Topological defect coarsening in quenched smectic-C films analyzed using artificial neural networks

Ravin A. Chowdhury, Adam A. S. Green, Cheol S. Park, Joseph E. Maclennan,[*] and Noel A. Clark

Department of Physics and Soft Materials Research Center, University of Colorado, Boulder, Colorado 80309, USA

Mechanically quenching a thin film of smectic-C liquid crystal results in the formation of a dense array of thousands of topological defects in the director field. The subsequent rapid coarsening of the film texture by the mutual annihilation of defects of opposite sign has been captured using high-speed, polarized light video microscopy. The temporal evolution of the texture has been characterized using an object-detection convolutional neural network to determine the defect locations, and a binary classification network customized to evaluate the brush orientation dynamics around the defects in order to determine their topological signs. At early times following the quench, inherent limits on the spatial resolution result in undercounting of the defects and deviations from expected behavior. At intermediate to late times, the observed annihilation dynamics scale in agreement with theoretical predictions and simulations of the 2D XY model.

I. INTRODUCTION

Topological defects, which are stable disclinations or dislocations in ordered physical systems, are typically formed as a result of spontaneous symmetry breaking during phase transitions [1,2]. The formation and evolution of such defects, which have been predicted, and in some cases observed, in such diverse contexts as cosmology [3] and condensed matter [4], is a classical phenomenon that has been studied in many physical systems, including thin magnetic films [5] and superfluids [6]. Liquid crystals (LCs) are a particularly convenient medium in which to study the behavior of such defects experimentally, with disclinations easily visualized in both the nematic and tilted smectic phases [1]. The structure and dynamics of topological defects in quasi-two-dimensional liquid crystals is broadly reviewed in [7].

Fluid smectics are fundamentally interesting because they can be drawn into extremely thin, freely-suspended films of the order of a few molecular layers thick, allowing the study of physics in two dimensions (2D) [8–10]. In the smectic-A (SmA) liquid crystal phase, the long axes of the molecules are oriented, on average, along the layer normal, while in the smectic-C (SmC) phase they are tilted from the layer normal, breaking the axial symmetry of the SmA phase and introducing topological complexity. The topology of freely-suspended SmC films may be described by projecting the average molecular long axis (the director) onto the plane of the layers, defining a vectorial orientation field called the c-director. When these films are viewed in reflection under crossed polarizers, this orientation field typically creates a schlieren texture, with characteristic, crosslike extinction brushes centered on any topological defects [11].

The visual appearance of defects in SmC films depends, in general, on their topological strength, the illumination conditions, and the relative locations and orientations of the other defects [12]. Several experimental studies of SmC defect dynamics have considered films with only a small number of defects [13–16]. When there are only a few defects in the field of view, and they are well separated, they can either be identified manually or tracked automatically by cross-correlating the images with synthetically generated templates of model defect textures [14].

However, in dense arrays of topological defects, such as those generated in the quenching experiments described here, the orientation fields around the defects produce irregular and complex schlieren textures in polarized light, making detecting and tracking the defects using the previously implemented techniques impractical. Machine learning has been shown to be a useful tool enabling object detection in images obtained in such diverse areas as solid-state physics [17], cellular biology [18,19], and in protein folding experiments [20]. We shall demonstrate here that deep learning can be used to solve the seemingly intractable problem of detecting topological defects in dense, two-dimensional arrays in LC films.

The analysis of coarsening dynamics in LC systems with large numbers of densely-packed topological defects has been found historically to be challenging in both experimental and numerical studies because of the practical difficulty of detecting the defects. The coarsening dynamics of model 2D SmC films with high defect densities have been studied extensively using simulations [21–25]. Experimental studies of defect dynamics in thin, quenched SmC films were reported by Muzny [26], who described the basic phenomenology of quenching, proposed a mechanism for defect generation, and measured the approach dynamics of defect pairs and the decay of defect number with time following the quench.

In more recent experiments [27], high-speed video microscopy was used to capture the textures of mechanically quenched smectic-C films with much better temporal resolution than in Muzny's experiments. A preliminary analysis of the evolution of the observed arrays of topological defects using a convolutional neural network (CNN), a type of artificial neural network that is ideal for image analysis, demonstrated

———————
*joseph.maclennan@colorado.edu

the utility of machine learning but also revealed the limitations imposed by training the network using simulated images, discussed further below.

In the present study, we analyzed the same experimental images but using a newer CNN trained on experimental rather than simulated images to determine the locations of the defects. The network was able to predict the defect locations starting at earlier times and with a much higher degree of accuracy than before, with measurements verified by comparison with manually determined (human-annotated) coordinates. In addition, a binary classification network was trained to distinguish between defects of opposite topological sign, allowing a comprehensive analysis of the coarsening dynamics in these films. The observed decay of defect density with time was compared with the *XY* model, with the predictions of a model proposed by Yurke and coworkers [21], and with the results of numerical simulations [22]. At early times ($t < 0.4$ s), when the defect density is high, many of the defects have a separation that is smaller than the imaging and machine learning spatial resolution limits, and the number of defects counted is lower than theoretically predicted. At later times, the defect density exhibits power-law decay with an exponent of 0.9, in agreement with theory.

II. EXPERIMENT

In the quenching experiments, smectic-*C* films are drawn across a circular aperture in a glass cover slip set in the opening of an otherwise airtight chamber. Increasing the air pressure in the chamber causes the originally flat film to be distorted into a dome. When the pressure is suddenly released, the film collapses rapidly to being planar again, a mechanical quench that increases the hydrostatic pressure in the film, causing a short-lived transition to the smectic-*A* phase. The subsequent return to the smectic-*C* phase results in the spontaneous appearance of thousands of 2π disclinations, topological defects of unit strength (i.e., with winding numbers ± 1), in the film. This initially dense array of defects then coarsens by the mutual annihilation of defects of opposite sign [26]. The experimental setup is shown in Fig. 1 and described in further detail by Green [28]. The evolution of the defect texture was captured using a high-speed video camera (Phantom v12.2) with a spatial resolution of 1104×800 pixels and a bit depth of 16, operating at 500 fps.

The liquid crystal material used in these experiments was PM2, a 50:50 mixture by weight of SYNTHON ST00552 (2-(4-*n*-hexyloxyphenyl)-5-*n*-octylpyrimidine) and SYNTHON ST00557 (5-*n*-decyl-2-(4-*n*-octyloxyphenyl)pyrimidine) [29], with the phase sequence SmC 52°C, SmA 68°C, N 72°C, Iso [30]. Films 5 mm in diameter and 20–30 molecular layers (60–90 nm) thick were drawn in the SmC phase at room temperature and the quenching experiments conducted at 35°C.

The *c*-director field in a ~ 0.6 mm^2 region of the film was visualized using polarized reflection microscopy. Under crossed polarizers, the film displays a schlieren texture, with each defect core surrounded by four alternating dark and light brushes. In order to be able to visualize the coarsening dynamics at high frame rates, the average intensity of the image was increased by decrossing the polarizers. This halves the number of brushes around each defect, transforming the
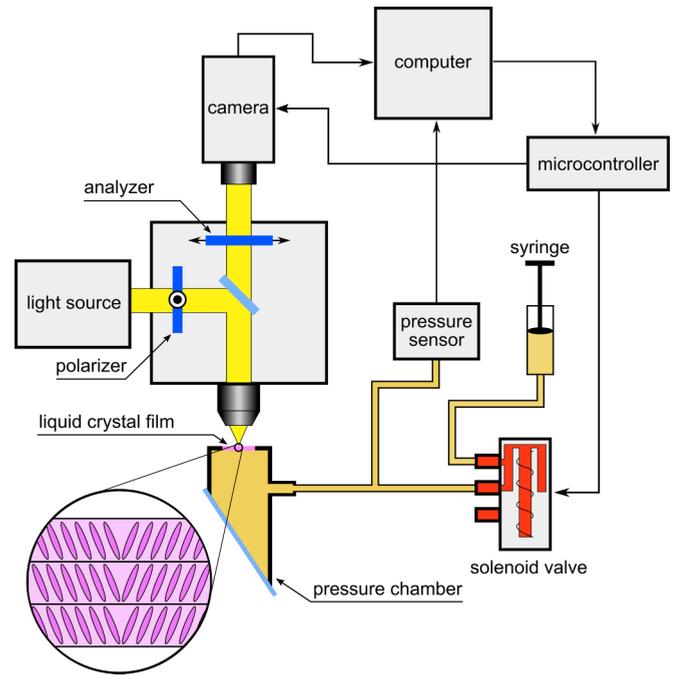
FIG. 1. Schematic of the film quenching experiment. A thin smectic-*C* film drawn across a small, circular opening in a sealed chamber is temporarily deformed to a dome by increasing the chamber pressure. A sudden release of the excess pressure then causes the film to return rapidly to its original planar geometry, a mechanical quench that results in the spontaneous formation of topological defects in the director field (inset). The defects are visualized using polarized reflected light microscopy and the coarsening of the defect texture recorded using a high-speed video camera.

crosslike brush textures to bow ties [31]. At the earliest times that defects can be observed following a typical quench, the average separation between defects is typically around 20 µm. Defects of opposite sign exert long-range $1/r$ attractive elastic forces on one another, while those of the same sign repel, interacting like infinite lines of electrical charge [26]. The defects also exhibit Brownian motion, diffusing laterally, in the plane of the film [13]. Over time, neighboring $+1$ and -1 defects approach each other and mutually annihilate, with most of the defects disappearing within a few seconds of the quench. A typical coarsening sequence is shown in Fig. 2.

III. TOPOLOGICAL DEFECT DETECTION

A rigorous analysis of the dynamics of topological defects in smectic-*C* films requires identifying their topological signs and tracking their locations as a function of time. Conventional feature-detection algorithms that use intensity thresholding or edge detection of object boundaries have been successfully used in many investigations of soft materials to detect features with regular shapes and well defined boundaries, such as colloidal particles in suspension [32] and islands and droplets on smectic films [33–36]. Detecting topological defects in liquid crystals is, however, a much more challenging task because the defects are identified principally by the diffuse schlieren textures surrounding them, which are typ-
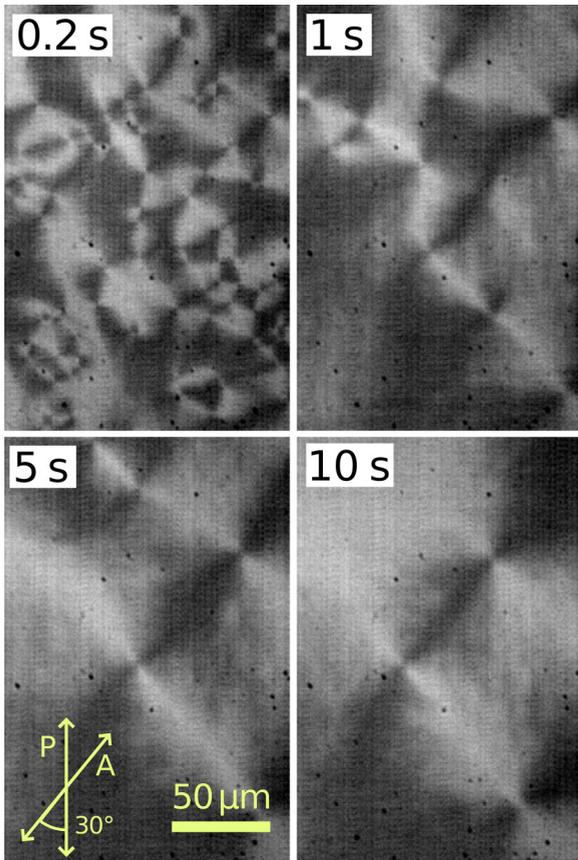
FIG. 2. Defect coarsening in a mechanically quenched SmC film. The topological defects are located at the points from which the pairs of light and dark brushes (resembling bow ties) emanate. Thousands of $+1$ and -1 defects are generated during a typical quench and these mutually annihilate over time. The polarizer and analyzer are decrossed by $60°$. The intensities of these images, which show only a small part of the film, have been normalized as described in the text.

ically irregular, have no well-defined boundaries, and vary in appearance. Quenching a film generates thousands of defects, resulting in a complex schlieren texture with defect core separations of as little as a few microns in our experiment. Accurate analysis of such textures is not feasible using our previous methods, in which we detected and tracked well-separated defects in equilibrium films by cross-correlating the experimental images with synthetically generated defect templates to determine their locations and brush orientations [14].

We recently demonstrated the utility of modern deep learning networks for defect detection, using the YOLOv2 network [37] trained on a large set of computed images of defect textures generated by Monte-Carlo simulations of the 2D *XY* model [27]. Although these training images were modified to emulate the experimental images more closely, nevertheless, several features present only in the experimental images regularly caused the trained model to predict false positives during analysis. For example, because dark speckles of the kind seen in many of the experimental images were not present in the training images, the network was not able to recognize these as being artifacts rather than defects. In addition, the black

mask of the microscope field stop was not considered in the computed training images, resulting in false positives along the edges of the field of view.

In the present analysis, we used YOLOv5, a deep learning network designed for fast object detection [38], that was trained on real experimental images, allowing us to achieve a much higher detection accuracy than before. YOLOv5 is (at this time) the newest iteration of YOLO, a lineage of neural networks that perform both bounding box predictions (object localization) and classification tasks for every object in the image, in a single instantiation of the network. Technical details of the neural network are given in Appendix A.

A. Image processing

Images captured during nine different quenching experiments were analyzed. To compensate for variations in the brightness and contrast of these data sets, all of the images were normalized to have the same average intensity and dynamic range. This ensured that all of the images input to the neural network had similar statistical properties, making it easier for the correct weights to be developed in training. This reduced the number of false positives and significantly reduced the number of training epochs required. A typical example of image normalization is shown in Fig. 3.

B. Model Training

The neural network was trained on experimental data, using a set of 141 images chosen at random from different film quenching experiments and divided into training and testing sets in the ratio 80:20. Gradient-descent calculations were carried out using the training set and the model performance was logged every cycle using the testing set. The locations of the defects in the training data (the 'ground truth' locations) were determined manually.

Training was carried out using Google's cloud research computing service, the Google Colaboratory, on an NVIDIA Tesla T4 GPU with 16 GB of memory [39]. Four YOLOv5 models of different sizes (listed in Appendix A) were trained for 500 epochs each on the same training set, using a batch size of 16. Model checkpoints (weights) were stored every epoch and the checkpoint yielding the highest mean average precision (mAP) on the testing set was selected. Choosing the optimal checkpoint of the neural network in this manner helped to prevent overfitting, which is generally a concern when the training data sets have fewer than 500 images, as in our case.

YOLOv5 uses several data augmentation techniques that further reduce the possibility of overfitting by changing the appearance of the training images in order to increase artificially the amount of training data. These modifications include carrying out vertical and horizontal flips, cropping, rotation, and a new method introduced with YOLOv5 called mosaic augmentation that meshes sections of multiple images together. In addition, the image quality may be altered intentionally using randomized exposure, saturation, or blurring [38].

Square bounding boxes were used to define the defect core locations. While keeping the bounding boxes small has
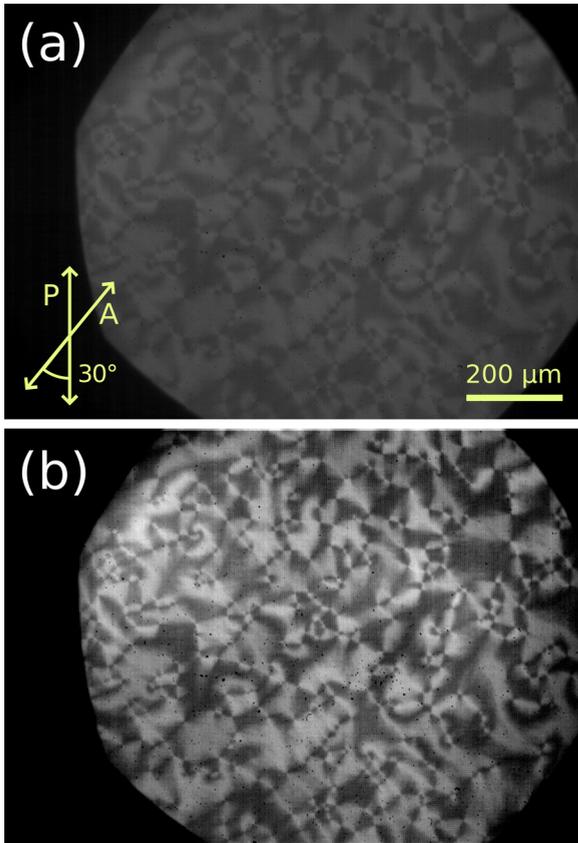
FIG. 3. Normalization of experimental images. (a) Raw experimental image ($t = 0.4$ s). (b) Normalized image. The raw images extracted from the quenching videos were normalized to have the same mean image intensity and dynamic range. This procedure ensured that the statistics of all of the experimental images from different quenching events would match those of the training set.

the benefit of increasing the precision of the detections and enabling defect detection even when the defects are close together, this also reduces the number of pixels associated with each defect, making training more difficult. The smallest bounding box size that our network could train on reliably was found to be 11×11 pixels.

C. Neural Network Performance

The performance of the four trained neural network models was evaluated using a control set of 48 normalized test images which were not included in the training data and hence had never been 'seen' before by the network. The detection results were compared to manually obtained 'ground-truth' locations using YOLOv5's built in performance metrics. Of particular significance are the mean Average Precision (mAP) values and the precision recall (F1) scores, since these metrics contain information about both false negatives and false positives [40]. The model with the highest recorded mAP score was the smallest model, YOLOv5s, which achieved a mAP of 0.970 and a peak F1 score of 0.96. The metrics computed for the various models are summarized in Table I in Appendix B.

Although using image cross-correlation to identify defects yielded accurate results in experiments where the density

of defects in the film was low [14], defect detection using this method becomes intractable in films with high defect densities, as is the case at early times in the quenching experiments described here. Cross-correlation is also sensitive to the presence of artifacts in the image, such as the liquid crystal deposits on the film chamber window visible as black speckles in most of our experimental images of quenched films. The resulting false predictions cause this method to have a low mAP, typically around 0.5 (see Table I in Appendix B). Deep learning networks, in contrast, can be trained to ignore such artifacts. The machine learning method developed by Minor *et al.* using YOLOv2 trained on synthetic images to analyze the quenching experiments occasionally misidentified the black speckles in the experimental images as topological defects [27], and the highest mAP achieved with this approach was around 0.8. In the present study, the YOLOv5 network trained on real images that included speckles and the edges of the field stop obtained significantly better mAP scores.

IV. DEFECT TRACKING AND SIGN CLASSIFICATION

A. Trajectory Linking

Once the defect locations had been determined in each frame, linked trajectories were constructed using TrackPy [41], a Python implementation of an algorithm originally developed for tracking colloidal particles [32]. The 'memory' feature of this program enables complete trajectories to be constructed even when there are occasional missed detections.

Visual comparisons of the computed defect trajectories following several different quenching events with manually determined defect locations demonstrated broad agreement, providing validation of the predictions of the neural network. A typical example of computed defect trajectories is shown in Fig. 4.

Two key assumptions about the defect dynamics were made in linking the defect locations to form trajectories. First, it was assumed that all of the topological defects present in the field of view were created at the time of the quench. Second, we assumed that defects of opposite topological sign annihilate and disappear in a pairwise manner. Defects that entered or left the frame during the coarsening experiment were exceptions that required special treatment when constructing their trajectories.

B. Defect Sign Classification from Brush Orientation Dynamics

When SmC films at equilibrium are observed in real time in the polarized light microscope, the signs of any topological defects may be readily determined by judiciously varying the decrossing angle of the polarizers and rotating the film. This procedure cannot, however, be followed in the quenching experiments, where most of the defects disappear less than a second after the quench occurs. The $+1$ and -1 defects are generally very similar in appearance and distinguishing them in static images alone is difficult. Near their cores, both kinds of defect resemble bow ties when the polarizers are decrossed but their orientations and the schlieren textures around them are highly variable, depending in a nontrivial way on their locations relative to other defects in the film. We have nevertheless been able to solve this fundamental classification
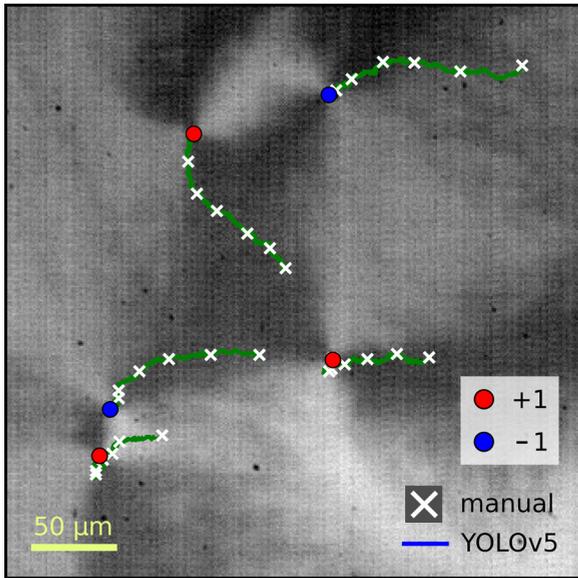
FIG. 4. Defect trajectories in a quenched film as determined by the trained neural network (green tracks) compared with defect locations identified manually (white crosses) every 500 frames (at 1 second intervals). The trajectories, which were determined over the course of 3000 frames (an elapsed time of 6 seconds), are superimposed here on the final image in the video sequence, with the surviving $+1$ and -1 defects shown respectively in red and blue. The dark speckles are liquid crystal droplets deposited on the chamber window by previously ruptured films. The neural network was trained to ignore these artifacts. The polarizer/analyzer settings are as in Fig. 3.

FIG. 5. Bright brush orientation as a function of time for a typical pair of annihilating defects. Quenching occurs at $t = 0$ and the pair annihilates at $t = 9.6$ s. Over the lifetime of the pair, the brushes around the $+1$ defect (red trace) fluctuate about an azimuth of $135°$ but do not change their average orientation significantly. The -1 defect (blue trace), in contrast, is more orientationally mobile. Initially oriented at $75°$ in this example, the brushes soon rotate to around $20°$, where they remain until annihilation. The insets show snapshots of the defects shortly after the quench ($t = 0.5$ s) and shortly before annihilation ($t = 8$ s). The variations in brush orientation over time are evaluated by a binary classification network in order to determine the topological signs of the defects.

problem by using their characteristic orientational dynamics to discriminate between the $+1$ and -1 defects.

First, the orientations of all of the defects, by which we mean the orientations of the bright brushes in the schlieren texture around the defects, were determined in every video frame. This was achieved by the previously mentioned technique of cross-correlating the region around every defect core with small, synthetic defect templates generated at different angles. When the computed defect orientations are plotted as a function of time, as in the example of Fig. 5, it is immediately apparent that the defects fall into two classes. The brush orientations of the first class of defect show little variation over time, with only small azimuthal fluctuations of less than $10°$. The mean orientation of these defects was found, in all of the quenched films, to be $\sim 135°$. The brush orientations of the second class of defect, in contrast, are not confined to a particular azimuth and vary substantially over time.

Since $+1$ defects have full axial (C_∞) symmetry, their brush orientation is relatively insensitive to thermal orientation fluctuations of the c-director. Defects of strength -1, on the other hand, are characterized by alternating bend and splay distortions of the surrounding c-director field and have only twofold (C_2) symmetry. A consequence of this anisotropy is that the appearance of the -1 defects is effectively more sensitive to thermal fluctuations and to their environment, responding readily to reorientations of the surrounding c-director field caused, for example, by distant defect annihilations or resulting from spatial rearrangements

associated with the approach of other defects in the film. The orientational mobility of the brushes around -1 defects was previously reported by Wachs [15], who observed that in thin films, an isolated -1 defect typically reorients shortly before annihilating with a $+1$ defect, maintaining the continuity of the director field along the line connecting the two defects. The appearance of the $+1$ defect, in contrast, seems to be relatively unaffected by the impending annihilation. This phenomenon has recently also been observed by Missaoui *et al.* [16] in thick SmC films and modelled theoretically by Tang and Selinger [42]. We conclude, therefore, that the first observed class of defects in our experiments has topological charge $+1$ and the second class topological charge -1.

Second, determination of defect sign on the basis of the brush orientation dynamics was achieved using a custom binary classification network comprising a fully connected model with one hidden layer. The hidden layer adds the complexity necessary to account for nonlinearities in the relationship between brush orientation dynamics and defect strength. The orientational data were reduced to three input features for each defect: the mean brush orientation, the standard deviation of the orientation, and the defect lifetime. The model determined the probability of each defect having strength $+1$ or strength -1. Defects with readily identifiable strengths were selected manually to generate training data for the model. 74 such defects were labeled in total, 60 for training and 14 for testing. The model was trained for 1000 epochs.
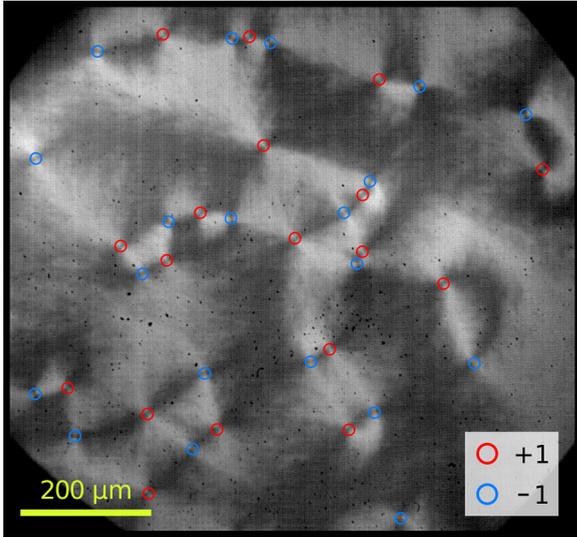
FIG. 6. Defect locations predicted by the neural network and classified according to topological sign. The image shows a SmC film six seconds after quenching, with the defects color-coded according to the topological sign predicted by the custom binary classification network. The defect locations determined by the deep learning model proved generally to be highly accurate. In this example, every defect was detected: 18 of strength $+1$ and 21 of strength -1. The film diameter is 5 mm, much larger than the image, the black borders here corresponding to the edges of the microscope field stop. The polarizer/analyzer settings are as in Fig. 3.

V. RESULTS

A. Detection and Classification of Defects

As we have seen in Fig. 4, the defect locations determined by YOLOv5 are in very good agreement with those determined manually. Starting at early times, soon after the quench ($t < 1$ s), the network even identifies defects in high-density regions that are missed by visual inspection. The numbers of $+1$ and -1 defects identified in each video frame were found to be roughly equal, as expected. Small deviations from parity are inevitable since only a finite region of the film is imaged, causing some of the partner defects to be located outside the field of view at any given time. Visual inspection confirmed that the number of false positives in any image was typically very small, even near the field stop defining the edges of the field of view or in the presence of artifacts in the image. An example of defects detected by the neural network and classified by topological sign is shown in Fig. 6.

B. Coarsening Dynamics

The decay of $N(t)$, the number of topological defects, as a function of time measured in nine different quenching experiments showed similar behavior. Defects could only be detected starting about 0.2 seconds after the quench was initiated, when the previously distorted film was flat and in focus again. In every experiment, N was observed to decrease slowly at first and then more quickly, decaying algebraically at longer times, with a roughly constant exponent. A measurement of $N(t)$ is shown in Fig. 7. Further examples are shown in the Supplemental Material [43].
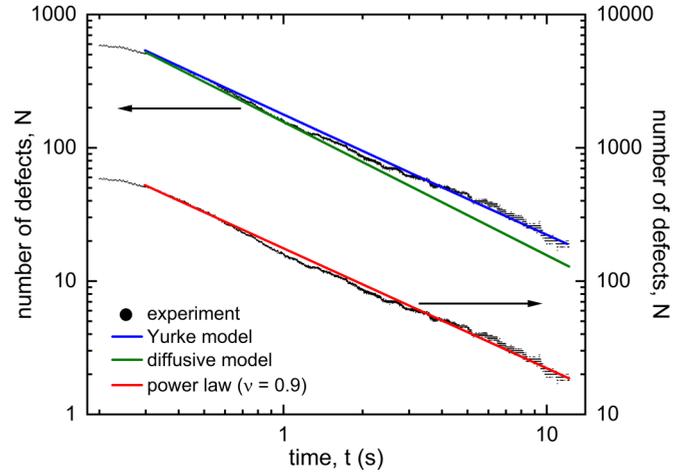
FIG. 7. Defect number vs time in a quenched smectic-C film. The large number of topological defects generated in a typical quenching experiment decreases over time by the mutual annihilation of $+1$, -1 defect pairs. Images of the film were analyzed starting when the film came back into focus, about 0.2 seconds after the mechanical quench. A fit with the diffusive model (green), which describes algebraic decay of the defect number N with an exponent of 1, predicts coarsening at long times that is faster than observed. The logarithmic correction term of Yurke's model (blue) results in a slower annihilation rate that closely approximates the experimental measurements. Power-law decay with an exponent of 0.9, as suggested by the simulations of Jang *et al.*, is in accord with the Yurke model and matches the data similarly well (red). The experimental data is duplicated here for visual clarity.

The overdamped, continuous XY model describing locally interacting, classical spins in 2D predicts an inverse power-law relation, $N \propto t^{-1}$ [1]. It is readily apparent from the graph, however, that the observed decay occurs more slowly overall than predicted by this model and deviates significantly from a simple power law at early times. Yurke *et al.* carried out computer simulations of the 2D XY model and described the observed coarsening behavior with $N \ln N \propto t^{-1}$, the logarithmic correction accounting for the effective drag on the defect cores [21]. Jang *et al.* also carried out numerical simulations including these frictional forces, finding that N varied as $N \propto t^{-0.9}$ [22], in agreement with the asymptotic scaling behavior predicted by the Yurke model. As is evident from Fig. 7, our experimental data are fit well by the Yurke and Jang predictions at intermediate and long times ($t \gtrsim 0.4$ s).

At early times, between 0.2 and 0.4 seconds after the quench, the observed dynamics deviate significantly from power-law behavior. Extrapolating the asymptotic slope to early times leads to the prediction that at 0.2 seconds there should be around 800 topological defects, about 200 more defects than were detected by the neural network. This discrepancy is attributed to the undercounting of defects that are closer than about 10 μm (11 pixels), which is the resolution limit for detection by the trained network. This limit is also manifest in defect pair correlation functions computed from the images, which are identically zero for defect separations closer than 10 μm, as shown in Fig. 10 in Appendix C. Comparing the pair correlation curves derived from experimental

data with those computed from simulations, also plotted in Fig. 10, supports the notion that the neural network is unable to discern topological defects that are very close together in the experimental images, resulting in undercounting at early times when the defect density is highest.

Finally, it has been suggested that there are circumstances in which the 2D XY model would be expected to exhibit exponential rather than power-law decay at early times (see, for example, Refs. [44,45]), which would clearly result in a knee of the kind shown in the plot of Fig. 7. However, the inherent uncertainty in our early-time data precludes any systematic consideration of this possibility.

VI. SUMMARY

In summary, we have demonstrated a deep-learning approach that allows us to detect topological defects in thin smectic-C liquid crystal films with a high level of accuracy. We have also developed a rigorous method for classifying the topological signs of the defects, using the distinctive orientational dynamic behavior of the director fields around them. A binary classification network trained to perform this task gives predictions consistent with the known physical properties of such arrays of defects, for example that mutual annihilation only occurs between defects of opposite sign and that in any given region of the film, the numbers of $+1$ and -1 defects are expected to be approximately equal.

We compared our experimental observations of the defect coarsening dynamics in films with several theoretical models. At long times, the number of defects was observed to decrease more slowly than predicted by the purely diffusive XY model, showing instead the power-law behavior with an exponent of 0.9 in agreement with the model of Yurke *et al.* and the simulations of Jang *et al.* The anomalously slow annihilation rate observed at early times is attributed to the undercounting of defects, an unavoidable consequence of the inability of the neural network to resolve defects with separations smaller than $10 \, \mu m$.

The deep-learning method described here could be improved by designing custom networks explicitly for detecting topological defects, which could allow a reduction in the number of required layers in the neural network, making defect detection faster and more transparent. The kind of machine learning implemented here could readily be applied to other detection tasks in soft materials, such as tracking a variety of inclusions in smectic films, colloidal particles in suspension, and bacterial cells in fluid media, or could be used to analyze the evolution of bubbles in foam coarsening experiments.

ACKNOWLEDGMENTS
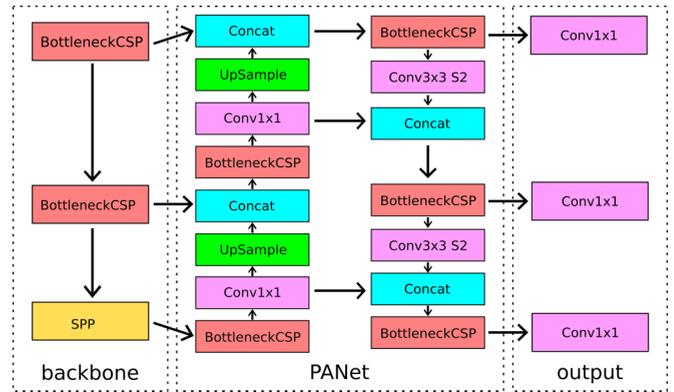
FIG. 8. Architecture of the YOLOv5 neural network. YOLOv5 comprises three parts, shown in the diagram. The model backbone is a Cross Stage Partial Network (CSPDarknet) for feature extraction. A Path Aggregation Network (PANet) combines these features etc.

APPENDIX A: THE NEURAL NETWORK

The architecture of YOLOv5, the neural network used in this study, is shown in Fig. 8. Unlike previous versions of YOLO that were developed in the Darknet framework, YOLOv5 is built using the PyTorch framework in Python. The backbone uses a Cross Stage Partial Network (CSPNet) to compress predicted features into fewer channels. A Spatial Pyramid Pooling Network (SPPNet) restructures the input to bypass fixed-size input constraints. The extracted features from the YOLOv5 backbone are then passed to a Path Aggregation Network (PANet) for feature fusion. Finally, the fused features are passed to the output layer (also called the YOLO layer), where the detection results are computed.

There are four popular YOLOv5 model sizes, with increasing numbers of layers: YOLOv5s, YOLOv5m, YOLOv5l, and YOLOv5x. We measured the training time, mAP value, and peak F1 score of all four models. The training times, shown in Table I, ranged from 50 minutes for the smallest model to 160 minutes for the largest.

APPENDIX B: PRECISION-RECALL CURVE AND F1 PLOT

In order to determine whether a detected result is a true or false positive, an intersection over union (IoU) of the detection bounding box with the ground truth bounding box is calculated. If the IoU ratio is above a prespecified confidence threshold, the detection is classified as a true positive, while if it is below the threshold, it is deemed a false positive. The precision-recall (PR) curve is created by sampling results from a range of confidence threshold values, usually from 0.5 to 0.95. This is the range used, for example, by the Microsoft COCO data set, a large annotated collection of images with only a few sizable objects per image [46]. In our case, in contrast, there could be hundreds of defects per image with bounding box sizes as small as 11×11 pixels. For bounding boxes as small as these, small pixelwise differences between manually created bounding boxes and those determined by YOLOv5 may impact the IoU dramatically. As a result, the likelihood of a true detection producing an IoU below 0.5 is relatively high. To obtain a better measure of
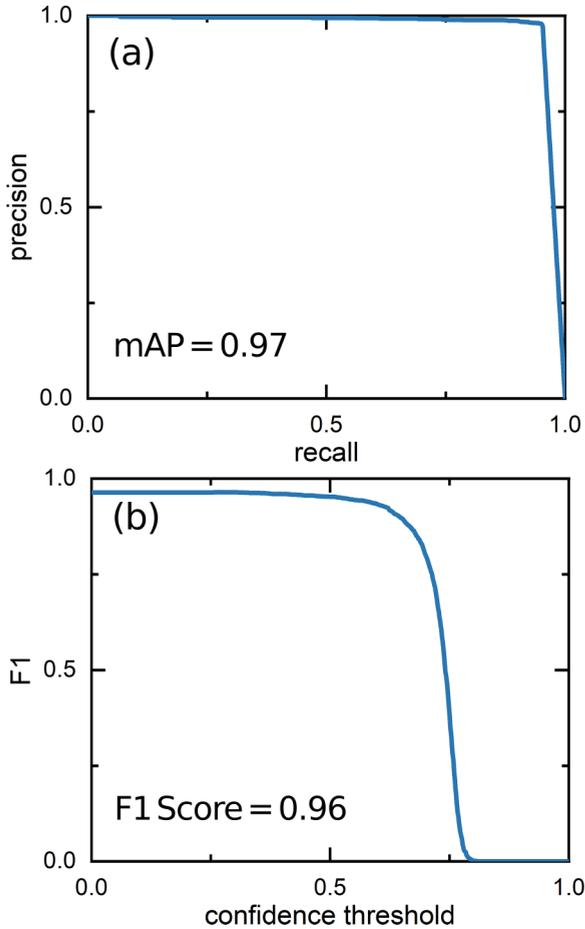
FIG. 9. Performance of the YOLOv5 model trained to detect topological defects. (a) The precision-recall curve yields a mean Average Precision (mAP) of 0.97, corresponding to a high degree of confidence in defect detection by the trained network. (b) The highest value on the F1 curve (the F1 score) was 0.96. The best F1 score was obtained at a confidence threshold of 0.27.

the network's results, we extended the threshold range from 0.0 to 0.95.

The mean average precision (mAP) is defined as the area under the precision-recall (PR) curve [Fig. 9(a)], a plot of the ratio of true positives to all detections (precision) vs the ratio of true positives to all ground truth values (recall). As seen in Table I, the smaller model sizes were found to yield slightly better mAP scores in this study, with the smallest model, YOLOv5s, achieving the highest mAP score of 0.970, marginally better than the score of 0.960 achieved by YOLOv5x, the largest model we looked at. These results indicate that the additional layers present in the larger YOLOv5 models do not result in better defect detection.

The F1 score is the harmonic mean of the precision and recall values, calculated for a range of confidence thresholds. Since every threshold value has an associated F1 score, the threshold that yields the highest F1 score is routinely reported. As is evident from Fig. 9(b), the highest F1 score (of 0.96) occurs when no detections are ignored, demonstrating that our model is unlikely to detect false positives or make double detections (two detections of the same defect). The F1 score

TABLE I. Neural network model training time and defect detection accuracy given by the mAP and F1 metrics. Shown are the metrics for topological defect detection using different YOLOv5 models, as well as those achieved when cross-correlating the experimental images with synthetic defect templates and using YOLOv2 trained using synthetic images.

Model	Training Time	mAP	F1 Score
Cross-Correlation	–	0.498	0.68
ForLL YOLOv2	1.1 hours	0.818	0.81
RealData11×11 YOLOv5s	0.8 hours	0.970	0.96
RealData11×11 YOLOv5m	0.9 hours	0.967	0.96
RealData11×11 YOLOv5l	1.4 hours	0.963	0.96
RealData11×11 YOLOv5x	2.6 hours	0.960	0.96

drops off rapidly above a confidence threshold of $\sim$0.70 and falls to zero before a threshold of 0.95 can be reached. This is a reflection of the sensitivity of the F1 score to small differences of a few pixels in bounding box locations and/or dimensions from the values manually determined in the control set. The highest F1 score for all four model sizes was 0.96, as shown in Table I. Also reported in the Table are the performance metrics for defect detection carried out both using template image cross-correlation and in the previous study using YOLOv2 [27].

APPENDIX C: DEFECT PAIR CORRELATIONS

The normalized pair correlation (radial distribution) function $g(r)$, a measure of the average defect density as a function of distance from any given defect, has been computed from the experimental data as a function of time. In these calculations, we used a radial bin size of 3 μm and the correlations were averaged over 10 frames ($\Delta t = 0.02$ s) in order to reduce the
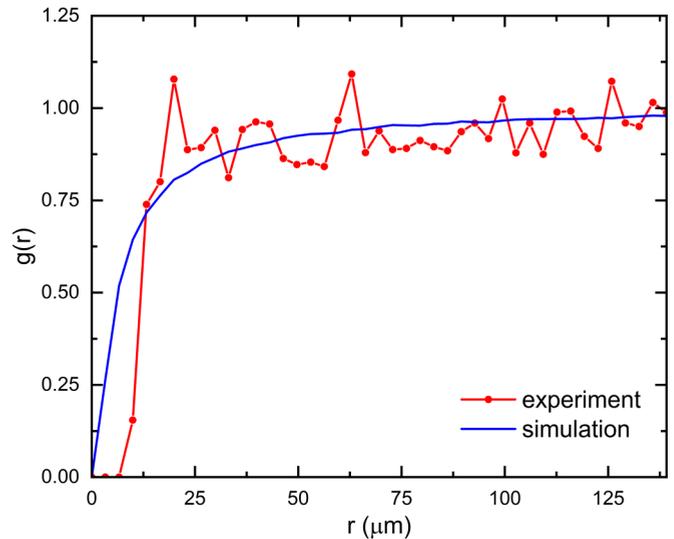
FIG. 10. Sign-agnostic topological defect pair correlation functions computed from experimental and simulation data shortly after a quench ($t = 0.2$ s). The correlation function derived from the simulation (blue) grows continuously from zero, while the experimental curve (red) is identically zero below 10 μm, the resolution limit for defect detection by the neural network.

statistical noise. We show in Fig. 10, by way of example, the sign-agnostic defect pair correlation function at $t = 0.2\,\text{s}$ for the quench event analyzed in this paper. This correlation function is identically zero below $10\,\mu\text{m}$ because no topological defects closer than than this distance were identified by the neural network. This is an artifact of the detection process: defects with separations smaller than the bounding box size of 11×11 pixels (or about $10\,\mu\text{m}$) could not be detected.

A numerical simulation similar to [22] of the evolution of a large number ($\sim$15,000) of diffusing, interacting topological charges initially generated with an average density similar to that observed in the film quenching experiments was carried out. The pair correlation function derived from the simulated data, also shown in Fig. 10, is seen to grow continuously from zero, as expected, saturating at long distances with a value $g(r) = 1$. Comparing this to the experimental correlation function confirms that defects in the experimental images with separations below $10\,\mu\text{m}$ are not detected by the neural network, leading to undercounting at early times following the quenching event.

[1] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics*, 7th ed. (Cambridge University Press, Cambridge, 2013).

[2] *Dynamics of Ordering Processes in Condensed Matter*, edited by S. Komura and H. Furukawa (Springer US, Boston, MA, 1988).

[3] I. Chuang, R. Durrer, N. Turok, and B. Yurke, Cosmology in the laboratory: Defect dynamics in liquid crystals, Science **251**, 1336 (1991).

[4] A. M. Srivastava, Topological defects in condensed matter physics, *Field Theories in Condensed Matter Physics*, Texts and Readings in Physical Sciences (Hindustan Book Agency, Gurgaon, 2001), pp. 189–237.

[5] N. S. Kiselev, A. N. Bogdanov, R. Schéfer, and U. K. Rößler, Chiral skyrmions in thin magnetic films: New objects for magnetic storage technologies? J. Phys. D **44**, 392001 (2011).

[6] R. Owczarek, Topological defects in superfluid helium, Int. J. Theor. Phys. **30**, 1605 (1991).

[7] K. Harth and R. Stannarius, Topological point defects of liquid crystals in quasi-two-dimensional geometries, Front. Phys. **8**, 112 (2020).

[8] C. Y. Young, R. Pindak, N. A. Clark, and R. B. Meyer, Light-Scattering Study of Two-Dimensional Molecular-Orientation Fluctuations in a Freely Suspended Ferroelectric Liquid-Crystal Film, Phys. Rev. Lett. **40**, 773 (1978).

[9] C. Rosenblatt, R. B. Meyer, R. Pindak, and N. A. Clark, Temperature behavior of ferroelectric liquid-crystal thin films: A classical *XY* system, Phys. Rev. A **21**, 140 (1980).

[10] P. Pieranski, L. Beliard, J.-P. Tournellec, X. Leoncini, C. Furtlehner, H. Dumoulin, E. Riou, B. Jouvin, J.-P. Fénerol, P. Palaric, J. Heuving, B. Cartier, and I. Kraus, Physics of smectic membranes, Physica A **194**, 364 (1993).

[11] R. Pindak, C. Y. Young, R. B. Meyer, and N. A. Clark, Macroscopic Orientation Patterns in Smectic-*C* films, Phys. Rev. Lett. **45**, 1193 (1980).

[12] D. R. Link, N. Chattham, J. E. Maclennan, and N. A. Clark, Effect of high spontaneous polarization on defect structures and orientational dynamics of tilted chiral smectic freely suspended films, Phys. Rev. E **71**, 021704 (2005).

[13] C. D. Muzny and N. A. Clark, Direct Observation of the Brownian Motion of a Liquid-Crystal Topological Defect, Phys. Rev. Lett. **68**, 804 (1992).

[14] N. M. Silvestre, P. Patrício, M. M. Telo da Gama, A. Pattanaporkratana, C. S. Park, J. E. Maclennan, and N. A. Clark, Modeling dipolar and quadrupolar defect structures generated by chiral islands in freely suspended liquid crystal films, Phys. Rev. E **80**, 041708 (2009).

[15] K. Wachs, Dynamics of smectic-*C* point disclinations in freely-suspended liquid crystal films, Undergraduate Honors Thesis, University of Colorado at Boulder, 2014.

[16] A. Missaoui, K. Harth, P. Salamon, and R. Stannarius, Annihilation of point defect pairs in freely suspended liquid-crystal films, Phys. Rev. Res. **2**, 013080 (2020).

[17] J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, Recent advances and applications of machine learning in solid-state materials science, npj Comput. Mater. **5**, 83 (2019).

[18] E. Moen, D. Bannon, T. Kudo, W. Graf, M. Covert, and D. Van Valen, Deep learning for cellular image analysis, Nature Methods **16**, 1233 (2019).

[19] J. G. Greener, S. M. Kandathil, L. Moffat, and D. T. Jones, A guide to machine learning for biologists, Nat. Rev. Mol. Cell Biol. **23**, 40 (2022).

[20] M. AlQuraishi, Machine learning in protein structure prediction, Curr. Opin. Chem. Biol. **65**, 1 (2021).

[21] B. Yurke, A. N. Pargellis, T. Kovacs, and D. A. Huse, Coarsening dynamics of the *XY* model, Phys. Rev. E **47**, 1525 (1993).

[22] W. G. Jang, V. V. Ginzburg, C. D. Muzny, and N. A. Clark, Annihilation rate and scaling in a two-dimensional system of charged particles, Phys. Rev. E **51**, 411 (1995).

[23] V. V. Ginzburg, P. D. Beale, and N. A. Clark, Scaling theory of particle annihilation in systems with a long-range interaction, Phys. Rev. E **52**, 2583 (1995).

[24] S. F. Burlatsky, V. V. Ginzburg, and N. A. Clark, Scaling model of annihilation-diffusion kinetics for charged particles with long-range interactions, Phys. Rev. E **54**, R1056 (1996).

[25] F. Rojas and A. D. Rutenberg, Dynamical scaling: The two-dimensional *XY* model following a quench, Phys. Rev. E **60**, 212 (1999).

[26] C. D. Muzny, Properties of defects in smectic-*C* thin films, Ph.D. thesis, University of Colorado at Boulder, University of Colorado at Boulder, 1994.

[27] E. N. Minor, S. D. Howard, A. A. S. Green, M. A. Glaser, C. S. Park, and N. A. Clark, End-to-end machine learning for experimental physics: Using simulated data to train a neural network for object detection in video microscopy, Soft Matter **16**, 1751 (2020).

[28] A. A. S. Green, Liquid and crystal: The applicability of the *XY* model to experimental systems of two-dimensional topological fluids; and revealing the nanoscale structure of the bent-core alpha phase, Ph.D. thesis, University of Colorado at Boulder, 2019.

[29] SYNTHON Chemicals GmbH & Co. KG.

[30] K. Harth, Episodes of the life and death of thin fluid membranes – patterns and dynamics at the cross-over from two to

three dimensions, Ph.D. thesis, Otto von Guericke University, Magdeburg, 2016.

[31] D. R. Link, Symmetry and structure of freely suspended liquid crystal films, Ph.D. thesis, University of Colorado at Boulder, University of Colorado at Boulder, 1998.

[32] J. C. Crocker and D. G. Grier, Methods of digital video microscopy for colloidal studies, J. Colloid Interface Sci. **179**, 298 (1996).

[33] Z. Qi, Z. H. Nguyen, C. S. Park, M. A. Glaser, J. E. Maclennan, N. A. Clark, T. Kuriabova, and T. R. Powers, Mutual Diffusion of Inclusions in Freely Suspended Smectic Liquid Crystal Films, Phys. Rev. Lett. **113**, 128304 (2014).

[34] Z. Qi, C. S. Park, M. A. Glaser, J. E. Maclennan, and N. A. Clark, Experimental realization of an incompressible Newtonian fluid in two dimensions, Phys. Rev. E **93**, 012706 (2016).

[35] Z. Qi, K. Ferguson, Y. Sechrest, T. Munsat, C. S. Park, M. A. Glaser, J. E. Maclennan, N. A. Clark, T. Kuriabova, and T. R. Powers, Active microrheology of smectic membranes, Phys. Rev. E **95**, 022702 (2017).

[36] E. Hedlund, K. Hedlund, A. Green, R. Chowdhury, C. S. Park, J. E. Maclennan, and N. A. Clark, Detection of islands and droplets on smectic films using machine learning, Phys. Fluids **34**, 103608 (2022).

[37] J. Redmon and A. Farhadi, YOLO9000: Better, Faster, Stronger, arXiv:1612.08242.

[38] G. Jocher, A. Stoken, J. Borovec, NanoCode012, A. Chaurasia, TaoXie, L. Changyu, A. V, Laughing, Tkianai, YxNONG,

A. Hogan, Lorenzomammana, AlexWang1900, J. Hajek, L. Diaconu, Marc, Y. Kwon, Oleg, Wanghaoyang0106 *et al.*, ultralytics/yolov5: v5.0 - YOLOv5-P6 1280 models, AWS, Supervise.ly and YouTube integrations (2021).

[39] Google Colaboratory, https://colab.research.google.com/.

[40] G. Salton and M. J. McGill, *Introduction to Modern Information Retrieval*, McGraw-Hill computer science series (McGraw-Hill, New York, 1983).

[41] D. B. Allan, T. Caswell, N. C. Keim, C. M. van der Wel, and V. W. Ruben, Trackpy: Fast, Flexible Particle-Tracking Toolkit, http://soft-matter.github.io/trackpy/.

[42] X. Tang and J. V. Selinger, Annihilation trajectory of defects in smectic-*C* films, Phys. Rev. E **102**, 012702 (2020).

[43] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.107.044701 for additional experimental measurements of the decay of the number of topological defects as a function of time.

[44] D. Toussaint and F. Wilczek, Particle–antiparticle annihilation in diffusive motion, J. Chem. Phys. **78**, 2642 (1983).

[45] R. Loft and T. A. DeGrand, Numerical simulation of dynamics in the *XY* model, Phys. Rev. B **35**, 8528 (1987).

[46] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, Microsoft COCO: Common Objects in Context, in *Computer Vision - ECCV 2014*, Vol. 8693, edited by D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Springer International Publishing, Cham, 2014), pp. 740–755.