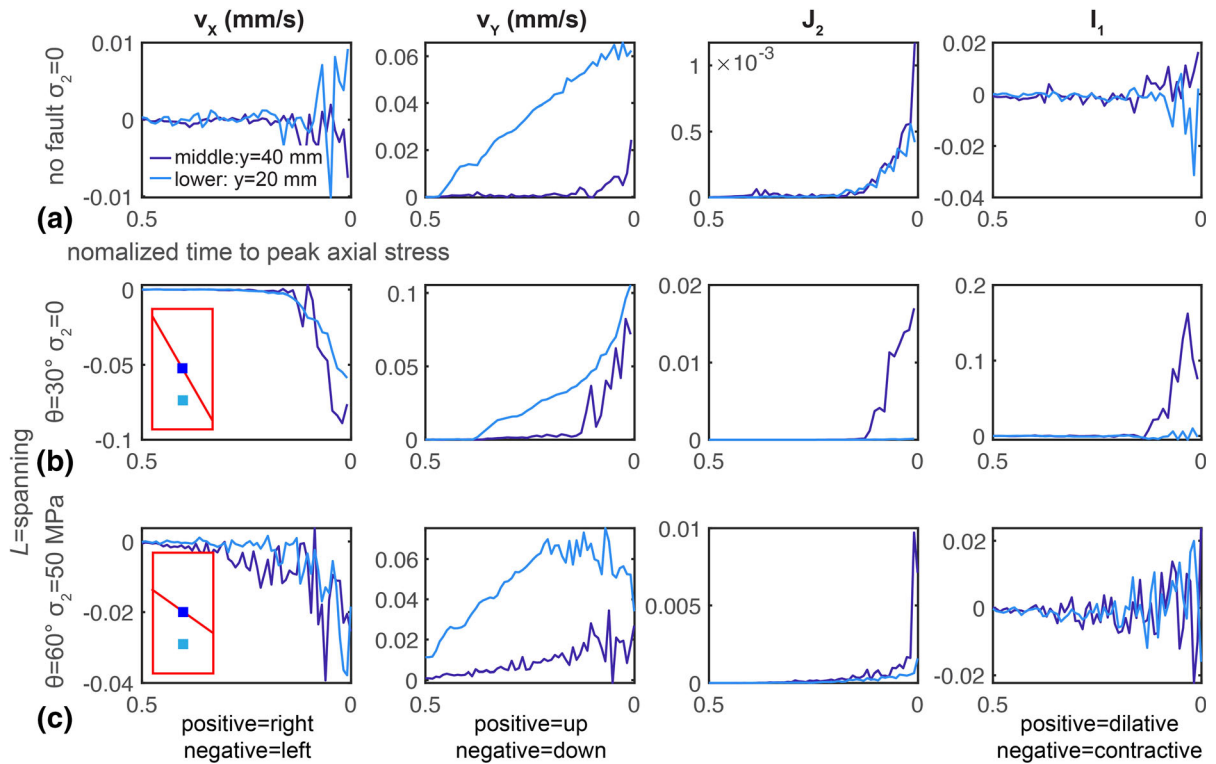


Figure 9

Evolution of three features that strongly control the model predictions relative to the normalized time to peak axial stress, $(t_f - t)/t_f$: v_x , v_y and J_2 , and one that does not, I_1 , in three example simulations: **a** when the system has no fault and $\sigma_2 = 0$, **b** when the fault is system-spanning, $\theta = 30^\circ$, and $\sigma_2 = 0$, and **c** when the fault is system-spanning, $\theta = 60^\circ$, and $\sigma_2 = 50$ MPa, at two example locations: in the middle of the simulation at $x = 20$ mm, $y = 20$ mm (dark blue), and in the lower portion of the simulation at $x = 20$ mm and $y = 20$ mm (light blue). The curves show the mean values of each of the four characteristics at the selected locations throughout the full thickness of the model, i.e., along the z -axis. The sketches in red in **b** and **c** show the preexisting faults, the boundaries of the simulations, and the sampling locations

and/or the velocity components v_y and v_x , may produce the most accurate models, while models that use the volumetric strain components, such as I_1 , may be the most inaccurate (Fig. 8). To test this idea, we develop models with three different subsets of the data: the components that describe the shear strain,

the components that describe the volumetric and normal strain, and the velocity components (Table 1).

Consistent with expectations, the models developed with the velocity components perform better than the models developed with only the shear strain components and the models developed with the normal strain components (Fig. 10). In particular, the mean R^2 scores of the models developed with the velocity components are 0.12–0.25 higher than the mean R^2 scores of the models developed with the volumetric strain components (Fig. 10e). Models developed from simulations with lower confining stress have the largest difference in the mean R^2 scores. This result occurs because the models developed for simulations with lower confining stress perform better than models developed for simulations with higher confining stress when they only have

Table 1

Velocity and strain characteristics used in the three types of models that use subsets of the data

Subset	Characteristics
Velocity components	x, y, z, v_x, v_y, v_z
Volumetric and normal strain components	$x, y, z, \epsilon_{xx}, \epsilon_{yy}, \epsilon_{zz}, I_1$
Shear strain components	$x, y, z, \epsilon_{xy}, \epsilon_{yz}, \epsilon_{zx}, J_2, \text{curl}$

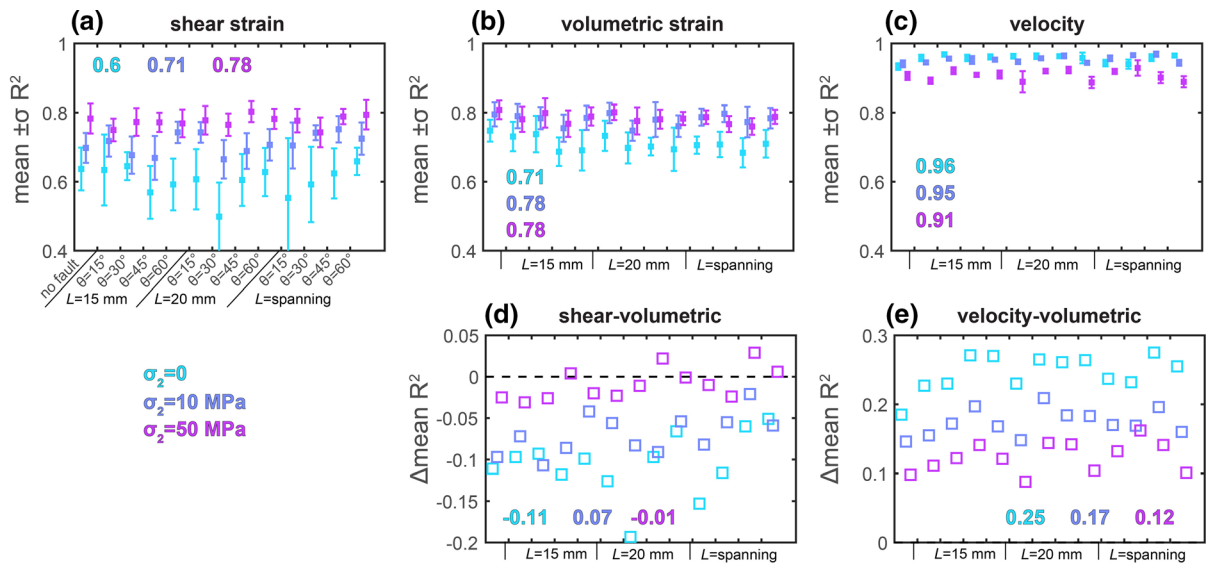


Figure 10

Comparing the performance of models developed with subsets of the strain and velocity components: the mean \pm one standard deviation of the R^2 scores of **a** models developed with the shear strain components, **b** models developed with the volumetric and normal strain components, and **c** models developed with the velocity components, and the difference between the mean R^2 score of each simulation for **d** models developed with the shear strain components and volumetric strain components, and **e** models developed with the velocity components and the volumetric strain components. Negative values in **d** indicate that the models developed with the volumetric strain components perform better than the models developed with the shear strain components. Positive values in **e** indicate that the models developed with the velocity components perform better than the models developed with the volumetric components

access to the velocity components. When the models only use the volumetric and normal strain components, the models developed for simulations with higher confining stress perform better than models developed for simulations with lower confining stress.

However, contrary to expectations, the models developed with only the shear strain components do not perform better than the models developed with only the volumetric and normal strain components. Instead, the mean R^2 scores of the models developed from the volumetric and normal strain components are 0.01–0.11 higher than the scores of the models developed from the shear strain components (Fig. 10d). The difference in the mean R^2 scores is the lowest when $\sigma_2 = 50$ MPa. Thus, when the confining stress is higher, the difference in the performance of the models is minor, and sometimes the models developed with the shear strain components perform better than the models developed with the volumetric strain components.

To better understand the key characteristics that control the three subsets of models, we now examine

the cumulative distribution of s_s across all the simulations (Fig. 11). Similar to the models developed with all of the strain and velocity components, the models developed with only the velocity components primarily depend on v_x , v_y and the y position (Fig. 11c). This result is consistent with the high s_s of these characteristics for the models developed with all of the data (Fig. 8). Similarly, the models developed with the shear strain components primarily depend on the curl of the velocity field, and on J_2 (Fig. 11a), consistent with the distribution of the s_s values of the models developed with all of the data. In contrast, the models that only use the normal and volumetric strain components primarily depend on characteristics that were not identified as highly important in the models that use all of the data. These models primarily depend on the normal strain components: ϵ_{xx} , ϵ_{yy} , and ϵ_{zz} . The higher cumulative s_s for ϵ_{yy} compared to the other normal strain components suggest why these models perform better or comparably to the models developed with the shear strain components. In particular, the models developed with all of the data and the models developed with the

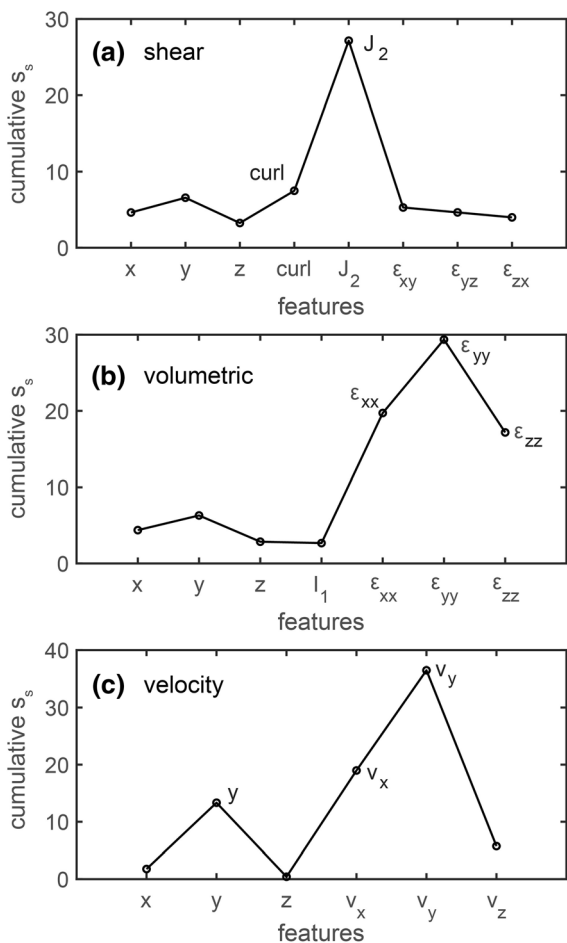


Figure 11

Distribution of the cumulative s_s for the models developed with the shear strain components (a), models developed with the volumetric and normal strain components (b), and models developed with the velocity components (c). The features labelled within each plot are those with cumulative s_s greater than 25% of the maximum cumulative s_s .

velocity components primarily depend on v_y . Because ϵ_{yy} is a function of v_y , it is not surprising that when we deprive the models of v_y , they are able to use ϵ_{yy} to perform the predictions with similar success.

4. Discussion

4.1. The Importance of the Displacement Field

We develop machine learning models to predict the proximity to macroscopic failure using the components of the strain and velocity tensors. The

models perform exceptionally well, with mean R^2 scores ranging from 0.93 to 0.98 across the simulations (Fig. 7). When the models have access to all of the velocity and strain components, the models primarily depend on the y -position, J_2 , v_y , and v_x (Fig. 8). The varying fracture networks produced in the different simulations (e.g., Figs. 3, 4, 5) suggest that different strain and velocity characteristics may help the models predict the timing of macroscopic failure. Indeed, models developed from simulations with lower confining stress do not depend strongly on J_2 , but only on the velocity components. Higher confining stress produces larger magnitudes of J_2 (Fig. 9), which allows the machine learning models to successfully predict the proximity of macroscopic failure using this measure of the shear strain field. The observed increasing dependence of J_2 with increasing confining stress is consistent with experimental observations that show a transition from tensile-dominated deformation to shear-dominated deformation with increasing confining stress (e.g., Paterson & Wong, 2005). In laboratory experiments, for example, during uniaxial compression, the brittle failure of rocks occurs through the propagation and opening of fractures aligned approximately parallel to the direction of σ_1 (Tapponnier & Brace, 1976). During triaxial compression, as the applied confining stress increases, the overall orientation of the dominant fracture network can trend more obliquely to the direction of σ_1 , closer to the orientation expected for shear failure (e.g., Renard et al., 2019). The orientation of individual fractures, however, within such system-spanning fault networks may remain oriented parallel to the direction of σ_1 , rather than along the apparent macroscopic orientation (McBeck et al., 2021a). Despite this difference between the true orientation of individual fractures and the overall fracture network, such fracture networks can accommodate shear strain (McBeck et al., 2019, 2020b; Renard et al., 2019).

The varying importance of J_2 , and consistency of the importance of v_y and v_x (Fig. 8) highlight that the components of the velocity field, and thus the displacement field, provide more reliable information about the proximity of macroscopic failure than the strain field. Moreover, when we develop models from subsets of the strain and velocity fields, the most

successful models are those developed with the velocity components, rather than the shear strain components or volumetric and normal strain components (Fig. 10). One may expect that the strain field may provide more valuable information about the timing of macroscopic failure than the displacement field. For example, fault slip and opening can produce sites of higher shear and dilative strain localized on the fault, producing systematic relationships between applied loading and cumulative strain. Because the displacement field within the host rock surrounding a preexisting fault can evolve in unexpected ways due to slip and opening on rough faults, one may expect that the relationship between the displacements produced by fault movement and loading, and the timing of macroscopic failure may be more difficult for machine learning models to recognize than the systematic relationships of cumulative strain and macroscopic failure. Contrary to expectations, the present analysis suggests that the displacement field also provides critical information about the proximity of macroscopic failure, as well as the components of the strain field (i.e., Fig. 12).

4.2. The Importance of the Deformation Field Parallel to the Loading Direction

The s_s distribution of the characteristics of the models developed with all of the strain and velocity components (Fig. 8), and this distribution for the models developed with subsets of the components (Fig. 11) highlight that the components of the deformation field that describe the displacement or strain parallel to the σ_1 direction are critical for successfully predicting the timing of macroscopic failure (i.e., Fig. 12). In particular, all of the models developed with all of the characteristics depend on v_y , the velocity component parallel to the axial loading, while only some of the models depend on v_x (Fig. 8). When the models are developed with only the normal and volumetric strain components, they primarily depend on the strain component parallel to the σ_1 direction, ϵ_{yy} (Fig. 11). Thus, despite the variety of fracture networks that develop with varying confining stresses and preexisting fault geometries, the deformation field parallel to the dominant loading direction provides the most systematic information

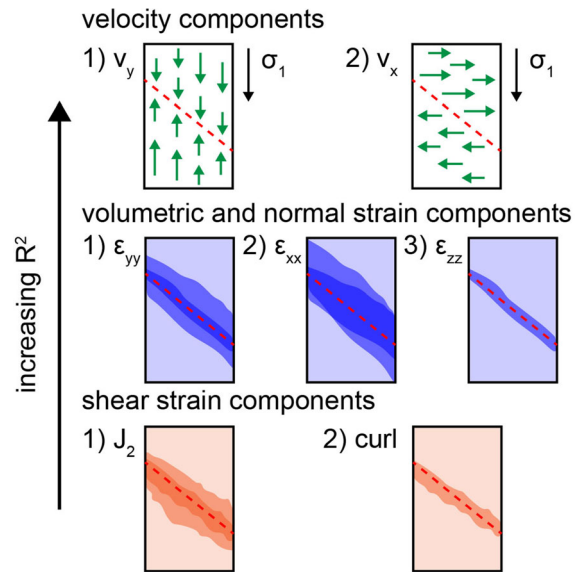


Figure 12

Sketch of the key conclusions. The models are the most successful, with the highest R^2 score, when they use the velocity components to predict the timing of macroscopic failure, and are more successful when they use the volumetric and normal strain components than when they use the shear strain components. Thus, tracking the velocity field may provide useful information about the timing of macroscopic failure, as well as tracking the strain field. The models that rely only on the velocity components primarily depend on v_y , and secondarily on v_x . Thus, the velocity component parallel to the maximum compression direction may provide the most useful information about impending failure. The models that rely only on the volumetric and normal strain components depend on the components in the following order: (1) ϵ_{yy} , (2) ϵ_{xx} , (3) ϵ_{zz} . Thus, the normal strain component that is dominated by axial-parallel contraction, ϵ_{yy} , may be more useful for predicting the timing of failure than the other normal strain components, which host both dilation and contraction in these systems. The models that rely only on the shear strain components depend primarily on J_2 , and secondary on the curl. Thus, efforts to recognize the onset of the precursory stage leading to large earthquakes may benefit from tracking J_2 rather than the curl of the displacement field

of the tested velocity and strain components (e.g., Fig. 9). This systematic information enables the machine learning models to make accurate predictions about the timing of macroscopic failure.

Previous machine learning analyses that predicted the proximity of macroscopic failure in triaxial compression experiments did not examine the individual nine components of the strain tensor, and instead only used the curl of the displacement field and the first invariant of the strain tensor, I_1 , to predict the timing of failure (McBeck et al., 2020a).

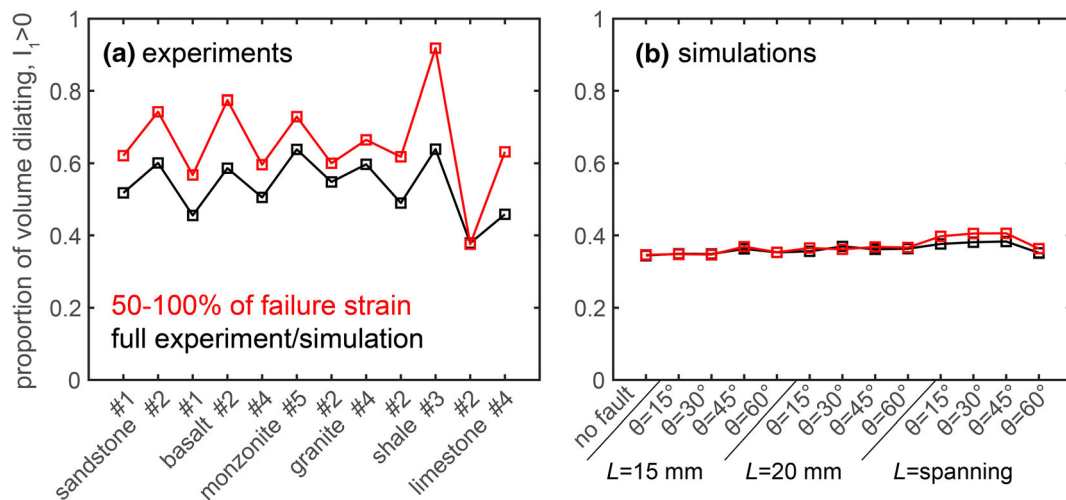


Figure 13

Proportion of the rocks deformed in triaxial compression experiments (a) and the simulations (b) that hosts dilation throughout the full experiment (black), and within 50–100% of the failure strain (red). The laboratory rocks host more dilation than the simulations throughout the full loading history, up to macroscopic failure, and in the time period between 50 and 100% of the axial strain at failure

This analysis found that the models depended primarily on the dilation, the positive I_1 , rather than the curl or the negative I_1 to predict the timing of failure. In contrast, in the present simulations, the models primarily depend on the velocity and strain components that characterize the deformation parallel to the σ_1 direction, v_y and ε_{yy} , and J_2 in models under higher confining stress. Due to the triaxial compression loading conditions in these simulations, preceding macroscopic failure ε_{yy} is generally negative throughout loading, indicative of contraction. Thus, when the models have access to all of the strain and velocity components, they do not use any of the volumetric or normal strain components indicative of dilation, and instead rely on y , v_x , v_y , and J_2 . When the models have access to only the volumetric strain components, they primarily depend on the normal strain component that is dominated by contraction, ε_{yy} . In contrast, the models developed from the triaxial compression experiments primarily depend on the dilation, rather than the curl.

In the simulations that include a preexisting fault well-oriented for shear failure, at 30° – 45° , we expect that shear deformation may dominate the system, and thus that the models should depend on components of the shear strain. However, even for simulations that lack a preexisting fault, the models do not depend on

measurements of the dilation, but primarily on J_2 and v_y . These models may not rely on the volumetric or dilative strain, I_1 , because the components of the velocity field and J_2 evolve with a more systematic evolution toward failure than I_1 (e.g., Fig. 9). In addition, the numerical simulations accumulate less dilative strain than the intact rocks deformed in the experiments. Indeed, comparing the proportion of the volume of the experiments and simulations that host dilation, where $I_1 > 0$, out of the total volume throughout each experiment and simulation indicate that the laboratory rocks host 40–60% dilation (McBeck et al., 2020b), whereas the simulations host about 30% (Fig. 13). Comparing this proportion in the final stages of the experiments and simulations, from 50 to 100% of the failure strain reveals an even larger difference in the volume proportion between the experiments and simulations. Thus, in the triaxial compression experiments, more and more of the rock volume experiences dilation as the rocks approach macroscopic failure. In contrast, as the simulations approach macroscopic failure, the proportion of the simulation that hosts dilation does not significantly increase. Whereas the magnitude of dilation provides systematic information that the machine learning models can use to predict the timing of macroscopic failure in the laboratory experiments, the dilative

strain does not follow such systematic relationships in the numerical simulations, causing the machine learning models to use other components of the velocity and strain tensor to successfully predict failure.

The simulations may experience less dilation than the experiments due to the initial porosity caused by the packing of the spherical particles. Whereas these experimental rocks generally have < 10% porosity (McBeck et al., 2020b), the initial porosity of these simulations is about 20%, although we use the maximum feasible range of particle radii to achieve the lowest initial porosity under the computational constraints. The higher initial porosity of the simulations promotes more contraction, and enables shear deformation to occur without significant dilation. The lower porosity of the laboratory rocks requires that the dilation of rough fracture surfaces occurs before these surfaces can slide. Indeed, dilative strain is often paired with shear strain in these experiments (McBeck et al., 2020b), highlighting the necessity of dilating before or concurrent with shearing.

5. Conclusions

We develop machine learning models to predict the timing of macroscopic failure in three-dimensional discrete element method simulations of triaxial compression. Comparing the predictive power of components of the velocity and strain fields in these models identify four characteristics that primarily control the model predictions: the y -position in the simulation, J_2 , v_y , and v_x . With increasing confining stress, the predictive power of J_2 increases, consistent with laboratory observations that show a transition from tensile-dominated to shear-dominated deformation with increasing confining stress (e.g., Paterson & Wong, 2005). In contrast to expectations from previous machine learning analyses (McBeck et al., 2020a; Shreedharan et al., 2021a), observations of varying seismic wave properties prior to some large earthquakes (e.g., Whitcomb et al., 1973), and laboratory observations of decreasing P-wave velocity prior to stick slip events (e.g., Shreedharan et al., 2021b), the volumetric strain component indicative of

dilation, I_1 , does not strongly control the model predictions. The models only depend on the normal strain components that quantify both dilation and contraction in these simulations, ε_{xx} and ε_{zz} , when we force the models to only use the normal and volumetric strain components. However, even these models primarily depend on the normal strain component that is dominated by contraction, ε_{yy} . The difference between the present analysis and previous machine learning analysis on the triaxial compression of rocks (McBeck et al., 2020a) may occur because the rocks deformed in the experiments are dominated by dilation, with increasing volume proportions toward failure (McBeck et al., 2020a, 2020b), whereas the numerical simulations do not experience similar strain accumulation (Fig. 13). The models developed with only the velocity components perform better than the models developed with only the volumetric and normal strain components, and better than the models developed with only the shear strain components. This result highlights that the velocity field, as well as the strain field, may provide valuable information about the onset of the precursory phase leading to catastrophic failure, such as earthquakes. In particular, the velocity field parallel to the maximum compression direction may provide the most valuable information about the timing of failure, regardless of the orientation of the preexisting fault.

Acknowledgements

The study was funded by the Norwegian Research Council (grant 300435), and UNINETT Sigma2 AS (project NN9806K). The codes to build and run the models are available on GitHub (https://github.com/jmbeck/Rock_Def_ESyS). François Renard provided valuable suggestions on drafts of this work. We thank Editor David and an anonymous reviewer for suggestions that improved this work.

Funding

Open access funding provided by University of Oslo (incl Oslo University Hospital). The Norwegian Research Council (grant 300435) funded this work.

Declarations

Conflict of interest There are no conflicts of interest.

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/>.

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

REFERENCES

- Abe, S., Place, D., & Mora, P. (2004). A parallel implementation of the lattice solid model for the simulation of rock mechanics and earthquake dynamics. *Pure and Applied Geophysics*, *161*(11), 2265–2277.
- Aggarwal, Y. P., Sykes, L. R., Armbruster, J., & Sbar, M. L. (1973). Premonitory changes in seismic velocities and prediction of earthquakes. *Nature*, *241*(5385), 101–104.
- Allam, A. A., Ben-Zion, Y., Kurzon, I., & Vernon, F. (2014). Seismic velocity structure in the Hot Springs and Trifurcation areas of the San Jacinto fault zone, California, from double-difference tomography. *Geophysical Journal International*, *198*(2), 978–999.
- Bakun, W. H., Aagaard, B., Dost, B., Ellsworth, W. L., Hardebeck, J. L., Harris, R. A., & Waldhauser, F. (2005). Implications for prediction and hazard assessment from the 2004 Parkfield earthquake. *Nature*, *437*(7061), 969–974.
- Bobet, A., & Einstein, H. H. (1998). Fracture coalescence in rock-type materials under uniaxial and biaxial compression. *International Journal of Rock Mechanics and Mining Sciences*, *35*(7), 863–888.
- Bühlmann, P., & Yu, B. (2003). Boosting with the L2 loss: Regression and classification. *Journal of the American Statistical Association*, *98*(462), 324–339.
- Candela, T., Renard, F., Klinger, Y., Mair, K., Schmittbuhl, J., & Brodsky, E. E. (2012). Roughness of fault surfaces over nine decades of length scales. *Journal of Geophysical Research: Solid Earth*, *117*(B8), B08409.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 785–794.
- Chiapparba, C., De Gori, P., Segou, M., & Cattaneo, M. (2020). Seismic velocity precursors to the 2016 Mw 6.5 Norcia (Italy) earthquake. *Geology*, *48*(9), 924–928.
- Donath, F. A. (1961). Experimental study of shear failure in anisotropic rocks. *Geological Society of America Bulletin*, *72*(6), 985–989.
- Faulkner, D. R., Jackson, C. A. L., Lunn, R. J., Schlische, R. W., Shipton, Z. K., Wibberley, C. A. J., & Withjack, M. O. (2010). A review of recent developments concerning the structure, mechanics and fluid flow properties of fault zones. *Journal of Structural Geology*, *32*(11), 1557–1575.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, *29*(5), 1189–1232.
- Giorgetti, C., Tesi, T., Scuderi, M. M., & Collettini, C. (2019). Experimental insights into fault reactivation in gouge-filled fault zones. *Journal of Geophysical Research: Solid Earth*, *124*(4), 4189–4204.
- Griffiths, L., Heap, M. J., Baud, P., & Schmittbuhl, J. (2017). Quantification of microcrack characteristics and implications for stiffness and strength of granite. *International Journal of Rock Mechanics and Mining Sciences*, *100*, 138–150.
- Jaeger, J., & Cook, N. (1979). *Fundamentals of rock mechanics*. Chapman and Hall.
- Karimabadi, H., Driscoll, J., Omelchenko, Y. A., & Omidi, N. (2005). A new asynchronous methodology for modeling of physical systems: Breaking the curse of courant condition. *Journal of Computational Physics*, *205*(2), 755–775.
- Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems, pp. 4765–4774.
- Lyakhovskiy, V., Ben-Zion, Y., & Agnon, A. (2001). Earthquake cycle, Fault Zones, and seismicity patterns in a rheologically layered lithosphere. *Journal of Geophysical Research*, *106*, 4103–4120.
- Madden, E. H., Cooke, M. L., & McBeck, J. (2017). Energy budget and propagation of faults via shearing and opening using work optimization. *Journal of Geophysical Research: Solid Earth*, *122*(8), 6757–6772.
- Main, I. G., Bell, A. F., Meredith, P. G., Geiger, S., & Touati, S. (2012). The dilatancy–diffusion hypothesis and earthquake predictability. *Geological Society, London, Special Publications*, *367*(1), 215–230.
- McBeck, J., Aiken, J. M., Ben-Zion, Y., & Renard, F. (2020a). Predicting the proximity to macroscopic failure using local strain populations from dynamic in situ X-ray tomography triaxial compression experiments on rocks. *Earth and Planetary Science Letters*, *543*, 116344.
- McBeck, J., Aiken, J. M., Cordonnier, B., & Renard, F. (2021a). Predicting fracture network development in crystalline rocks. *Pure and Applied Geophysics*. <https://doi.org/10.1007/s00024-021-02908-7>
- McBeck, J., Ben-Zion, Y., & Renard, F. (2020b). The mixology of precursory strain partitioning approaching brittle failure in rocks. *Geophysical Journal International*, *221*(3), 1856–1872.

- McBeck, J., Cordonnier, B., Vinciguerra, S., & Renard, F. (2019). Volumetric and shear strain localization in Mt. Etna Basalt. *Geophysical Research Letters*, *46*, 2425–2433.
- McBeck, J., Mair, K., & Renard, F. (2021b). Decrypting healed fault zones: How gouge production reduces the influence of fault roughness. *Geophysical Journal International*, *225*(2), 759–774.
- Niu, F., Silver, P. G., Daley, T. M., Cheng, X., & Majer, E. L. (2008). Preseismic velocity changes observed from active source monitoring at the Parkfield SAFOD drill site. *Nature*, *454*(7201), 204–208.
- Paterson, M. S., & Wong, T. F. (2005). *Experimental rock deformation—the brittle field*. Springer Science & Business Media.
- Pio Lucente, F., De Gori, P., Margheriti, L., Piccinini, D., Di Bona, M., Chiarabba, C., & Piana Agostinetti, N. (2010). Temporal variation of seismic velocity and anisotropy before the 2009 MW 6.3 L'Aquila earthquake, Italy. *Geology*, *38*(11), 1015–1018.
- Renard, F., McBeck, J., Cordonnier, B., Zheng, X., Kandula, N., Sanchez, J. R., & Dysthe, D. K. (2019). Dynamic in situ three-dimensional imaging and digital volume correlation analysis to quantify strain localization and fracture coalescence in sandstone. *Pure and Applied Geophysics*, *176*(3), 1083–1115.
- Scholz, C. H., Sykes, L. R., & Aggarwal, Y. P. (1973). Earthquake prediction: A physical basis. *Science*, *181*(4102), 803–810.
- Scuderi, M. M., Marone, C., Tinti, E., Di Stefano, G., & Collettini, C. (2016). Precursory changes in seismic velocity for the spectrum of earthquake failure modes. *Nature Geoscience*, *9*(9), 695–700.
- Shreedharan, S., Bolton, D. C., Rivière, J., & Marone, C. (2021a). Machine learning predicts the timing and shear stress evolution of lab earthquakes using active seismic monitoring of fault zone processes. *Journal of Geophysical Research: Solid Earth*, *126*(7), e2020JB021588.
- Shreedharan, S., Bolton, D. C., Rivière, J., & Marone, C. (2021b). Competition between preslip and deviatoric stress modulates precursors for laboratory earthquakes. *Earth and Planetary Science Letters*, *553*, 116623.
- Tapponnier, P., & Brace, W. F. (1976). Development of stress-induced microcracks in Westerly granite. *International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts*, *13*(4), 103–112.
- Tenthorey, E., & Cox, S. F. (2006). Cohesive strengthening of fault zones during the interseismic period: An experimental study. *Journal of Geophysical Research: Solid Earth*, *111*(B9), B09202.
- Whitcomb, J. H., Garmany, J. D., & Anderson, D. L. (1973). Earthquake prediction: Variation of seismic velocities before the San Francisco earthquake. *Science*, *180*(4086), 632–635.

(Received January 12, 2022, revised June 8, 2022, accepted September 11, 2022, Published online September 21, 2022)