

Machine learning and predicting the time-dependent dynamics of local yielding in dry foams

Leevi Viitanen , Jonatan R. Mac Intyre , Juha Koivisto, Antti Puisto, and Mikko Alava
Aalto University, School of Science, Department of Applied Physics, P.O. Box 11100, 00076 Aalto, Finland



(Received 17 September 2019; accepted 24 May 2020; published 15 June 2020)

The yielding of dry foams is enabled by small elementary yield events on the bubble scale, “T1”s. We study the large-scale detection of these in an expanding two-dimensional (2D) flow geometry using artificial intelligence (AI) and nearest neighbor analysis. A good level of accuracy is reached by the AI approach using only a single frame, with the maximum score for vertex centered images highlighting the important role the vertices play in the local yielding of foams. We study the predictability of T1s ahead of time and show that this is possible on a timescale related to the waiting time statistics of T1s in local neighborhoods. The local T1 event predictability development is asymmetric in time, and measures the variation of the local property to yielding and similarly the existence of a relaxation timescale after local yielding.

DOI: [10.1103/PhysRevResearch.2.023338](https://doi.org/10.1103/PhysRevResearch.2.023338)

I. INTRODUCTION

Dry foams are assemblies of gas pockets separated by thin films of liquid, forming a connected polygonal film structure [1] at a configuration globally minimizing the surface energy [2]. A finite external stress is needed to make foams flow and, due to the absence of thixotropy, foams are usually rheologically characterized as simple yield stress fluids [3]. The steady-state, global flow curves of foams are considered to be typical examples of Herschel-Bulkley fluid-like behavior [4,5].

At bubble scale, the viscoplastic flow of dry foams is enabled by small elementary topological yield events referred to as T1’s and T2’s, analogous, for instance, to shear transformation zones (STZs) in amorphous solids [6,7]. The T2 events involve the disappearance of three-sided bubbles, while the T1 events refer to a neighbor swap between four bubbles. Both events enable the system to jump from one metastable surface energy minimum to another, including a local relaxation of the stored elastic energy of the foam [8,9]. T2 events require either film breakup or gas exchange to occur, making them less frequent compared to T1s in systems under continuous deformation. A quasi-2D setting or simplifying geometry, is often used to study the flow of foams and T1s therein [10–14]. The T1 events can be either reversible or nonreversible depending on the geometric configuration and the stress direction [15,16].

In the macroscopic flow and yielding of foams, the plastic deformation arises from a complex dynamics of T1 events with localization and clustering [17–19]. Here, we study T1 events from large-scale statistics. We first detect tens of

thousands of events from bubble raft flow dynamical data by comparing two sequential frames using a nearest neighbor detection algorithm [10]. Second, we feed single frames to a convolutional neural network (CNN) and successfully predict the T1 events by the film structure and extract the essential features. We use this tool to explore T1 dynamics and the role of the film structure by varying the region of interest (ROI). The relation of the T1 dynamics to macroscopic flow have been studied and interestingly, the time-dependent deformation tensor [20,21] and the strain rate [22], both dependent on the local velocity, were found to correlate with the T1 frequency, while it has been suggested there is no correlation between the snapshot of the film structure and the T1 frequency [10]. This would imply T1 detection based on a single structure snapshot should fail most of the time.

We find that T1 events can be predicted by the localized changes in film vertices, while the bubble shape itself does not contain the same information. The local predictability evolution is shown by time-dependent AI predictions to be asymmetric in time, and we interpret the results as a measure of the decrease in local yield stresses and as a manifestation of the inherent local relaxation after a T1. The foam structure exhibits weaker and stronger spots, be it due to the network structure or film properties [23]. In amorphous solids soft spots have been studied using measures for nonaffine deformation [24], Voronoi cell anisotropy [25], and machine learning tools [26,27] and the distribution of local yield stresses is an inherent property of plasticity models [28], as is the local relaxation dynamics after a yield event [29].

We compare two methods of identifying T1 events from a video of a two-dimensional flow. First, we detect tens of thousands of events from dynamical data by comparing two sequential frames using well-established detection methods. Second, we feed single frames to neural network and successfully predict the T1 events by the snapshot of the film structure. Our main findings show that T1 events are correlated with the localized change in film vertices, while the bubble perimeter itself does not contain similar information.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

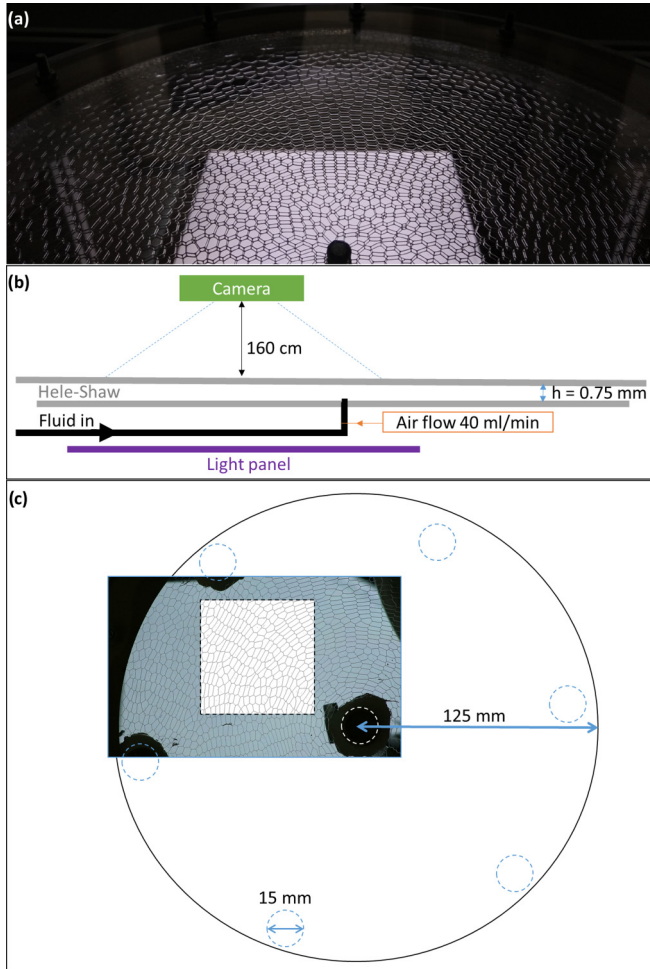


FIG. 1. Radially symmetric 2D Hele-Shaw cell creates an expanding flow field. (a) An angled photograph of the Hele-Shaw cell shows dry foam with typically hexagonal bubbles. The 15-mm inlet hose is located at the bottom center of the device. (b) Schematic side view illustrates how the fluid and air are foamed at the inlet. (c) The schematic drawing of the top view of the circular Hele-Shaw cell is overlaid by a single raw image from the video stream. The dashed line illustrates the analysis area that is skeletonized to a binary image. The inlet at the center and outlet holes have diameter $d_h = 15$ mm shown as dashed blue circles.

II. METHODS

Figure 1 illustrates the circular Hele-Shaw cell with a bubble inlet at the center used in the experiments. The flow created is symmetric with a radial expansion that ensures T1 events will be present due to shear [30]. In addition, due to the decreasing flow velocity toward the edges of the cell, a single experiment inherently probes the system at a range of shear rates. The cell diameter is $d = 250$ mm and the gap height is $h = 0.75$ mm. The inlet has the diameter of 15 mm. Foaming is facilitated by reducing the water surface tension by mixing common household dishwashing liquid (Fairy) with the weight fraction of 1 to 39. The solution viscosity is indistinguishable from that of water, $\eta = 1$ mPa s, confirmed using Anton Paar 302 rheometer in a bob-and-cup geometry in the shear rate range of $\dot{\gamma} = 0.1$ to $\dot{\gamma} = 100$ 1/s.




FIG. 2. The foam studied is rather monodisperse, having the average effective radius $R_{eff} = 2.0 \pm 0.1$ mm, where the error is the standard deviation. The histogram contains bubble sizes from three snapshots separated by 200 s with data from 378 bubbles.

Both the fluid and air are driven into the cell using a constant pressure set by two SMC ITV0010 pressure controllers and a manual flow limiter. The inlet flow rates are set to give an average bubble velocity of $v = 1.0$ mm/s. The air is injected through a 26-gauge needle. The outer rim outlets are open to the ambient pressure. The imaging is implemented using a Canon M3 camera, which takes compressed video at 1920×1080 pixels using the h.264 encoding resulting in image resolution corresponding to 0.09 mm/pixel. The resulting exposure time of one frame is 40 ms, which is significantly shorter than the temporal separation between two sequential T1 events *at the same node*. Figure 1(c) displays an example of a grayscale snapshot overlaid to the schematic top view of the device. The video is then interpolated to 13 000 grayscale images at 10 Hz frequency, resulting on average 0.1-s time between subsequent frames and eventually cropped to a square corresponding to the region of interest. The cropped square can be then used as a grayscale image as is or it can be processed to a skeletonized binary image with exactly one-pixel-wide lines separating bubbles from each other. The bubbles are thus identified from images as connected 2D regions with an area A . Their size distribution using effective radius $R_{eff} = \sqrt{A/\pi}$ is illustrated in Fig. 2, highlighting the monodispersity of our foam.

In the skeletonized image, the nearest neighbors of a bubble are exactly 1 pixel apart. This information is further utilized in identifying the time and location of the T1 events: If there is a change in nearest neighbors, there is either one or more T1 or T2 events. Here, we did not observe any T2 events as no significant coarsening occurs at the timescales during which the bubbles occupy the imaging area. The raw images and the identified locations of the T1 events are then divided into training and test sets as detailed in Supplementary Tables 1 and 2 [31].

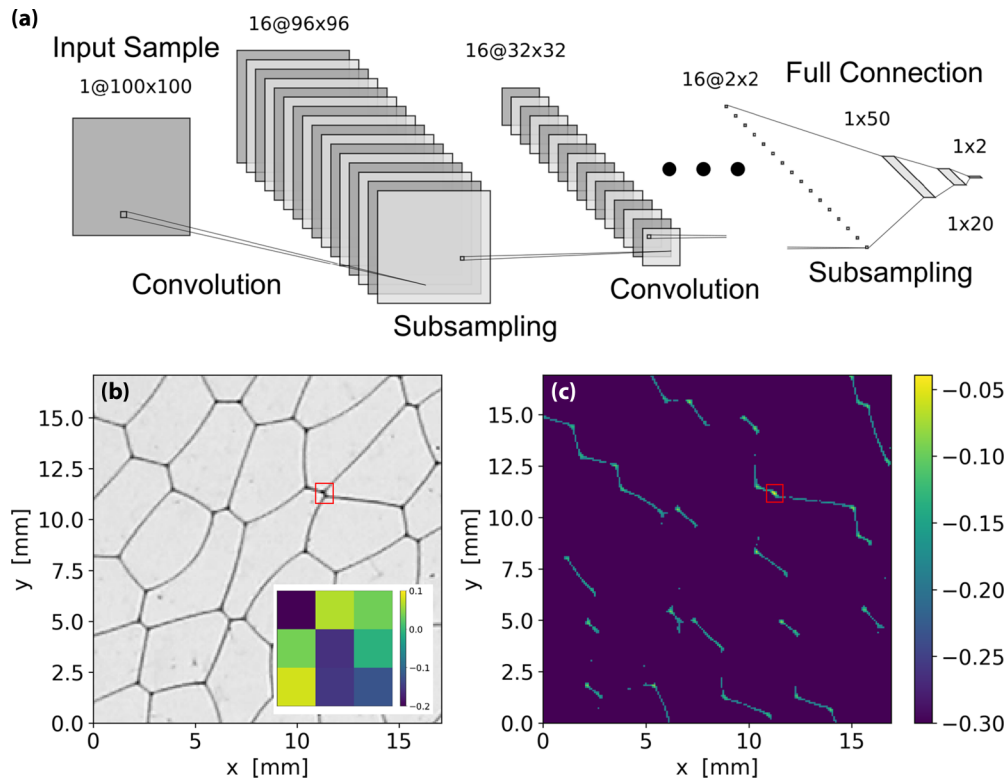


FIG. 3. (a) Architecture of neural network (illustrated using a tool from Ref. [32]) consists of three subsequent convolutional layers and subsampling layers and finishes with two fully connected layers. (b) Example of a raw grayscale image with one vertex that produces a T1 event highlighted with a red square. The inset shows one of the 3×3 convolution filters of a fully trained CNN. The dark diagonal of the filter captures the similarly aligned films of the raw images. (c) The convolution of the filter and the raw image illustrates the captured elements of the raw image. The convolution is thresholded for clarity.

We pose the following question: Given a single frame, are we able to predict the occurrence of a T1 event in the subsequent frame or frames? We start with a convolutional neural network (CNN) [33,34] and modify it to accept grayscale and skeletonized image as an input. Figure 3(a) depicts a schematic illustration of the CNN application leading to a single binary classification: Either T1 event occurs after Δt seconds or not. The neural network convolves and coarse grains the input image in three layers. The values are then run through fully connected neural network with two hidden layers until one value is left that is thresholded to produce the prediction. In the training, we use Adam optimizer [35]. In total, 10 000 training steps are used with learning rate of 5×10^{-5} , and the prediction results were found to saturate around 2000 training steps. The network was trained using sample batches at each step, rather than using the whole data set at every step. The batch size was 1000 samples and it was tested that increasing the batch size did not improve the prediction results anymore. An example of a grayscale input and a trained filter is depicted in Fig. 3(b) and the first convolution layer is depicted in Fig. 3(c). Here, the filter is from a fully trained network. Initially the filters contains randomly chosen values that converge to ones capturing essential features by reinforcing the features producing correct prediction using back propagation during training. The other filters capture different orientations of the films, while the function of some filters are not clear.

The input data is split into two, roughly equally large, sets for training and testing. In the training, a sequence of small regions of images are used. The essential structural features for T1 detection and prediction may be studied by limiting the amount of input data given to the network, here by limiting the size and location of regions of interest (ROI) around the possible T1 event. First, we generate a training set consisting of small ROIs that precede a T1 event in the next frame. Then, we balance the data set by adding randomly picked nonoverlapping regions which do not precede a T1 event until the data set contains 45–55% portions of each ROI type. This results in data sets that contain roughly equal amounts of both sample types so the AI does not benefit by favoring either of the outcomes. This prescreening is necessary since T1 events are rare. Without it, the algorithm scores high by always predicting “no T1.” The quality of the prediction is measured with a score parameter ξ that is the ratio of correct predictions (T1 or no T1) per all predictions. Some examples of input images given for CNN are shown in Fig. 4 with the outcomes of the prediction. In Fig. 4, the left column has samples that do not precede a T1 event within 0.1-s time frame while the samples on the right column precede a T1 event. The figure demonstrates that better accuracy of the predictions can be reached by using the vertex centered sample set than by using the bubble centered sample set. The CNN trained with bubble centered samples highlights features such as bubble shapes and stable vertices that are rather irrelevant for T1 prediction.




FIG. 4. Examples of two different kind of samples input for the CNN and the prediction results. On the left are negative samples that do not precede a T1 event within 0.1 s and on the right are positive samples which precede a T1 event. The sample types are starting from the first row: vertex centered grayscale image and bubble centered grayscale image. The labels “Correct” and “Incorrect” in the figure indicate the success of the CNN prediction.

III. PREDICTION RESULTS

A. Machine learning

Next, we explore the T1 predictions produced by the CNN. For this purpose, we use the score parameter defined earlier. We study the CNN for five different combinations of locations and grayscale or skeletonized frames for the full range of ROIs. The combinations selected allows us to compare (i) the effect of film or vertex thicknesses by comparing the predictions produced using grayscale (filled symbols) and skeletonized images (unfilled), (ii) the effect of focusing on the vertices (triangles) or bubble centers (circles), or (iii) concentrating on the shape of the bubble perimeter without its neighbors (diamonds). These choices of ROIs allow us to explore the image features, such as bubble shape, liquid fraction, or neighbor bubbles, by which the CNN is able to produce the best predictions. Figure 5(b) shows the score parameter ξ for different sizes ($4 \leq L \leq 13$ [mm]) and locations (bubble or vertex centered) of ROIs.

The best score $\xi = 98\%$ is obtained using vertex centered ROIs of 4×4 mm illustrated by the red triangle and a corresponding example input in Fig. 5(a). As a sanity check, we get the same score $\xi_+ = 98\%$ if we restrict to true positives by defining ξ_+ as correctly predicted T1 events per T1 predictions only. This indicates that there is no bias in the classification. Increasing the size of the ROI does not improve the score ξ and thus we conclude that the information in the local surroundings of the vertex provides a good indicator of a T1 event (filled triangles).

We considered that the change in the local liquid fraction is the reason for lower yield point and T1 events, which would appear as darker nodes. Thus, removing the information about the local liquid fraction by skeletonizing, we expected that the score parameter would decrease its value significantly. We performed the same CNN analysis on the skeletonized

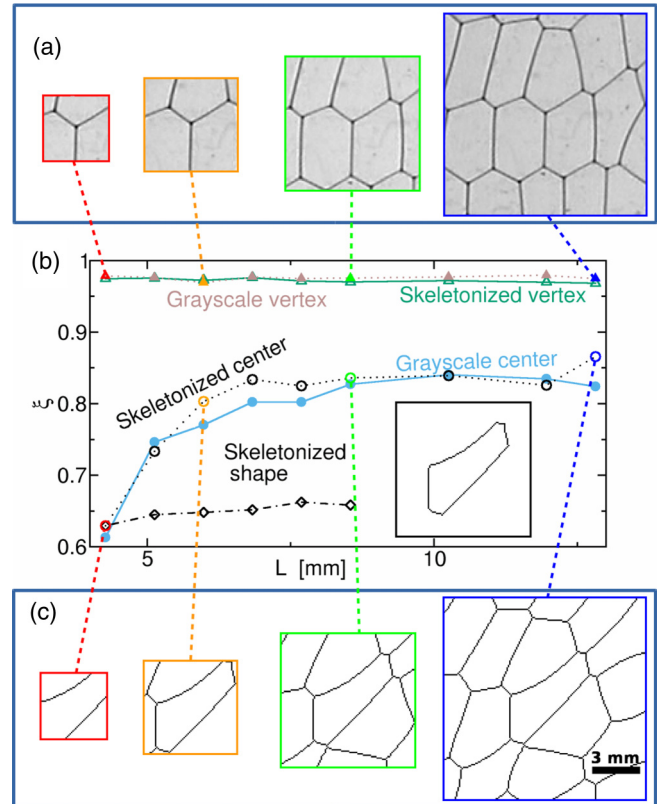


FIG. 5. Comparing the score parameter ξ to the size of regions of interest virtually reveals that for vertex centered images the score parameter is virtually independent of the ROI size. However, for bubble centered images the accuracy of the algorithm increases with the size of the region of interest ($L \times L$ mm) until the area is at least equivalent to that of four bubbles. There the performance is similar for both the grayscale (filled symbols) and skeletonized (unfilled symbols) images with the same centering. For the largest ROI, the algorithm works best for skeletonized network (unfilled circles), reaching almost 90% accuracy. The result is worse for grayscale images (blue) or if the data are reduced down to a single bubble shape [inset in pane (b), open diamonds].

frames, which result roughly the same scores with the grayscale images as witnessed in Fig. 5(b) (unfilled triangles). This is indicative that the CNN is not in fact capturing the grayscale levels in the images, but rather is predicting solely based on the orientations of the films surrounding the vertices. In other words, local liquid fraction or liquid motion in the films plays no significant role in the local yielding. Similar to the grayscale data, the skeletonized data show no significant dependence on the ROI size.

Moving the ROI center from the vertex to the bubble center of volume has a significant impact on the score parameter. First, let us concentrate on the skeletonized bubble shape [black diamonds and inset in Fig. 5(b)], where all the information about the neighboring bubbles is removed from ROIs. The score is close to $\xi = 65\%$ with only weak increase with ROI size L . This is significantly lower compared to the ROIs that focus on the vertex indicating that the bubble shape does not contain enough information for accurate prediction of T1 events. The score calculated only for the positive predictions

is even lower $\xi_+ = 60\%$ showing that the CNN does a slightly better job at predicting samples without a T1 event than predicting the events.

Now we include the films of neighboring bubbles but keep the ROI center at the center of bubble. These data are plotted as a function of the bubble centered ROI size for both grayscale (filled circles) and skeletonized (unfilled circles) images in Fig. 5(b). The predictions on the smallest ROIs for skeletonized or grayscale images gives similar scores as the predictions based on bubble shapes considered above. For both skeletonized and grayscale ROIs, the score ξ increases with the size of the ROI. As the ROI increases to include the entire bubble and parts of its neighbors (green circle), the score reaches over 80%, lower than the one that focused on the vertices, yet significantly higher than the one excluding all the bubble's neighbors. This supports our previous interpretation that the essential information on the T1 event is encapsulated by the structures and locations of the vertices and not by the bubble shape.

Comparison of the pure skeletonized data with and without the neighboring bubbles offers a visual confirmation to these observations. These structures are shown in the inset of Fig. 5(b) and highlighted by green color in Fig. 5(c). In the frame in question, a T1 event will take place in the lowest vertex, a location that is not obvious by only looking at the bubble shape. Thus, if the information on the neighboring bubbles is removed, the score parameter understandably drops dramatically.

B. Comparison to established methods

In the simplest view, the CNN could be predicting the T1 events simply by applying Plateau's laws. This seems to be the case in the sense that the CNN obtains the best predictions using the vertex centered samples and Plateau's rules apply to the vertices here. As shown in Fig. 6, the number of neighbors, on the other hand, only correlates weakly with the probability of the occurrence of a T1 event. As the figure shows, the most probable number of neighbors for a bubble going through a T1 event is actually six. This observation may explain why bubble centered samples yield worse predictability compared to vertex centered. Here, however, one must note that we have a finite accuracy in determining if Plateau's rules are obeyed. The Supplementary Video [31] highlights the lifetimes of these metastable states apparently violating Plateau's rule lasting up to several seconds. The actual violation of Plateau's rule lasts only a few milliseconds, observed in a three-dimensional (3D) case with high magnification [8]. Therefore, a more rigorous statement is that the CNN bases the predictions on apparent violation of Plateau's rules within the spatial and temporal resolutions of the measurement.

To further confirm the irrelevance of the bubble shape, we checked for any correlation with the textural tensor

$$\hat{\ell} = \left\langle \frac{\vec{\ell} \otimes \vec{\ell}}{\ell} \right\rangle \tag{1}$$

capturing the bubble shape [21] and the T1 event rate which has already been used for various bubble monolayers with similar results [10,20]. Here $\vec{\ell}$ represents the vector joining




FIG. 6. Histogram of number of nearest neighbors for all bubbles and bubbles right before a T1 event. Clearly the distribution is wider for the bubbles about to have a T1 event, indicating that local configurations where a bubble has exactly six neighbors are more stable than configurations where bubbles have more or less than six neighbors.

two neighboring Plateau borders within a single bubble. Using a rectangular mesh of 8×8 boxes dividing each image into 64 subregions, Fig. 7 shows no correlation between T1 events and the bubble shape for single frames.

Supplementary Table 3 shows the score parameter ξ along with other benchmark parameters of our CNN with various sanity checks [31]. These included a virgin data set without any reduction of vertices, comparison to a handmade algorithm, and prescreening the input data to contain only vertices involving four films, those appearing to violate Plateau's 120-deg rule. The last method has the best performance,




FIG. 7. Textural tensor and T1 event rate averaged over 8×8 sub-regions of imaged area and time. Although, both quantities show variation between different locations the quantities are not correlated.

reaching $\xi = 99\%$ score and capturing 34% of all the T1 events.

Based on our extensive testing, we conclude that there exist a significant subset of T1 events triggered by the changes of local microstructure and film orientations in the vertex that can be captured with a CNN without any information on the time dependence. Thus, already a small region of interest enables the CNN to predict the T1 events. In the future, to further improve the predictions we will train other neural networks using parameters which describe features of the local bubble geometry. We suspect adding some information about the magnitude and direction of the velocity or recent T1 history could improve the predictability.

IV. TIME-DEPENDENT PROPERTIES OF ELEMENTARY YIELD EVENTS

We next study to which extent in time the T1 events are predictable. To answer this, we study the temporal development of T1 events using the data plotted in Fig. 8.

We identify the score parameter ξ as a measure of predictability. Essentially, predictability here means the existence of a heterogeneous feature or “a defect”—since it leads to yielding—unlike a featureless material, allowing the CNN algorithm to make a prediction. We perform the analysis for vertex and bubble centered grayscale samples separately for the ROI size $9 \times 9 \text{ mm}^2$, corresponding to sample image highlighted green in Figs. 5(a) and 5(c). The size is chosen as the smallest area where the score parameter is saturated to its limiting value. Also, it is not necessary to include analysis with skeletonized images since the skeletonization only has small effect on the prediction score (<5%), if any.

Figure 8(a) depicts predictability ξ against the time difference $\Delta t = t_{\text{frame}} - t_{\text{event}}$ between the T1 event and the single input frame used by the CNN. This allows to examine the structural signatures of a T1 event before ($\Delta t < 0$) or even after ($\Delta t > 0$) the event. The brown triangles refer to data sets with vertex centered ROI, where the CNN is retrained and evaluated for each measurement point. Here, the vertex is always at the center of the ROI even if the lag Δt is increased.

We find that the prediction probability follows an exponential fit (solid line) and drops to one half at $\Delta t = -1.9 \text{ s}$ (before) the event. Using the ROIs from frames after the T1 event, the probability drops to one half faster at $\Delta t = 1.5 \text{ s}$, indicating that there is an asymmetry in predictability. Interestingly, a similar breakdown of time reversibility related to local geometry of T1 events has been reported previously in bubble raft shearing experiment measuring orientation of T1 events [22]. The asymmetry seems reasonable as shear rate drives the film shrinkage while force balance at the liquid-air interfaces drives the film growth for small shear rates (see Supplemental Video [31]) [36]. Although this picture is only a crude simplification as evident based on the video and the previous studies, T1 events also show clustering due to redistribution of stress [37], causing nonlocal deformation in the foam [38,39].

The bubble centered predictions are symmetric in Δt , although the prediction probability is not as good as in the vertex centered case. This might be due to the CNN focusing more on other parts than the film created during a T1 event.

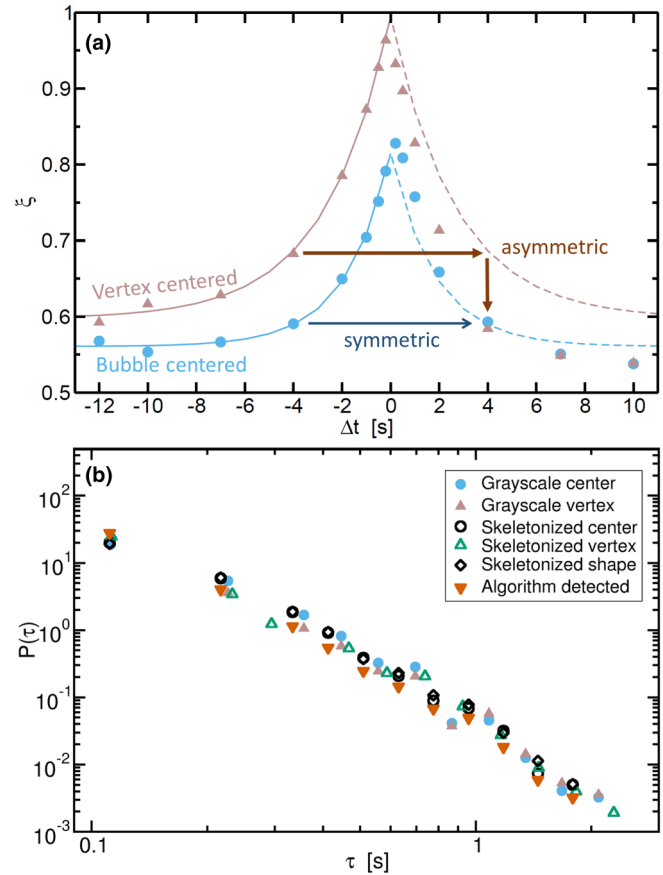


FIG. 8. (a) Temporal evolution of the score parameter ξ as a measure for structural inhomogeneity for different lag times between the CNN input frame and the T1 event $\Delta t = t_{\text{frame}} - t_{\text{event}}$, for the vertex (brown triangles) and bubble (blue circles) centered cases. The solid curves are the exponential fits to the data before the T1 event. The same fits are also plotted at the positive side (dashed). This shows the time-reversal asymmetry with the vertex centered case highlighted with arrows. (b) The characteristic timescale of subsequent T1 events is of the order of 1 s from the waiting time distributions $P(\tau)$.

Therefore, while the film instability triggers the event, the local neighborhood still plays a significant role. This can be understood in the context of slow energy dissipation, where a relaxation of the energy landscape enables reversible T1 events [15]. Thus, the configuration does not essentially change in the scale of few bubbles, even if the local film geometry can become completely different.

Figure 8(b) shows the T1 event waiting time distributions using the interevent arrival time in a restricted area of $50 \times 50 \text{ mm}^2$. The data based on all the CNN prediction algorithms agree with that obtained using the nearest neighbor detection algorithm capturing all the events (red triangles). Moreover, the data show that the typical T1 timescales are short, of the order of 1 s, in agreement with Fig. 8(a). To compare other dry foams with synthetic surfactants that do not affect the viscosity of the carrier liquid [40], the timescales for stress jumps were measured $t_{\text{jump}} = 1.5 \pm 0.5 \text{ s}$ [41] and the average time to reach 90% of the final film length was measured $t_{90} = 0.5 \text{ s}$ [42], that is in both cases of the same order of magnitude

as here, suggesting that film growth is driven by interfacial tension. With the average bubble velocity of $v = 1.0$ mm/s, the timescales are in the ballpark of the average bubble diameter of $d = 4$ mm, making T1 events highly localized.

V. CONCLUSIONS

We have successfully trained a general purpose CNN to recognize the neighbor swap T1 events in radial 2D foam flow using only snapshots of the structure with no time-dependent information. We capture the essential features of the images, namely vertices and film orientations. Using these features, we show that typically T1 events initiate from the unstable vertices that appear to violate Plateau's rules while the bubble shape is a less relevant quantity (Fig. 5). This highlights the importance of local film geometry and microstructure in rearrangements similarly to bubble coalescence [43]. The development of the shape or perimeter of a bubble and its neighborhood is symmetric in time for the bubble participating in a T1 event. In contrast, the changes in films are asymmetric in time. Here, the emphasis is on the neighbor swap aspect; the bubble is still a relevant unit for other processes such as coarsening [44] and recoverable elastic response [15]. The elastic energy stored in the system does not vanish instantly; a major part of it is stored in a different film configuration, making it possible to have reversible local rearrangements.

On average, a vertex about to yield looks very different before than after the event, as seen, for example, in Supplementary Video 1; that is, one can see whether the configuration under scrutiny is close to the local yield threshold [31]. This manifests as temporal asymmetry in the score ξ

that drops relatively fast to the baseline after the event. In other amorphous solids, e.g., granular pillars [26], the local configuration, namely the particle contacts, correlate with the yield stress and exhibit local variations (soft spots).

Our work has focused on the experimentally amenable 2D case considering the case of constant driving pressure and liquid fraction. It would now be quite interesting to investigate how one can change the foam yield in terms of the local predictability of T1 events and their spatiotemporal correlations. This could be achieved altering the physical properties of the sample foam, such as liquid fraction, polydispersity, geometry, or shear rate. However, changing any of these parameters may impose difficulties in maintaining a stable flow and keeping coarsening negligible. One particular direction would be particle-laden foams where the reinforcing particles would induce other dynamical timescales. This leads to a wide variety of industrial applications with tunable orientation-dependent properties such as tensile strength or heat conductivity [45,46].

ACKNOWLEDGMENTS

J.K. and A.P. acknowledge the 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 Aalto Science IT project for computational resources. L.V. acknowledges the funding from the Vilho, Yrjö and Kalle Väisälä Foundation via personal grant and Academy of Finland (Grant No. 278367). We acknowledge Henri Salmenjoki for sharing CNN code and helpful discussions.

-
- [1] D. L. Weaire and S. Hutzler, *Foam Collapse* (Oxford University Press, New York, 1999).
 - [2] K. A. Brakke, *Experimental Mathematics* **1**, 141 (1992).
 - [3] P. Moller, A. Fall, V. Chikkadi, D. Derks, and D. Bonn, *Philos. Trans. A Math. Phys. Eng. Sci.* **367**, 5139 (2009).
 - [4] G. Ovarlez, S. Cohen-Addad, K. Krishan, J. Goyon, and P. Coussot, *J. Non-Newtonian Fluid Mech.* **193**, 68 (2013).
 - [5] J. D. Barry, D. Weaire, and S. Hutzler, *Rheol. Acta* **49**, 687 (2010).
 - [6] M. L. Falk and J. S. Langer, *Phys. Rev. E* **57**, 7192 (1998).
 - [7] M. L. Manning and A. J. Liu, *Phys. Rev. Lett.* **107**, 108302 (2011).
 - [8] P. Petit, J. Seiwert, I. Cantat, and A.-L. Biance, *J. Fluid Mech.* **763**, 286 (2015).
 - [9] S. Cohen-Addad, R. Höhler, and O. Pitois, *Annu. Rev. Fluid Mech.* **45**, 241 (2013).
 - [10] B. Dollet and F. Graner, *J. Fluid Mech.* **585**, 181 (2007).
 - [11] M. Van Hecke, *J. Phys. Condens. Matter* **22**, 033101 (2010).
 - [12] T. Chevalier, J. Koivisto, N. Shmakova, M. J. Alava, A. Puisto, C. Raufaste, and S. Santucci, *J. Phys.: Conf. Ser.* **925**, 012025 (2017).
 - [13] L. Viitanen, J. Koivisto, A. Puisto, M. Alava, and S. Santucci, *Eur. Phys. J. B* **92**, 38 (2019).
 - [14] L. Viitanen, A. Halonen, E. Friström, J. Koivisto, M. Korhonen, A. Puisto, and M. Alava, *BioResources* **14**, 5716 (2019).
 - [15] M. Lundberg, K. Krishan, N. Xu, C. S. O'Hern, and M. Dennin, *Phys. Rev. E* **77**, 041505 (2008).
 - [16] N. C. Keim and P. E. Arratia, *Soft Matter* **9**, 6222 (2013).
 - [17] A. D. Gopal and D. J. Durian, *Phys. Rev. Lett.* **75**, 2610 (1995).
 - [18] D. J. Durian, *Phys. Rev. E* **55**, 1739 (1997).
 - [19] A. Abd el Kader and J. C. Earnshaw, *Phys. Rev. Lett.* **82**, 2610 (1999).
 - [20] B. Dollet, *J. Rheol.* **54**, 741 (2010).
 - [21] P. Marmottant, C. Raufaste, and F. Graner, *Eur. Phys. J. E* **25**, 371 (2008).
 - [22] Y. Wang, K. Krishan, and M. Dennin, *Philos. Mag. Lett.* **87**, 125 (2007).
 - [23] N. D. Denkov, S. Tcholakova, K. Golemanov, and A. Lips, *Phys. Rev. Lett.* **103**, 118302 (2009).
 - [24] S. S. Schoenholz, A. J. Liu, R. A. Riggelman, and J. Rottler, *Phys. Rev. X* **4**, 031014 (2014).
 - [25] J. M. Rieser, C. P. Goodrich, A. J. Liu, and D. J. Durian, *Phys. Rev. Lett.* **116**, 088001 (2016).
 - [26] E. D. Cubuk, S. S. Schoenholz, J. M. Rieser, B. D. Malone, J. Rottler, D. J. Durian, E. Kaxiras, and A. J. Liu, *Phys. Rev. Lett.* **114**, 108001 (2015).
 - [27] T. A. Sharp, S. L. Thomas, E. D. Cubuk, S. S. Schoenholz, D. J. Srolovitz, and A. J. Liu, *Proc. Natl. Acad. Sci. USA* **115**, 10943 (2018).
 - [28] A. Nicolas, E. E. Ferrero, K. Martens, and J.-L. Barrat, *Rev. Mod. Phys.* **90**, 045006 (2018).

- [29] D. Rodney, A. Tanguy, and D. Vandembroucq, *Modell. Simul. Mater. Sci. Eng.* **19**, 083001 (2011).
- [30] B. Dollet and C. Bocher, *Eur. Phys. J. E* **38**, 123 (2015).
- [31] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.2.023338> for a video of T1 events and further details about machine learning data sets and prediction results.
- [32] A. LeNail, Publication-ready *nn*-architecture schematics, <http://alexlenail.me/NN-SVG/LeNet.html>.
- [33] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, Berlin 2006).
- [34] H. Salmenjoki, M. J. Alava, and L. Laurson, *Nat. Commun.* **9**, 5307 (2018).
- [35] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [36] A.-L. Biance, S. Cohen-Addad, and R. Höhler, *Soft Matter* **5**, 4672 (2009).
- [37] K. W. Desmond and E. R. Weeks, *Phys. Rev. Lett.* **115**, 098302 (2015).
- [38] J. Goyon, A. Colin, G. Ovarlez, A. Ajdari, and L. Bocquet, *Nature* **454**, 84 (2008).
- [39] F. Elias, C. Flament, J. Glazier, F. Graneri, and Y. Jiang, *Philos. Mag. B* **79**, 729 (1999).
- [40] N. D. Denkov, S. Tcholakova, K. Golemanov, K. P. Ananthpadmanabhan, and A. Lips, *Soft Matter* **5**, 3389 (2009).
- [41] I. Cantat and O. Pitois, *Phys. Fluids* **18**, 083302 (2006).
- [42] M. Durand and H. A. Stone, *Phys. Rev. Lett.* **97**, 226101 (2006).
- [43] E. Forel, B. Dollet, D. Langevin, and E. Rio, *Phys. Rev. Lett.* **122**, 088002 (2019).
- [44] K. Khakalo, K. Baumgarten, B. P. Tighe, and A. Puisto, *Phys. Rev. E* **98**, 012607 (2018).
- [45] B. Wicklein, A. Kocjan, G. Salazar-Alvarez, F. Carosio, G. Camino, M. Antonietti, and L. Bergström, *Nat. Nanotechnol.* **10**, 277 (2014).
- [46] M. Yang, N. Zhao, Y. Cui, W. Gao, Q. Zhao, C. Gao, H. Bai, and T. Xie, *ACS Nano* **11**, 6817 (2017).