







Predicting and following T1 events in dry foams from geometric features

Oskar Tainio , Leevi Viitanen *, Jonatan R. Mac Intyre [†], Mehmet Aydin , Juha Koivisto ,
Antti Puisto , and Mikko Alava

Aalto University, School of Science, Department of Applied Physics, P.O.B 11100, 00076 Aalto, Finland



(Received 14 January 2021; revised 5 May 2021; accepted 28 June 2021; published 15 July 2021)

Machine learning techniques have been recently applied in predicting deformation in amorphous materials. In this study, we extract structural features around liquid film vertices from images of flowing 2D foam and apply a multilayer perceptron to predict local yielding. We evaluate their importance in the description of the T1 events and show that a high level of predictability may be achieved using well-chosen combinations of features as the prediction data. The most relevant features are extracted by performing the predictions separately for isolated sets of features, and these findings are verified using principal component analysis. Using this approach, we determine which properties of the images are most important with regard to the physics of the processes. Our findings indicate that film lengths and angles between the liquid films joining at the vertex are the most important features that predict the local yield events. These two features describe 83% of the yield events. As an application, we extract the statistics of event waiting times from the experiment.

DOI: [10.1103/PhysRevMaterials.5.075601](https://doi.org/10.1103/PhysRevMaterials.5.075601)

I. INTRODUCTION

A mixture of gas bubbles and liquid constitutes a liquid foam at bubble volume fractions surpassing a critical one at the jamming point [1]. Both the surface tension at the liquid-gas interfaces as well as the volume conservation of the gas inside the bubbles create a rigid structure with elastic responses to finite deformation stresses, making foam a yield stress fluid [2,3]. When sheared beyond the yield stress, foams begin to flow, exhibiting shear thinning behavior; their constitutive steady behavior is often described by the Herschel-Bulkley relation [4]. When the foam deformation exhibits plasticity, the shear enforces topological changes to the bubble structure, known as T1 events, as shown in Fig. 1 [5]. The continuum of mechanical properties would correspond to meso/macroscopic elastoplastic models that are widely used to study amorphous solids [6].

In reality, however, the first plastic activity is observed in the form of T1 or yield events already well below the macroscopic yielding [7–9]. The relevance of plastic shear transformation zones, analogous to T1 events, on deformation was originally suggested based on molecular dynamics simulations of metallic glasses [10]. These plastic events have been found to determine the global flow of amorphous materials by forming shear bands, which may be nucleated by single events [11] that subsequently dominate the flow [12,13]. Similarly, in foams, the local dissipation rate has been related to the T1 event activity [14], which would explain the velocity profiles of foams under simple shear deformation [15]. Indeed, observations of foam deformation reveal the effect of single T1 events on the shape of the surrounding bubbles [16]. Another

way of interpreting the same nonlocal changes is via redistribution of stress when a droplet cluster relaxes in a T1 event in emulsions [17,18]. As the structure relaxes in a T1 event, new events may be triggered nearby, resulting in avalanches of T1 events with power law size distributions [19,20].

Due to the crucial role of the yield events in the deformation of foams and other amorphous materials, the ability to predict where and when the yield events occur is critical for anticipating the behavior of foams under mechanical load. Solving complex problems, such as finding structural features that indicate yielding from amorphous materials, is well suited for machine learning tools, which continue to gain popularity in science [21–23].

The present paper builds off our earlier work in Ref. [24], where we identified vertices and bubbles about to yield in T1 events in flowing 2D foam using convolutional neural networks (CNNs) [24]. Although this technique successfully focused on unstable states [Fig. 1(b)] in a snapshot of the foam flow, the complexity of the CNN has hidden the physics behind the T1 events. Here we focus solely on unstable vertices [Fig. 1(b)] that connect four liquid films setting aside the stable vertices with three liquid films. Hence, instead of using a snapshot of the foam flow, we use various structural features extracted from snapshots of the foam as prediction data. We first detect potential T1 events by inspecting tens of thousands of fourfold vertex formations within the foam bubble raft flow. Dynamical data are created by inspecting each frame of a video and extracting explicit T1 event-related features. Next, we feed these fourfold structure variables to a multilayer perceptron (MLP) and successfully predict the occurrence of T1 events. A recent study used a similar approach to explore the correlation of different structural features with future plastic deformations in two-dimensional glasses [25]. However, we measure the correlation indirectly by using prediction metrics of the MLP rather than correlation of physical

*leevi.viitanen@aalto.fi

[†]jonatan.macintyre@aalto.fi




FIG. 1. Schematic of a T1 event. The initial configuration (a) evolves to an unstable fourfold configuration (b). From this unstable state, a new film is created, and a new metastable configuration (c) is reached.

quantities. By choosing essential features, we can steadily increase the accuracy of the model. Henceforth, such a tool can be used to explore T1 dynamics and their effect on the overall transformation of the foam. To showcase this application, we study the waiting times between subsequent T1 events. Our findings indicate that characterizing the ranges of film angles and curvatures represents a major factor for the successful prediction of T1 events.

We first introduce our methods detailing the experimental procedure and how the features are extracted and MLP implemented to allow prediction of T1 events in each frame. Next, we present the main results of the machine learning pipeline model, along with the weight of each feature, the principal component analysis (PCA) component adjusted probability density map, and the detection, extraction, and calculation of several features in the foam with PCA 2-dimensional point clouds. Once the 65-dimensional feature space is reduced to 8, we interpret the black-box AI results as a tangible physical model. Finally, we describe the waiting time distributions and conclude with the discussion and future prospects.

II. METHODS

A. Experimental procedure

The experiments were performed in a circular Hele-Shaw cell, as illustrated in Fig. 2. The same data set has been analyzed previously in Ref. [24] using a convolutional neural network, where we also report the detailed experimental setup. Briefly, the foam enters the cell from an inlet located at the center of the cell and expands toward the edges. The setup is one of the various flow schemes where the edges or obstacles in the flow channel cause shear and T1 events due to that [26–28]. In our geometry, the lack of side edges causes the T1 events to be distributed rather arbitrarily all over the cell, unlike in channels where clear yield zones are detected by bubble ordering caused by the cell walls [29].

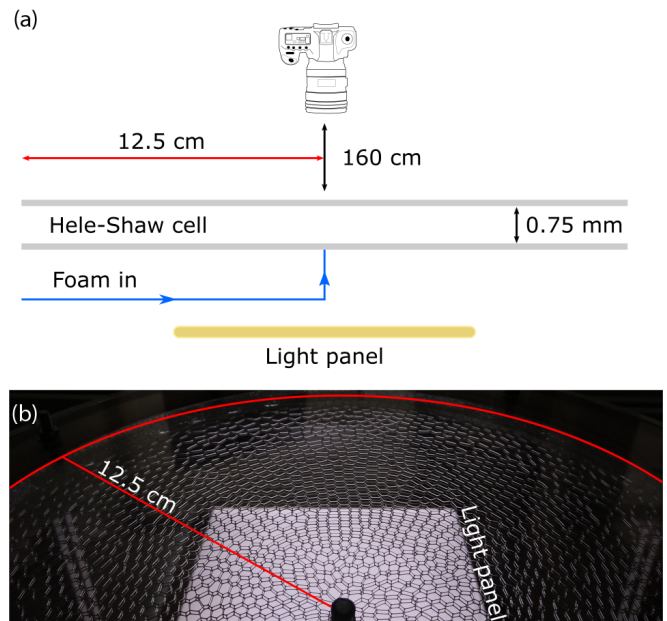


FIG. 2. A two-dimensional circular Hele-Shaw cell creates a radially symmetric expanding flow field. (a) Schematic side view illustrates the experimental setup and how the foam is injected. The foam flow is recorded using a camera placed perpendicular to the cell. (b) A photograph of the Hele-Shaw cell shows a dry foam with typically hexagonal bubbles. The 15-mm inlet pipe is located at the bottom center of the device.

B. Feature extraction

The flowing foam is filmed and the resulting video is interpolated to 13 000 images. Since the background light passes freely through the bubbles, they appear as lighter areas compared to the darker films. Hence, thresholding the gray levels of all the video frames separates the images into black and white pixels representing liquid and gas, respectively (Fig. 3). Skeletonization is used to reduce the widths of the liquid films between bubbles to one pixel. Therefore, each bubble is identified as a region of white pixels in the frames, and in the skeletonized image, the nearest neighbors of a bubble are exactly one pixel apart. For each frame, two lists are generated [Fig. 3(a)]: one contains all the bubbles (blue dots) and the other contains all the neighbors for each bubble (red dashed lines). Comparing the lists between two consecutive frames allows us to detect the T1 events over space and time.

Films are identified as black pixels in the frames. Vertex points are extracted from the film net as the points where three or more films meet [Fig. 3(b)]. We focus on vertices with four films and label each vertex by coordinates (x_v, y_v) as a possible T1 event location. If the point is going to have a T1 event, it is tagged with a number 1 (called positive sample); otherwise, it is tagged with a 0 (negative sample). Once a 4-fold vertex is located, we extract the *end point* (x_e, y_e) and *middle point* (x_m, y_m) from each associated film, as indicated in Fig. 3(b) by the purple and red dots, respectively. Based on these three points, we calculate a series of geometrical features to evaluate the respective weights each one of them has in the T1 event.

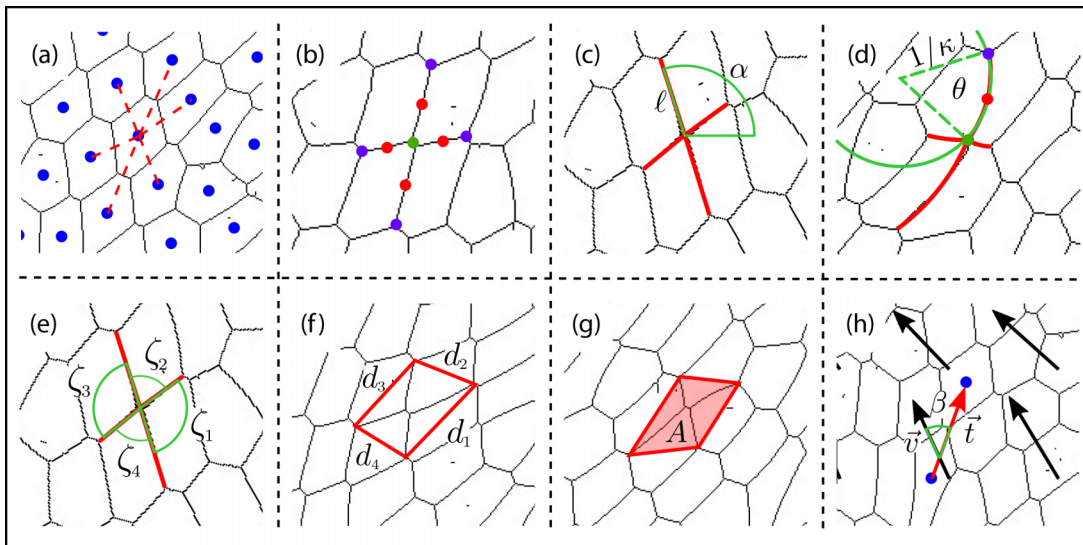


FIG. 3. Detection, extraction, and calculation of several features in the foam flow (note the division into subplots indicated by the dashed lines). (a) Identification of bubbles (blue dots) and neighbors (red dashed lines). The difference between two consecutive frames allows T1 events to be detected. (b) Detection of 4-fold vertices of coordinates (x_v, y_v) (green dot). Eight points are selected from each of the four films: four middle points (red dots) and four end points (purple dots) of coordinates (x_m, y_m) and (x_e, y_e) , respectively. (c) α is the angle between the red vector and an arbitrary coordinate system placed at the vertex point (line $y = y_v$). The vector length is equal to ℓ . (d) The curvature is defined as the osculating circle of three points on the film: vertex (green dot), middle (red dot), and end (purple dot) points. The film length is given by the arclength $s = \theta/\kappa$. (e) Angle ζ between two consecutive films (red lines). (f) Polygon of side d formed using two consecutive end points. (g) The area A of the polygon is marked in red; the perimeter of the polygon is P . (h) The orientation of the T1 event (red arrow) with respect to the velocity field direction (black arrows) is represented by the angle β .

To determine the vertex angles, we define an arbitrary rectangular coordinate system centered at the vertex under consideration (x_v, y_v) . We then calculate the angles α between the position vector of the other vertices and the positive x axis of the coordinate system with the formal definition

$$\alpha = \arctan\left(\frac{y - y_v}{x - x_v}\right), \quad (1)$$

(x, y) being the coordinates of an end or middle point. Figure 3(c) shows the four vectors defined using the film end points (red lines) and one of the four angles (green arc). The length of each vector is called ℓ . In total, eight α angles and eight ℓ lengths (four angles and four lengths both for end and middle points) are defined for each vertex.

Film curvature κ is calculated by determining the osculating circle at the middle point (x_m, y_m) . To calculate this, we use the vertex, end, and middle points as shown in Fig. 3(d). The center of the osculating circle allows us to first calculate the central angle θ and then define the film extension as $s = \theta/\kappa$. As a consequence, four curvatures κ , four angles θ , and eight arclengths s are linked to each vertex point.

Next, we use two consecutive films to define a series of features. Figure 3(e) replots the vectors from vertex to end points in Fig. 3(c). Here, instead of a reference coordinate system, we compute the angle ζ between two consecutive vectors. Note that ζ , unlike α , is independent of the 4-fold vertex orientation, thus allowing us to compute angles ζ based on angles α , but not vice versa. Additionally, we define the polygon side d as the distance between two end points, as shown in Fig. 3(f). Similarly, we obtain another characteristic P for two consecutive middle points. The area A and perimeter P

of the polygon are also computed [Fig. 3(g)]. Finally, the area a that is formed between the film and the chord is calculated.

To finalize the features calculation, we determine the velocity field \vec{v} of the foam flow using the bubble position and average over the 13 000 frames. Figure 3(h) shows the velocity field in black arrows within a smaller area on the frame. For groups of four bubbles involved in possible T1 events, we calculate two vectors between opposite bubbles defined as the vector difference between the centroids. We then take the larger vector \vec{t} and define the T1 orientation β with respect to the foam flow. Figure 3(h) shows the larger vector for 4-fold vertices (red arrow) and the bubble centroids (blue dot). The orientation β with respect to the velocity field is then defined as

$$\beta = \arccos\left(\frac{\vec{v} \cdot \vec{t}}{\|\vec{v}\| \|\vec{t}\|}\right). \quad (2)$$

We then analyze and predict T1 events combining multiple features. Table I summarizes the features used as input to the MLP algorithm. In what follows, if the feature is calculated using middle points instead of end points, a subscript m is added on the feature symbol. Note that for each *vertex* point detected, 65 features are linked to it. The following subsection explains the artificial intelligence method used here.

C. Machine learning pipeline

Supervised learning is applied to study the binary classification problem of predicting T1 events from the local structure of a sample foam. First, we detect all 4-fold vertices [the possibly unstable vertices, Fig. 1(b)] from the 13 000 video frames showing foam flowing in the Hele-Shaw cell. In

TABLE I. List of features computed and defined in Fig. 3.

Feature	Symbol
Angle	α
Chord of circle	ℓ
Curvature	κ
Film extension	θ
Arclength	s
Internal angle	ζ
Orientation	β
Polygon side	d
Polygon area	A
Perimeter	P
Area film and chord	a

total, 172 000 4-fold samples were detected, all of which are labeled with binary values (1,0) indicating whether the cross section precedes a T1 event (1) or not (0). Of the 172 000 detected cases, 15 000 samples are real T1 events; i.e., they are tagged with a number 1. 30 000 samples were chosen for training and testing purposes, half of which represented true T1 events. To optimize the parameters of the MLP, 24 000 of these samples are used as a training set and the remaining 6000 samples are reserved for testing the predictions. In both the training and test sets, an equal amount of positive and negative samples are used. Finally, in order to verify the wider generalization, the model was tested with all extracted cases.

The resulting data were first duplicated into two separate data frames. One half was normalized to zero mean and unit variance. This normalization process was conducted due to the sensitive nature of the latter pipeline estimators. Our tests showed that using the scaled input data improves the prediction scores by a few percent. The cross-validation results shown in Table II (see Appendix) show that without scaling, the prediction results display possible overfitting. However, the cross-validation tests confirm that overfitting is not present when scaling is applied. Scaling is further justified since the average values of different features vary by over three orders of magnitude.

To predict T1 events from the features, we use an MLP algorithm. The analysis is done using the *scikit-learn* Python library [30]. The MLP model is constructed using three hidden layers with the composition of $14 \times 11 \times 8$ neural nodes [31]. The results were confirmed to be robust for the number of nodes per layer and even for the number of layers; however, the chosen values resulted in the highest accuracy (see Appendix). A rectified linear unit was used as the activation function for the model. During the training, the parameters of the neural network are optimized according to batches of data points. The training requires multiple epochs, where each sample is used once during an epoch. The training was continued until the cross-entropy loss function between the subsequent epochs fell below the threshold of 10^{-4} , usually resulting in between 100 and 200 epochs depending on the number of features used as input data. The weights and biases were initialized with random numbers. The weight optimization procedure was conducted using the *Adam optimization method* [32]. In addition to this MLP analysis, we performed PCA for all of the 65 extracted structural features using the




FIG. 4. Probability density for each binary class according to the PCA components. (a) Probability density for positive samples, T1 events. (b) Probability density for negative samples, no T1 event. (c) The separation density [i.e., the difference between the PDFs of panels (a) and (b)] reveals the values of components 1 and 2 that separate the positive samples from the negative ones. The x and y axes are the values of first and second PCA components. The axes and their values are the two principle components produced by the PCA method.

scikit-learn library [30] to reduce the dimensionality of the feature space to only two orthonormal components. The two resulting features were also tested with the MLP.

The performance of the individual models was evaluated by examining the ratio of true positives and true negatives to false positives and false negatives with model training and test accuracies. The final model evaluation was conducted using cross validation with five separate iterations and F1 scores. To check for possible overfitting or bias in the data, we also performed cross-validation tests.

III. RESULTS

A. Evaluating single features

As a reference control experiment, the model was run with a mixture of both 3-fold and 4-fold vertex samples. A static value of -1 was used in the place of the fourth feature column values, representing the fourth film for each 3-fold feature. The model achieved a training score of 1.0 and a test score of 0.997, similar to the results where T1 events were predicted from images [24]. This high degree of predictability was achieved with only small training and test sets containing 500 and 350 samples, respectively. Thus, the MLP model has the capability of distinguishing a 4-fold cross section of bubbles representing a potential T1 event from a core foam structure. This result was achieved with both unmodified and scaled data. Since special cases of 3-fold and 5-fold cross joints can be clearly classified, they were neglected in order to focus on the complex 4-fold data. Further performance results can be seen in the Appendix.

The dimensionality of the scaled training data was reduced using PCA resulting in a transformation matrix which contains the weights of each physical feature. According to these vector component significance values, the following features contribute the most to the values of the two principal components shown in Fig. 4: angles α , chord length ℓ , arclengths s , and angles between two films ζ . By utilizing principal-component-derived values, it is possible to visualize differences of T1 event and non-T1 event samples, even within environments containing only 4-fold liquid films. The

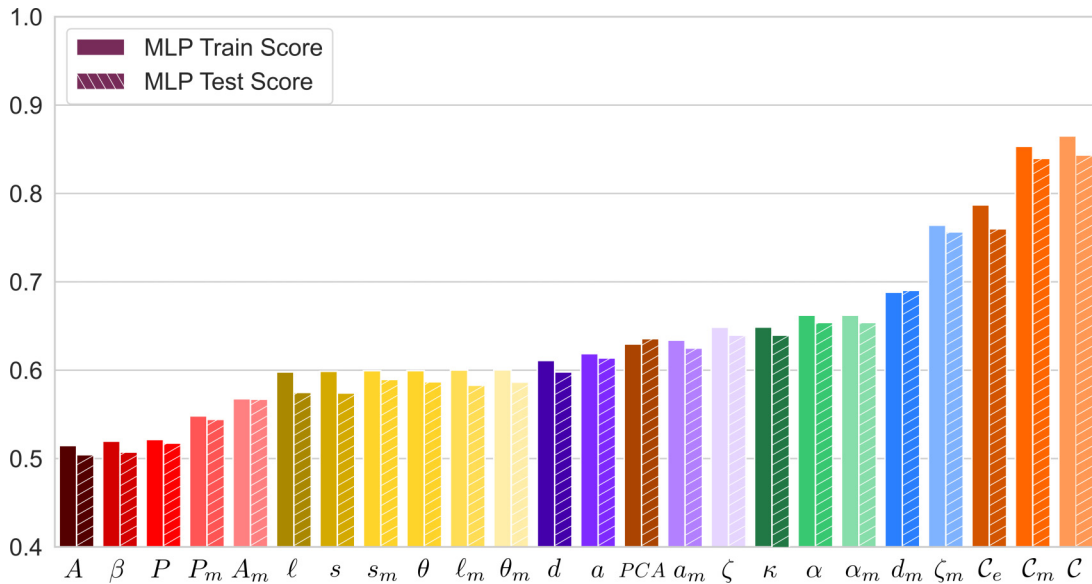


FIG. 5. Bar chart depicting MLP training (solid) and test (pattern) scores for each extracted feature. Features obtaining the best scores have been placed on the right, whereas the worst-scoring features are placed on the left. Subscript m stands for middle point used for feature extraction. The combination of all features extracted using end and middle points is shown with C_e and C_m , respectively, and all features together with C . The score using the PCA components is labeled as PCA.

probability densities of principle components for positive (T1) and negative (no T1) samples are represented as heat maps in Figs. 4(a) and 4(b), where the axis represent the value of the principle components obtained from the PCA transformation. Figure 4(c) shows that the clouds of true T1 events are concentrated in separate areas compared to the cloud of false T1 events.

The importance of each feature to the T1 predictability was measured by fitting and testing the MLP model with reduced input data sets. Here, the input contained only data associated with each individual feature as listed in Table I. Training and test scores are shown in Fig. 5, where every derived feature has been ranked from worst to best based on the obtained MLP model scores. It is clearly seen that features dropped out by utilizing PCA performed the worst. These include variables of both end and middle point types. Model performance was poor for polygon perimeter P and area A , and T1 orientation β with test score close to 0.5 comparable to a coin toss. The best scores were obtained by utilizing angles α_m , angles between consecutive films ζ_m , and polygon side d_m using middle points. The model having only 36 of the 65 input variables (see C_m and C_e in Fig. 5) performs nearly as well as the model with all possible parameters. As a result, an efficient T1 prediction model can be achieved by implementing a fraction of all possible descriptive features within the foam network. This result is supported by comparison with the weight coefficients of the PCA components.

Finally, according to the PCA analysis and single feature model performance, we combined all features together into an input matrix using two components and fed this to the MLP model. In this case, the performance of the MLP decreases by 25% compared to the best case with all features, as shown in the Fig. 5 bar labeled as PCA . This implies that the number of input features for the MLP can be reduced dramatically

without significantly compromising the predictability of T1 events. With this in mind, we next start to explore the model performance with restricted input data.

B. Evaluating features in pairs

Until this point, the T1 events were predicted either using all of the structural features or only one type of feature. However, the MLP can take any kind of set of data as an input. Next, we pass the input features to the MLP by pairs of features listed in Table I and predict the yielding. It could be expected that the pairs of features with the highest individual scores in Fig. 5 yield the best combined scores. This is not the case, however, as shown in Fig. 6(a) (blue square). A combination of the angles between films ζ_m and film length l_m achieves the best prediction score of 0.83. Interestingly, although angles ζ_m result in the highest predictability alone, the feature l_m alone has a prediction score of 58%. Another notable example from prediction scores is the angle orientations α and α_m , which do not improve each others' score at all.

The nontrivial prediction performance can be attributed to the dependence of different features on each other. For example, the angles α and α_m contain almost the same information, so using both as MLP input does not add any useful information compared to only using one. On the contrary, the features ζ_m and l_m are completely independent; that is, the angles between films do not give any information about film length. This indicates that the predictability based on l_m (8 percentage points better than a coin toss) can be simply added to the score of predictions with ζ_m when the two are combined.

Next, we take the two best performing features, the angles ζ_m and l_m , for the four adjoint films in a vertex (8 features total) and perform a dimensionality reduction using PCA visualized in Fig. 6(b). The colored areas represent the decision

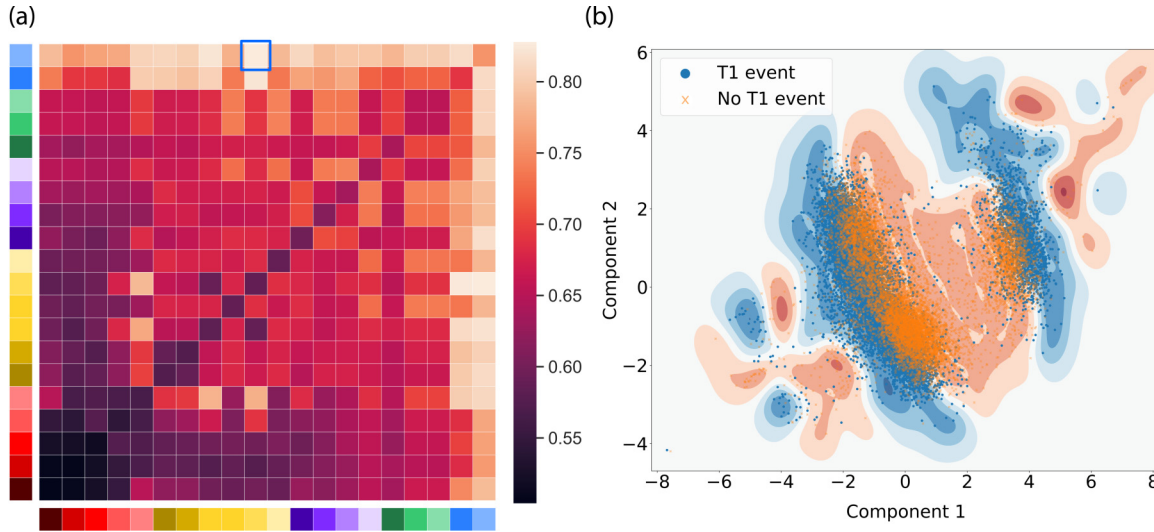


FIG. 6. (a) The images show matrix representation of T1 prediction scores using features as input pairs (vertical axis is identical to horizontal axis). Diagonal values are scores of each single feature and labels follow the color code as in Fig. 5. The best score (0.83) is marked with a blue frame for the combination ζ_m and ℓ_m . (b) To visualize the prediction data, we reduce the dimensionality of input data with features ζ_m and ℓ_m from 8 to 2 using PCA, and plot the samples that precede a T1 event (blue) and those that do not precede a T1 event (red). We also divide the space using the support vector classification in two areas where we predict T1 to occur (blue) or not (red). This division achieves 76% prediction score.

boundary given by the support vector classification, with blue areas indicating samples preceding a T1 event. The division presented in Fig. 6(b) achieves a 76% prediction score, which is lower than the score in Fig. 6(a) for ζ_m and ℓ_m because of the coarse-grained decision boundary. The evaluation of features in pairs suggests that the MLP can learn the important local mechanical properties at the vertex from the given data and deduce the stability of the vertex. The mechanical information of the foam structure—such as local stresses—is encoded in the geometry descriptors and their correlations.

C. Converting black-box AI to physics

A closer inspection of the essential data produced by our automated data collection and processing algorithms described above reveals the following: (i) From the 65-dimensional feature space, the two best descriptors are $\zeta_{m,i}$ and $\ell_{m,i}$, $i = 1 \dots 4$, and (ii) the T1 events can form tilted elliptic-like clusters seen with blue color in Fig. 6(b). Hence, we create two elliptical decision boundaries to classify the nodes as ones that produce T1 events, and those that do not. The nodes that produce the T1 event reside within either of the ellipses, while the nodes that do not produce the T1 event fall outside of both ellipses. For this, the two-variable system produced by the PCA algorithm [similar to Fig. 6(b)] is needed,

$$x = 0.40(\zeta'_{m,1} + \zeta'_{m,3}) - 0.43(\zeta'_{m,2} + \zeta'_{m,4}) - 0.28(\ell'_{m,1} + \ell'_{m,2} + \ell'_{m,3} + \ell'_{m,4}), \quad (3)$$

$$y = 0.17(\zeta'_{m,1} + 2\zeta'_{m,3}) - 0.29(\zeta'_{m,2} + \zeta'_{m,4}) + 0.43\ell'_{m,1} - 0.36\ell'_{m,2} + 0.33\ell'_{m,3} - 0.51\ell'_{m,4}, \quad (4)$$

where $\zeta'_{m,1}$ is the largest angle between neighboring films normalized to zero mean and unit variance. A similar normalization process is repeated for other angles between films $\zeta'_{m,i}$ and chord (film) lengths $\ell'_{m,i}$. The values of constants are rounded to two decimals for convenience. Grouping even and odd $\zeta'_{m,i}$ together indicates that the opposite angles between films should be alike, in a “squeezed x”-like geometry. The asymmetry in chord lengths in $\ell'_{m,i}$ indicates that the “x” is sheared. This typical asymmetry in the vertex geometry is also visible in Figs. 8 and 9.

The two elliptical decision boundaries separating the vertices that produce T1 and those that do not in the coordinate system described by (3) and (4) are determined by using a brute force table seek algorithm with the following constrictions: (i) the probability of T1 events inside the ellipses matches the probability of no-T1s outside the ellipses, i.e., equal probabilities for true positives and true negatives, and the (ii) tilt and (iii) aspect ratios of the ellipses are the same. This is motivated by the symmetry seen in Fig. 6(b). The result is two ellipses with a tilt of $\alpha = -35$ degrees, an aspect ratio $dy/dx = 0.6$, and center points located in $(x_1, y_1) = (-2.0, 0.4)$ and $(x_2, y_2) = (4.0, -0.7)$ with radii $r_1 = 1.3$ and $r_2 = 0.6$. With this information, the coordinate system can be manipulated to a simpler form,

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} x \\ 0.6y \end{pmatrix}, \quad (5)$$

$$1.3^2 = (x' + 2.0)^2 + (y' - 0.4)^2, \quad (6)$$

$$0.6^2 = (x' - 4.0)^2 + (y' + 0.7)^2, \quad (7)$$

where we have essentially two center points and a radius, shown in Fig. 7. Here, the circular decision boundary divides the events such that both the true negatives and true positives




FIG. 7. The two elliptical decision boundaries (green), seen as circular in this coordinate system, capture 71% of the T1 events (blue, true positives) and exclude 71% of the non-T1 events (red, true negatives).

have 71% probability; i.e., 71% of blue crosses are inside the green circles and 71% of the red crosses are outside the circles. The physical interpretation of the black-box AI is that the probability of finding a T1 event increases as one moves closer to the center of either of the circles. With this information, we can draw a typical T1-event-producing node shown in Fig. 8 in the footsteps of Fig. 1. Here, 1000 vertices from the experimental data are drawn with opaque coloring, where the intensity of color increases with the probability of encountering a film in that location. The colors indicate different films, where yellow is the closest to the horizontal axis set by the camera, and the colors change counterclockwise for other films. The diagonal panels of Figs. 8(a) and 8(d) represent 71% of the data (true positives and negatives) while

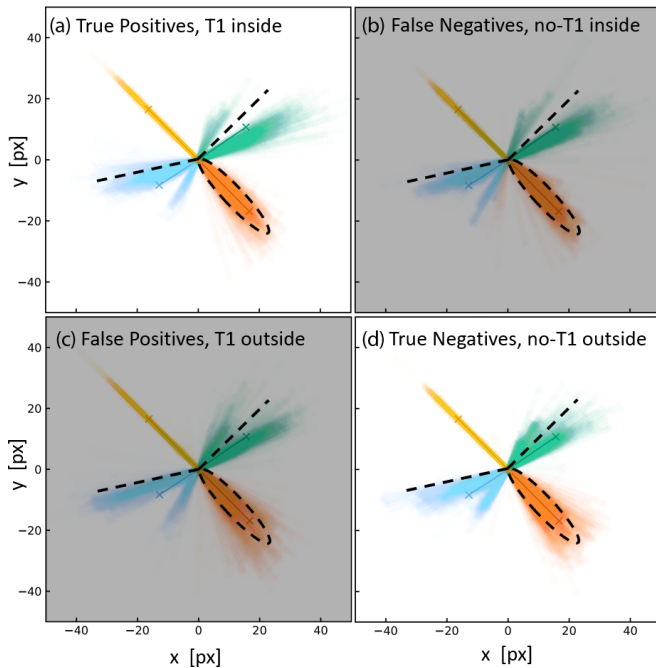


FIG. 8. The four color-coded films are drawn for each possible case of the AI classifier: true (white background) or false (gray background), positive (left) or negative (right). Each panel contains 1000 films drawn on top of each other. The repeating dashed construction is a guide to visualize the differences. The dark colored lines with x symbol are the average location and length of all the films. The labels inside and outside refer to whether the vertex at hand falls inside or outside the elliptical decision boundary.

the off-diagonal panels of Figs. 8(b) and 8(c) represent 29% of the data (falsely categorized vertices). Thus, it is meaningful to compare the diagonal panels of Figs. 8(a) and 8(d) and treat the off-diagonal elements as errors of the method.

The dashed construction is the same for all panels, highlighting the differences. One of the film orientations (for green films) is missing in Figs. 8(a) and 8(b); some of the (blue) films are above the dashed line, indicating that nodes with this kind of film structure produce a T1 event. On the other hand, if the (red) films are outside the dashed ellipse, the node does not produce the T1 event. Hence, when a node with this film orientation is encountered, it is stable and does not produce a T1 event with 71% probability. The asymmetry between the green and blue sets of lines indicates that the “squeezed x” geometry preferred by the T1 events prefers one direction over the other. This is most likely due to the camera orientation (laboratory coordinates) relative to the flow direction.

This demonstrates the possibility of creating an automated algorithm that highlights the main differences of a high 65-dimensional data cube and provides the result in a form that is easily interpreted, either as a graphical representation or in the form of an equation. If the measured film construction approaches the “squeezed x” construction visually (Fig. 8) or the data falls within one of the ellipses described in Eqs. (3)–(7), there is a 71% chance of a T1 event. In fact, the equations describe the energy landscape of the system $E = -(x - x_1)^2 - (y - y_1)^2$, which maximizes the elastic energy of the system at the center point of the ellipses. In such geometry, it is more favorable to create a plastic yielding event than to further deform the bubbles. For illustrative purposes, one can simply use the angles without the lengths to obtain a simple formula for energy landscape

$$E = -(\zeta_{m,1} - 112^\circ)(\zeta_{m,2} - 68^\circ), \quad (8)$$

where the two constants are the angles (in degrees) between yellow and green and yellow and blue films illustrated in Fig. 9. A more detailed energy landscape could be obtained by expanding Eqs. (3)–(7) leading to, for example, an energy landscape near the first ellipse (6) as

$$\begin{aligned} E = & -\sum (c_i \ell'_{m,i} + d_i \zeta'_{m,i})^2 + (e_i \ell'_{m,i} + f_i \zeta'_{m,i})^2 \\ = & -0.01(3.53 + \ell'_{m,1} - 2.78\ell'_{m,2} + 0.52\ell'_{m,3} - 3.50\ell'_{m,4} \\ & + 2.33\zeta'_{m,1} - 3.01\zeta'_{m,2} + 3.14\zeta'_{m,3} - 3.01\zeta'_{m,4})^2 \\ & - 0.13(5.50 + \ell'_{m,1} + 0.44\ell'_{m,2} + 0.93\ell'_{m,3} + 0.34\ell'_{m,4} \\ & - 0.87\zeta'_{m,1} + 0.86\zeta'_{m,2} - 0.75\zeta'_{m,3} + 0.86\zeta'_{m,4})^2, \quad (9) \end{aligned}$$

where the c_i , d_i , e_i , and f_i are obtained from the PCA analysis. Here, the values and signs of the constants reveal the importance of the films. The values of angles are understood as the violation of the plateau principle; i.e., energy minimum is achieved when the angle between the films is 120 degrees. The observation here is that if we include the film lengths, we can improve the accuracy of the T1 prediction by 8% (Fig. 5) and we can say that the lengths of the films obey the same symmetry as the angles. The opposite films in the vertex ($\ell_{m,1} \approx \ell_{m,3}$ and $\ell_{m,2} \approx \ell_{m,4}$) have similar weights, indicating that their lengths are the same and the node is sheared. This is also visible in Fig. 9.




FIG. 9. Angles between films that produce the T1 event with highest probability.

D. Waiting time distribution

Finally, we examine the inter-event arrival times and the effect of the region-of-interest (ROI) size on the observed statistics. The ROI is square shaped and its center is fixed halfway between the inlet and the cell edge. Vertices that yield a T1 event inside the ROI form a time series where each point represents a time instance when the T1 event was detected inside the ROI. The waiting times τ are differences in time between subsequent points of the time series. Figure 10 shows the waiting time distribution $P(\tau)$ for ROIs $L \times L$ from $L = 1$ cm to $L = 5$ cm. The data are fitted to a




FIG. 10. The waiting time distribution is a power law with a ROI-scale dependent exponential cutoff, $P(\tau) = A\tau^{-1}e^{-\lambda\tau}$. The data for small ROIs exhibit a power law regime with exponent -1 , while data for the largest ones are dominated by the exponential cutoff. The inset shows a scaling $\lambda \propto L^3$ indicating that in addition to a typical size dependence of L^2 the radial expansion plays a role in the scaling.

power law with an exponential cutoff $P(\tau) = A\tau^{-1}\exp(-\lambda\tau)$, where A and λ are fitting parameters and the exponent of the power law is fixed to -1 . The form of the gamma distribution $P(\tau)$ originates from the interplay between event triggering avalanche dynamics and the system's intrinsic background event rate [33,34]. The linear power law region indicates event clustering, most evident for the smallest ROI (black dots). As the ROI size increases, the power law region narrows down due to the cutoff. This indicates that the event waiting times result from observing noninteracting Poisson processes in subareas of the larger ROI. A naive scaling would thus scale with ROI area as $\lambda = BL^2$. The actual $\lambda = BL^3$ (inset of Fig. 10) results from another scale L entering through radial slowing of the foam flow—thus also T1 production—in the larger ROIs. Since waiting times become uncorrelated with large ROI size, there is a cooperativity length of a few bubble diameters, comparable to that observed in particulate systems [35].

IV. CONCLUSIONS

The purpose of this study was to apply machine learning methods to image sets of flowing foams in order to understand local yield events. To this end, we extracted structural features and used machine learning methods to try to predict such local yield events from the structure. We achieved a high level of prediction accuracy by inputting all of our structural features in a multilayer perceptron classifier, meaning that the most relevant features have been found. The application of such ML techniques proceeds by using the MLP classifier for selected sets of the features. Using this approach, we find that the most relevant features are the angles between the liquid films joining at the vertex ζ at any location where this is attempted. Likewise, we may conclude that the most irrelevant features include the orientation of cluster β and the polygon area A .

It is reasonable that the angles between films are a strong indicator of T1 events. In a stable state, the films form 120 degree angles, but when two vertices become close and form a 4-fold vertex, the angles decrease such that the total sum of the four angles is 360 degrees. Many of the other well-performing features, such as film orientations and polygon side lengths, are highly correlated with the angle between films; hence adding them leads to no or modest improvements to the predictions.

The accuracy of prediction reached by this approach is about 80%. The result leaves some space for further development as around one fifth of the predictions are incorrect. As a methodological approach to improve the situation, one could test different machine learning algorithms for the problem using the same input data. For instance random forests could provide a potential alternative. A more physical approach would be to refine the input data set. This could be done by determining new features from the experiment not related to the structure or the affine motion of the foam, but instead on the nonlocal effects [14], such as the stress redistribution occurring after previous yield events [6,17,18]. However, defining such features might emerge ambiguous and difficult to justify. It may also be such that the yielding—the T1—events are also affected by random, thermal noise, rendering such improvement impossible.

The results are verified with a PCA analysis, where the features are projected to two principal orthogonal components. These components align with the success of MLP predictions, and similarly imply that angles and film lengths are the most relevant components. Such tools allow information about the dynamics of T1s to be extracted. We demonstrate this by considering the T1 waiting time distributions in varying ROIs as a function of the ROI size. These distributions have a broad, power law character with an exponential cutoff. We note that the waiting time exponent becomes clear with a small enough ROI having a value close to unity. The cutoff scales with the ROI scale as its third power, which likely results from natural geometric considerations. The above exponent value remains an empirical observation without a theoretical explanation.

Although the possibility of improvement remains, our analysis shows that machine learning may be exploited to detect the essential parameters related to plastic deformation from the structural features it receives as input. This makes it an attractive tool to analyze foam dynamics, analogous to amorphous materials plasticity. This conclusion is thus in line with other applications of machine learning techniques to extract predictions of deformation from structure in amorphous glass simulations [25]. A further line of research could be to search for collective dynamics among T1 events, as a generalization of our waiting time analysis. Likewise, more complex foams such as particle-laden or Pickering foams present another exciting avenue for this approach.

ACKNOWLEDGMENTS

J.K. and A.P. acknowledge funding from Academy of Finland (Grants No. 308235 and No. 278367), Business Finland (Grant No. 211715), and Aalto University (Grant No. 974109903), as well as the Aalto Science IT project for computational resources. L.V. acknowledges funding from the Vilho, Yrjö and Kalle Väisälä Foundation via personal grants, and Academy of Finland (Grant No. 278367).

APPENDIX: MACHINE LEARNING PERFORMANCE

The multilayer perceptron (MLP) algorithm used was extensively analyzed. The best performance was achieved for a three-layer architecture $14 \times 11 \times 8$. We used a cross validation with 5 groups and F1 score

$$F1 = \frac{TP}{TP + (FN + FP)/2}, \tag{A1}$$

TABLE II. Performance of the architecture employed. RL: Randomly labeled.

Data set	No. samples	Scaled	RL	Train score	Test score	No. true positive	No. true negative	F1	Cross validation
3-fold+4-fold	6000	Yes	No	1	0.998	600	5880	0.967	0.985
	6000	Yes	Yes	0.5	0.5	1500	1500	0.111	0.55
4-fold	6000	Yes	No	0.853	0.840	2631	2407	0.845	0.82
	172000	Yes	No	0.860	0.832			0.487	0.918
	6000	No	No	0.814	0.812	2339	2532	0.806	0.59
	172000	No	No	0.865	0.817			0.471	0.916
	6000	Yes	Yes	0.555	0.508	1653	1395	0.528	0.5

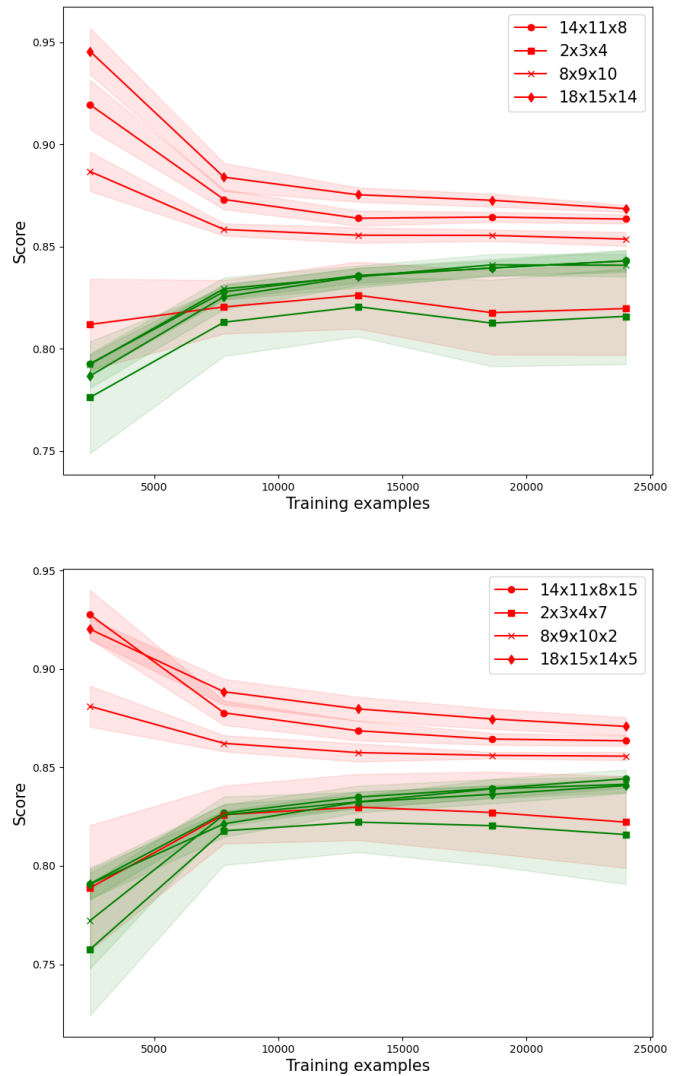


FIG. 11. Learning curve for the architecture employed. As the number of samples increases, the standard deviation decreases (shadow). The small gap between both curves indicates that the model is not overfitted. When the number of layers and nodes increases, the score does not change significantly.

where TP is true positives, FP false positives, and FN false negative. We also randomize the T1 label to further explore the accuracy in the prediction. Table II summarizes our analysis. The scaled method corresponds to zeroing the mean and scaling samples to unit variance, i.e., $(x - u)/s$, where x is

the data, and its mean and standard deviation are u and s , respectively.

Figure 11 shows the learning curve, i.e., the training and test scores for an increasing number of samples. The scores

are high for all the different architectures plotted, but the one displayed here produces less variance and better generalizes the model, as seen by the gap size between training and test scores.

-
- [1] A. J. Liu and S. R. Nagel, The jamming transition and the marginally jammed solid, *Annu. Rev. Condens. Matter Phys.* **1**, 347 (2010).
- [2] G. Ovarlez, S. Cohen-Addad, K. Krishan, J. Goyon, and P. Coussot, On the existence of a simple yield stress fluid behavior, *J. Non-Newtonian Fluid Mech.* **193**, 68 (2013).
- [3] D. Bonn, M. M. Denn, L. Berthier, T. Divoux, and S. Manneville, Yield stress materials in soft condensed matter, *Rev. Mod. Phys.* **89**, 035005 (2017).
- [4] S. Cohen-Addad, R. Höhler, and O. Pitois, Flow in foams and flowing foams, *Annu. Rev. Fluid Mech.* **45**, 241 (2013).
- [5] D. L. Weaire and S. Hutzler, *The Physics of Foams* (Oxford University Press, New York, 2001).
- [6] A. Nicolas, E. E. Ferrero, K. Martens, and J.-L. Barrat, Deformation and flow of amorphous solids: Insights from elastoplastic models, *Rev. Mod. Phys.* **90**, 045006 (2018).
- [7] M. Dennin, Statistics of bubble rearrangements in a slowly sheared two-dimensional foam, *Phys. Rev. E* **70**, 041406 (2004).
- [8] E. N’Gouamba, J. Goyon, and P. Coussot, Elastoplastic behavior of yield stress fluids, *Phys. Rev. Fluids* **4**, 123301 (2019).
- [9] B. Shang, P. Guan, and J.-L. Barrat, Elastic avalanches reveal marginal behavior in amorphous solids, *Proc. Natl. Acad. Sci. USA* **117**, 86 (2020).
- [10] M. L. Falk and J. S. Langer, Dynamics of viscoplastic deformation in amorphous solids, *Phys. Rev. E* **57**, 7192 (1998).
- [11] A. Barbot, M. Lerbinger, A. Lemaitre, D. Vandembroucq, and S. Patinet, Rejuvenation and shear banding in model amorphous solids, *Phys. Rev. E* **101**, 033001 (2020).
- [12] A. Kabla and G. Debrégeas, Local Stress Relaxation and Shear Banding in a Dry Foam under Shear, *Phys. Rev. Lett.* **90**, 258303 (2003).
- [13] A. Ghosh, Z. Budrikis, V. Chikkadi, A. L. Sellar, S. Zapperi, and P. Schall, Direct Observation of Percolation in the Yielding Transition of Colloidal Glasses, *Phys. Rev. Lett.* **118**, 148001 (2017).
- [14] J. Goyon, A. Colin, G. Ovarlez, A. Ajdari, and L. Bocquet, Spatial cooperativity in soft glassy flows, *Nature (London)* **454**, 84 (2008).
- [15] G. Katgert, B. P. Tighe, M. E. Möbius, and M. van Hecke, Couette flow of two-dimensional foams, *Europhys. Lett.* **90**, 54002 (2010).
- [16] F. Elias, C. Flament, J. Glazier, F. Graner, and Y. Jiang, Foams out of stable equilibrium: Cell elongation and side swapping, *Philos. Mag. B* **79**, 729 (1999).
- [17] K. W. Desmond and E. R. Weeks, Measurement of Stress Redistribution in Flowing Emulsions, *Phys. Rev. Lett.* **115**, 098302 (2015).
- [18] D. Chen, K. W. Desmond, and E. R. Weeks, Experimental observation of local rearrangements in dense quasi-two-dimensional emulsion flow, *Phys. Rev. E* **91**, 062306 (2015).
- [19] Y. Jiang, P. J. Swart, A. Saxena, M. Asipauskas, and J. A. Glazier, Hysteresis and avalanches in two-dimensional foam rheology simulations, *Phys. Rev. E* **59**, 5819 (1999).
- [20] S. Tewari, D. Schiemann, D. J. Durian, C. M. Knobler, S. A. Langer, and A. J. Liu, Statistics of shear-induced rearrangements in a two-dimensional model foam, *Phys. Rev. E* **60**, 4385 (1999).
- [21] E. D. Cubuk, S. S. Schoenholz, J. M. Rieser, B. D. Malone, J. Rottler, D. J. Durian, E. Kaxiras, and A. J. Liu, Identifying Structural Flow Defects in Disordered Solids using Machine-Learning Methods, *Phys. Rev. Lett.* **114**, 108001 (2015).
- [22] M. Harrington, A. J. Liu, and D. J. Durian, Machine learning characterization of structural defects in amorphous packings of dimers and ellipses, *Phys. Rev. E* **99**, 022903 (2019).
- [23] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. Schoenholz, A. Obika, A. Nelson, T. Back, D. Hassabis *et al.*, Unveiling the predictive power of static structure in glassy systems, *Nat. Phys.* **16**, 448 (2020).
- [24] L. Viitanen, J. R. Mac Intyre, J. Koivisto, A. Puisto, and M. Alava, Machine learning and predicting the time-dependent dynamics of local yielding in dry foams, *Phys. Rev. Research* **2**, 023338 (2020).
- [25] D. Richard, M. Ozawa, S. Patinet, E. Stanifer, B. Shang, S. A. Ridout, B. Xu, G. Zhang, P. K. Morse, J. L. Barrat, L. Berthier, M. L. Falk, P. Guan, A. J. Liu, K. Martens, S. Sastry, D. Vandembroucq, E. Lerner, and M. L. Manning, Predicting plasticity in disordered solids from structural indicators, *Phys. Rev. Mater.* **4**, 113609 (2020).
- [26] B. Dollet and C. Bocher, Flow of foam through a convergent channel, *Eur. Phys. J. E* **38**, 1 (2015).
- [27] T. Chevalier, J. Koivisto, N. Shmakova, M. J. Alava, A. Puisto, C. Raufaste, and S. Santucci, Foam flows through a local constriction, *J. Phys.: Conf. Ser.* **925**, 012025 (2017).
- [28] D. Vecchiolla and S. L. Biswal, Dislocation mechanisms in the plastic deformation of monodisperse wet foams within an expansion-contraction microfluidic geometry, *Soft Matter* **15**, 6207 (2019).
- [29] Y. Gai, A. Bick, and S. K. Y. Tang, Timescale and spatial distribution of local plastic events in a two-dimensional microfluidic crystal, *Phys. Rev. Fluids* **4**, 014201 (2019).
- [30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [31] J. Schmidhuber, Deep learning in neural networks: An overview, *Neural Networks* **61**, 85 (2015).
- [32] D. Kingma and J. Lei Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980v9](https://arxiv.org/abs/1412.6980v9).

- [33] J. Baró, Á. Corral, X. Illa, A. Planes, E. K. H. Salje, W. Schranz, D. E. Soto-Parra, and E. Vives, Statistical Similarity between the Compression of a Porous Material and Earthquakes, *Phys. Rev. Lett.* **110**, 088702 (2013).
- [34] T. Mäkinen, Collective phenomena in deformation, Ph.D. thesis, Aalto University, 2020.
- [35] E. D. Cubuk *et al.*, Structure-property relationships from universal signatures of plasticity in disordered solids, *Science* **358**, 1033 (2017).