

Methods and Supplementary Information

Active Foam: The Adaptive Mechanics of 2D Air-Liquid Foam under Cyclic Inflation

Laurel A. Kroo, Matthew Storm Bull and Manu Prakash

SI Movie Files 1-5:

SI Movie 1: Depicts inflation and deflation experiments. Raw data example.

SI Movie 2: Example of segmentation and feature extraction. Features overlay on raw experimental data demonstrates automated experimental analysis.

SI Movie 3: Qualitative strain field visualization. Video generated by compiling a video constructed of only the first frame of each cycle (digital stroboscopic sampling), and then uses a moving time-window convolution technique to highlight regions of high strain rate (See: flowtrace).

SI Movie 4: Demonstrates T1 transitions in the foam; examples are shown of both reversible and irreversible transitions.

SI Movie 5: Example of 4-bubble simulation to study the cooperation between vertex swirl.

<https://drive.google.com/drive/folders/1CoslaMBrYyhZqpAHL9Flx4pTufXtVnC?usp=sharing>

Contents

1 Comparison between Foam and an Ideal Potential Flow Source/Sink	16
2 Sample Preparation	18
3 Properties of Initial Foam Structures	18
A Quenched disorder and Effective Temperature	18
B Polydispersity and Coordination Number	20
C Isolated large bubbles	21
D Edge Dispersion and distance to yield	22
4 Simulation for Micro-structural Simulation of Vertex Motion	22
5 Multi-bubble Simulation Methods	23
A Galilean Transform on Edge Motion Equations	23
B Bubble pressure modelled as an ideal gas	24
6 Spring Network Models	24
7 Methods for Segmentation of Experimental Data	25
8 Measuring CW/CCW "Swirl" Directionality in Vertex Trajectories	25
A Lock-In Amplifier for Binary Directional Measurement	25
9 Cellular Segmentation (based on Watershed algorithm)	26
A Network Connectivity	26
B Vertex Identification	26
C Edge Identification	26
10 Methods for T1 Transition Identification	27
11 Data and Code Availability	27

1. Comparison between Foam and an Ideal Potential Flow Source/Sink

For the purpose of identifying interesting signatures in the foam, we can compare any point in the material to a simple potential flow with an unsteady boundary condition. The result is shown in figure 7; where the black lines represent the predicted trajectory from this analytical model. Deviations from this prediction indicate non-newtonian mechanisms at play.




Fig. 7. On the left is the predicted tracer pathlines from an ideal potential flow. On the right is an overlaid version of the prediction (in black) with the actual experimental data in color.

This analytical model is standard and is also cited in well-known texts in fluid mechanics (56). Consider a velocity potential in Cartesian coordinates (where u and v are components of \mathbf{v}):

$$u(x, y, t) = \frac{d\phi}{dx} \quad [8]$$

$$v(x, y, t) = \frac{d\phi}{dy} \quad [9]$$

We assume incompressibility:

$$\nabla \cdot \mathbf{v} = 0 \quad [10]$$

And for each timestep, the flow is irrotational such that for a typical closed contour not containing a singularity:

$$\nabla \times \mathbf{v} = 0 \quad [11]$$

Another way of saying this is that there is not a net vorticity, (unless it is constant and generated by a singularity in the flow, which constitutes an exception).

Equations 10 and 11 lead to Poisson's equation for the velocity potential:

$$\nabla^2 \phi = 0 \quad [12]$$

Velocity potentials (ϕ) are solutions that satisfy the Poisson equation 12 and obey continuity 10. These solutions are linearly superpositionable.

Equation 5 becomes in 2D polar coordinates:

$$\frac{1}{r} \frac{d}{dr} \left(r \frac{d\phi}{dr} \right) + \frac{1}{r^2} \frac{d^2 \phi}{d\theta^2} = 0 \quad [13]$$

where r is the radius from the origin.

Now, say we would like to solve for the velocity field (u, v) that occurs because of an expanding incompressible volume (source), such as our active bubble. We will solve this in the unsteady domain such that the radius of the expanding circle is equal to a function $a(t)$ (defined later).

We see that ϕ is a function of r and t only.

The boundary condition will be a velocity condition: the fluid at the expanding bubble edge needs to match the driving velocity:

$$\frac{\partial \phi}{\partial r} = \frac{da}{dt} \text{ at } r = a(t) \quad [14]$$

Equation 13 then simplifies to:

$$\frac{\partial}{\partial r} \left(r \frac{\partial \phi}{\partial r} \right) = 0 \quad [15]$$

$$r \frac{\partial^2 \phi}{\partial r^2} = - \frac{\partial \phi}{\partial r} \quad [16]$$

Integrating both sides, we see:

$$\int r \frac{\partial^2 \phi}{\partial r^2} dr = \int - \frac{\partial \phi}{\partial r} dr \quad [17]$$

$$r \frac{\partial \phi}{\partial r} - \phi = -\phi + c_1 \quad [18]$$

Apply boundary condition given in equation 14 to solve for $c_1(t)$:

$$c_1(t) = a \frac{da}{dt} \quad [19]$$

$$\frac{\partial \phi}{\partial r} = \frac{a}{r} \frac{da}{dt} \quad [20]$$

Integrate a second time:

$$\phi = a(t) \frac{da}{dt} \ln(r) + c_2 \quad [21]$$

Set $c_2 = 0$ to set velocity potential relative to origin.

To solve for the radial component of the velocity, take the derivative of the velocity potential with respect to r :

$$u_r = \frac{a}{r} \frac{da}{dt} \quad [22]$$

Now say the source is oscillating such that:

$$a(t) = \frac{r_{max} - r_{min}}{2} \sin 2\pi t + \frac{r_{max} + r_{min}}{2} \quad [23]$$

Where r_{min} is the minimum active bubble radius, and r_{max} is the maximum active bubble radius. Alternatively, (for comparison to our actual experimental data), we can define $a(t)$ as a piece-wise ramp and hold function similar to the one shown in figure 2c.

For a known maximum radius and minimum radius of the active bubble, we can approximate the projected position of any initial coordinate in the 2D potential flow for any point in time.

2. Sample Preparation

Preparation of the foams discussed in this paper were done by hand, by cleaning and then wetting a glass plate with a surfactant-laden solution. Bubbles were generated by moving a steel blunt-tip needle, partially submerged, in the wetting layer, dispersing nitrogen gas at a gauge pressure of about 0.75 psi. A second wetted glass plate was then set 1 millimeter above the bubble layer, with spacers controlling the gap distance between plates. The resultant foams have the properties that they are polydisperse (see example distributions in Appendix B) and single-layer (2D), with volume fractions under ten percent. Because of our preparation technique, these are frustrated structures far from any global energy minima (unlike hyperuniform foams which have experienced significant annealing through long-term relaxation and coarsening (57)). Our initial foams are internally diffusive in character, and have potential structural energy that relaxes over very long timescales (order 10-30 hours), in the absence of other perturbations. The preparation of the initial foam introduces quenched disorder in the structure, that in the absence of perturbation, only relaxes on timescales far longer than that of the experiment.

The experiment was illuminated with a low-power LED backlight sheet with an acrylic diffuser in a large optically-isolated box. A wide-field lens camera imaged the foam from below. The illumination source was very low power, and over 25 cm from the foam (very long optical axis). The temperature in the box was held approximately constant (thermostat-controlled the room temperature at approximately 25 deg C). The humidity was intentionally kept high within the box (using dishes of evaporating water) to minimize any evaporative effects at the edges of the Hele-Shaw cell on the timescale of experiments. These precautions were in an effort to reduce any variability in the rheology of the surfactant or the surface tension of the fluid on the timescale of these experiments. A typical experiment duration was around 37 minutes for typical single-foam experiment for 70 inflation/deflation cycles. Additionally, all newly prepared foams were set within the chamber to equilibrium for 5-10 minutes prior to commencing to experiments. This ensured that any changes in environmental conditions stabilized fully prior to experiments beginning. This soak time also allowed the material to relax to a somewhat metastable jammed state (no active TIs visible in the structure over 5+ minutes), despite being a highly frustrated and reasonably high-energy structure.

Additionally, an intensive cleaning protocol was used on the plates prior to each experiment, which included washing the plates with water, followed by extensive cleaning with solvents (isopropanol followed by acetone) to breakup / remove any surfactant or lingering solution left on the plate. We experimented with using a plasma cleaner on the plates as part of our cleaning protocol, but found that it was somewhat unnecessary given the extensiveness of the protocol. The cleaning steps to ensure smooth, clean surfaces took vastly longer than the experiments themselves.

3. Properties of Initial Foam Structures

A. Quenched disorder and Effective Temperature. Because of the way in which these foams are made (see: section 2), the structure can be quite fragile and soft. The placement of the top plate over the bubble raft forces the geometry into a new, more frustrated configuration over a short period of time (much like "quenching"). This locks in a fragile, disordered state that is mechanically quite irregular.

Although there is very little true noise / or thermally-induced fluctuations in our system, the prepared foams in this study are interestingly not in a particularly ultra-low or zero temperature state. Although they are jammed (the foams prior to an experiment are not undergoing spontaneous TIs or exhibiting any directed motion over timescales on the order of many minutes), the centroids of the cells move as if they had an effectively diffusive nature. The evidence for this is that the mean squared displacement of the centroids of the cells prior to a perturbation has a diffusive character, characterized by the slope of approximately one on the log-log MSD plot in figure 8.

These foams are highly frustrated and the vertex motion is "fluid-like", despite their volume fractions under 10 percent, and the confluent, globally jammed nature of the bubbles. The initial structure is storing energy, and the internal motion of this structure dissipates this energy very slowly (via shear on lubrication films at the top and bottom plate), as it moves down a highly frustrated landscape.

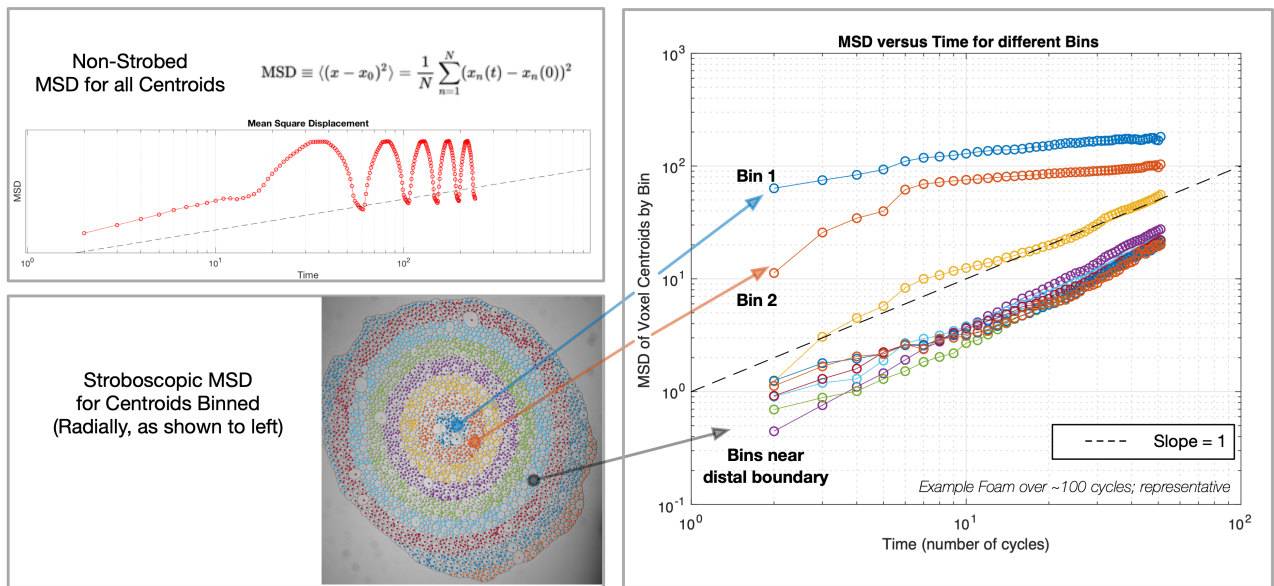


Fig. 8. Top Left: Mean Squared displacement as a function of time is shown. Bottom Left: We can define bins, moving radially outward from the site of activity. Right: Mean Squared Displacement is shown for each bin, where the x-axis is sampled stroboscopically (only includes the first frame of each cycle). This highlights long-term adaptation of the structure through many cycles.

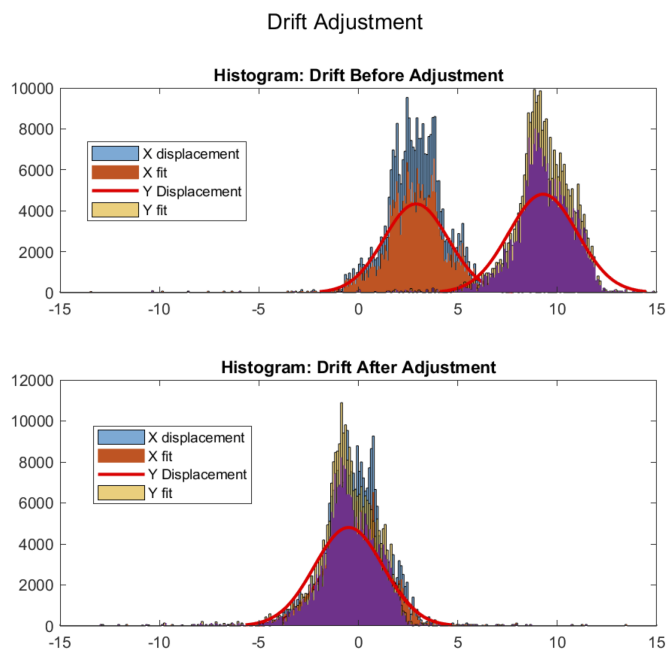


Fig. 9. Prior to computing MSD, it is necessary to correct the displacements first for any background drift (caused by gravity / occasional error in the flatness of the experiment) as shown in this plot

B. Polydispersity and Coordination Number. Histograms are shown here for the polydispersity (area distributions) for the 3 videos (15 cycles) used for the experimental analysis in Figures 1 through 6. As expected, the median and average coordination is 6 (a robust topological feature of foam). The blue line corresponds to Foam 1 in figures in the paper, The orange line corresponds to Foam 2; yellow line to Foam 3.

Because of the handmade nature of these foams, not all foams had the same average bubble size, and the distributions (while similar in shape), were not identical.

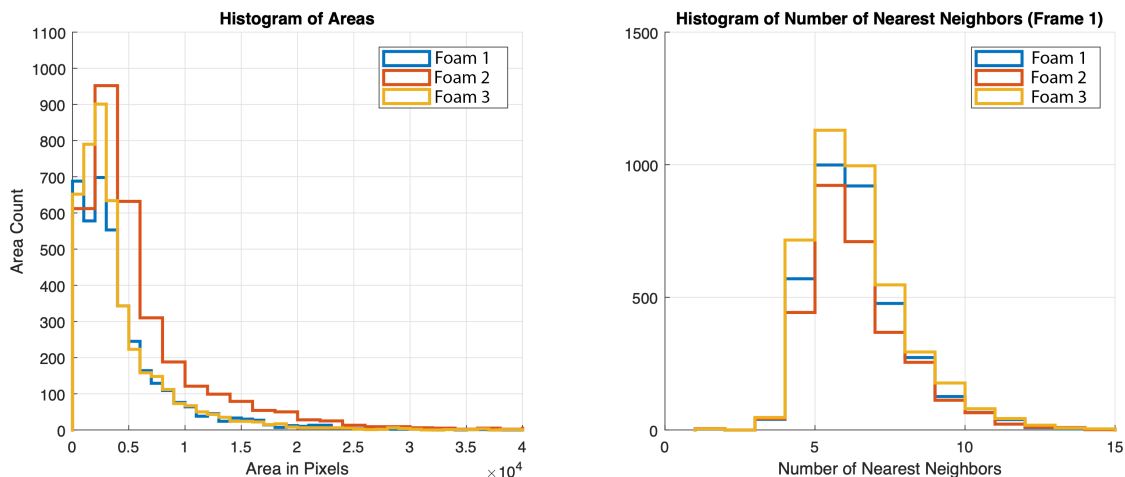


Fig. 10. Polydispersity histogram of the 3 example datasets discussed in paper

C. Isolated large bubbles. From our experiments, it is somewhat unclear what the effect of very large cells in the far-field have on the local material response (e.g. the effect does not seem trivial/straightforward from the data that large cells either encourage or explicitly discourage T1s in their immediate vicinity). We hypothesize the reason for this may be because of two (or more) competing effects:

We are guessing that a very large cell (say embedded in an otherwise relatively monodisperse foam) would cause a local region of heterogeneous mechanical properties — specifically, a lower potential energy regionally and potentially a faster propagation of mechanical signals due to the reduced dissipation per unit area (from the material lacking additional structure/fluid/mass in that region). One might expect this to introduce a meso-scale discontinuity into the material that might make the region particularly susceptible to T1s.

However, a large bubble embedded within a disordered foam will also locally increase the polydispersity in the foam. The local polydispersity tends to be directly influence the the local structural disorder (not always, but there is a general correlation). As demonstrated by both our simulations and experiments, we see that abnormally large local disorder encourages cooperation of swirl and larger swirl ‘grain’ sizes (somewhat counter-intuitively). While further analysis would be required for a definitive solution in this space, this finding implies that exceptionally large bubbles might have less of an effect than one might expect at “breaking-up” regions of coherent swirl domains because of effective material heterogeneity - due to this more subtle role of disorder apparently stabilizing the coherency of swirl grains.

The data we have on this topic is inconclusive; we don’t see T1s localizing to abnormally large cells, nor do we see T1s explicitly avoiding them. We suspect that it may be because of the two effects we describe above somewhat cancelling each other out. We look forward to examining this more closely in future experiments.

However, we have seen (via experiments) that in the limit of high area polydispersity, the centroids of abnormally large cells in the immediate vicinity of activity (less than 10 average cell diameters) tend to migrate, through sequences of T1 transitions, toward the source of activity. Because the activity is trying to minimize the regional structural energy, attracting proximal large bubbles is an effective solution in the limit of high polydispersity. From qualitative observations, this size-sorting appears to be a short-range effect.




Fig. 11. Edge Dispersity PDF and CDF of the first 5 cycles of each of the 3 example datasets discussed in paper

D. Edge Dispersity and distance to yield.

4. Simulation for Micro-structural Simulation of Vertex Motion

In an effort to understand the physical, mechanical origins of the phenomenological ‘loop’ nature of the vertex trajectories close to the site of injection, we have developed a simple model that demonstrates that these trajectories are breaking symmetry fundamentally due to the 3-fold symmetry of the microstructure, and a discrepancy in relaxation timescales between vertices and edges.

Our reduced order model takes several approximations. The first of which is that each edge in the system is, at all times, a circular arc segment. The arc segment is fully defined by a single degree of freedom, (defined as h), the height of the arc segment relative to the chord at the center-point between the two vertices. In addition to each edge having a single degree of freedom, each vertex has 2 degrees of freedom (position in x and y). For the simple model system, we studied a single mobile vertex attached to 3 edges — a system of five 2nd order, coupled nonlinear ODEs (3 edges each with a single degree (h_{1-3}), 1 vertex with two spatial degrees (in x and y).

Each coupled ODE is given by a simple force balance which naturally give rise to the familiar equilibrium laws such as Plateau’s law and Laplace’s law.

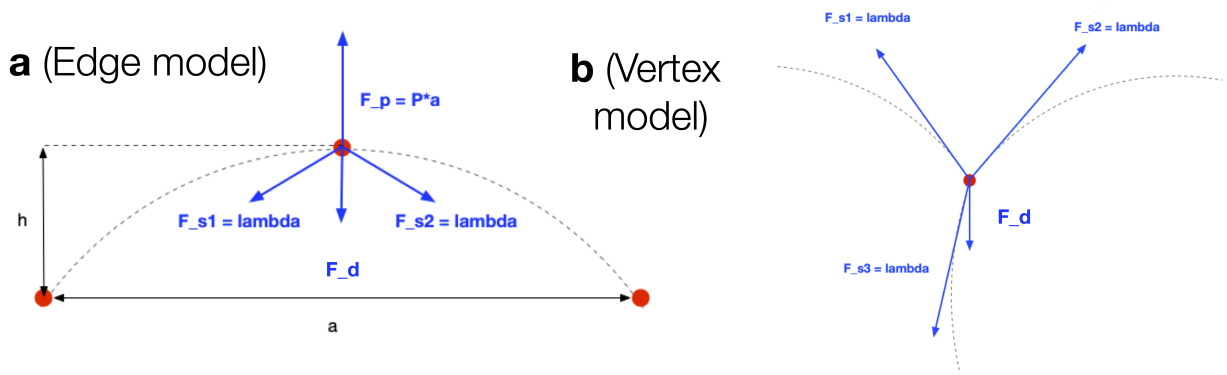


Fig. 12. Free body diagrams for the force balance at of the edge objects versus the vertex objects

The force balance on the vertex is given by

$$\vec{F}_{vertex} = \vec{F}_{\lambda,1} + \vec{F}_{\lambda,2} + \vec{F}_{\lambda,3} + \vec{F}_{D,vertex} \quad [24]$$

Where $\vec{F}_{\lambda,1}, \vec{F}_{\lambda,2}$ and $\vec{F}_{\lambda,3}$ are forces with a magnitude of λ (surface tension), and a direction exactly tangent to the edge at the position of the vertex. The orientation of these three forces defined by the edge curvature through this tangent constraint, coupling the vertex motion to the edge motion.

The force balance on the edge node in the center of the edge is given by:

$$\vec{F}_{edge} = (\vec{P}_{Diff})c + \vec{F}_\lambda + \vec{F}_{D,edge} \quad [25]$$

Where \vec{P}_{Diff} is the pressure difference across the edge, c is the chord length between the edges, \vec{F}_λ is the force normal to chord from surface tension, and $\vec{F}_{D,edge}$ is a nonlinear damping term for the edge curvature, proportional to the velocity to the 2/3 power, as suggested by (43, 44).

The boundaries of the system are fixed / pinned in this model, such that each edge has a distal vertex that cannot move its position.

In this model we can study the motion of the mobile vertex in response to a step function in pressure of the 3 surrounding bubbles.

The total force on each node (given in equations 24 and 25) is used to simultaneously solve the second-order system numerically, using a Runge-Kutta integration method.

The purpose of building an inertial solver is to clearly demonstrate the mechanics of the problem in terms of known underlying mechanisms that drive micro-structural dynamics (e.g., pressure, surface tension). We chose not to use a gradient-descent analog of this model, specifically because it slightly obfuscates the underlying forces involved by enforcing a cost function based on volume conservation. However, it is noted that an energy-based minimization method should also work to demonstrate these phenomena, if the motion of the vertices and the edges have different relaxation timescales. In other words, the looping and non-closure phenomena we demonstrate with this model do not appear to be inherently inertial in origin, as the behavior is extremely robust to large changes in the "mass" terms.




Fig. 13. Simulations of three-edge model demonstrate that the edges have faster relaxation time than vertices in addition to a subtle phase lead.

5. Multi-bubble Simulation Methods

We extend the dynamical simulations (the framework of which is described in the previous section) of edge and vertex dynamics to a more realistic system of bubbles.

A. Galilean Transform on Edge Motion Equations. When we study edges in a foam, each edge's motion is complicated by the fact that the entire edge may be moving significantly (moving reference frame because of the fact that the two vertices are translating). Before extending the edge model to full foam systems (as discussed in figure 5h-g), we must be careful to account for the motion of the reference frame of the edge.

This is accomplished with a Galilean transform correction on the acceleration of the edge in the direction normal to the chord of the edge \vec{F}_{edge} :

$$\frac{d^2h}{dt^2} = \frac{1}{m_{edge}} ((\vec{P}_{Diff})c + \vec{F}_\lambda + \vec{F}_{D,edge}) + \frac{dV_{frame,\hat{k}}}{dt} \quad [26]$$

The acceleration of the edge height, $\frac{d^2h}{dt^2}$, is corrected by a small term to account for the reference frame acceleration in the \hat{k} direction (perpendicular to c , the chord of the edge). This correction term uses the current acceleration of the vertex nodes, computed earlier within the same timestep. This reference frame correction is a common issue when dealing with height-function treatments on free-boundaries(45).

The limitations of imposing this reference frame correction is that errors will accumulate in the edge shape approximation if significant rotations and or significant shear forces are applied to the edge. For the modelling of foam in response to predominantly normal forces (inflation and deflation; tension or compression) these approximations are adequate.

Thus, this model would perform relatively poorly in modelling foam flows with significant shear forces (which are often visible in the foam by the observation of non-circular arc edge segments). In such cases, one is advised to consider using other conventional methods (46, 58) which use fewer approximations.

B. Bubble pressure modelled as an ideal gas. Bubble pressures were modelled as an ideal gas:

$$P_i V_i^\gamma = n_i RT \tag{27}$$

Where at each timestep, the pressure in each bubble P_i is computed from the volume of the bubble V_i . At the start of each simulation, $n_i RT$ is given (e.g. assigning a fixed number of gas molecules to each bubble). It is assumed gas molecules do not migrate across films in this model. The volume is computed at each timestep from the x, y position of the vertices and the height of the edges. γ is the adiabatic exponent which in these simulations is set to 1. One could vary the gas compressibility by modifying this exponent.

Peak-to-trough maximum variation in pressure on the passive bubbles is approximately $3 * 10^{-4}$ psi. Volume variation is $2.8 * 10^{-5}$ for $\frac{\Delta V}{V_{initial}}$ or stated differently, peak-to-trough volume change is only 0.003 percent of the initial original volume. This is shown in figure 14.

The point here is to note that although we treat the bubbles as an ideal gas, the net volume variation within a bubble is exceptionally small. Because we are working with only polygons and circular-arc segments, the computation of bubble volumes is an analytical function not dependent on a mesh or grid size. This gives us the volume and pressure in an analytical form (independent of a mesh size).




Fig. 14. Simulation Method Sanity Check: Peak-to-trough maximum variation in pressure on the passive bubbles is approximately $3 * 10^{-4}$ psi. Volume variation is $2.8 * 10^{-5}$ for $\Delta V / V_{initial}$ or stated differently, peak-to-trough volume change is only 0.003 percent of the initial original volume.

On benefit to such a model is that more complex processes that are highly dependent on these subtle changes in pressures could be easily added to this model. For example, because at every step we track the volume and pressure of each bubble, it would be a simple extension to add the dynamics of coarsening to such a model – where n_i for each bubble is a function of time based on the pressure differences and length of the edges.

One could also represent rough empirical behaviors of interior liquid flows in the foam films by allowing the mass terms of the edges of vertices to be functions of time, governed by the properties of the fluid. Such extensions are feasible additions to this framework.

6. Spring Network Models

As shown in Figure 15, we can apply the same lock-in amplifier method (SI A) to a simple linear spring network to compare with foam-like systems. We constructed a network where force balances on the nodes are given by the attached springs and a drag force (Fig 15a). The springs varied in length, but all had the same linear spring constant, k . An additional boundary force was applied to those nodes adjacent to the active bubble.

The resulting displacement of the network in response to a single inflation-deflation is shown in Fig 15b.

The lock-in amplifier technique is shown in figure 15c, where red dots correspond to nodes with a CW swirl and blue nodes correspond to nodes with a CCW swirl. We do not observe coherent large-scale CW/CCW in these linear systems, unlike the foam we study in section .




Fig. 15. A simple spring model network demonstrating very minimal local cooperation in an inflating and deflating network. In this system, cells do not conserve volume, which screens the inflation and deflation very locally to the active site.

7. Methods for Segmentation of Experimental Data

We have developed a set of segmentation and tracking methods for these cellular systems, which we provide upon reasonable request. These image-processing methods are relatively standard in the study of foams and emulsions, and centrally utilize the MATLAB built-in "Watershed" algorithm, which uses the Fernand-Meyer method (59).

8. Measuring CW/CCW "Swirl" Directionality in Vertex Trajectories

To measure the direction of swirl of the vertex trajectories, we implemented a lock-in amplifier technique as described in detail below.

A. Lock-In Amplifier for Binary Directional Measurement. By treating the y-axis as an imaginary parameter and the x-axis as a real parameter, then we can represent point motion using a complex variable,

$$P_{exp} = K e^x, \quad [28]$$

that varies as a function of time. P is the position of the point, K is an amplitude variable which is considered strictly positive, t is time, and x is a variable with the form of

$$x = i(\omega_{exp} t + \phi). \quad [29]$$

ω_{exp} is either approximately $-\omega$, representing clockwise motion, or $+\omega$, representing counterclockwise motion. The motion of the point also starts off in an unknown orientation which can be taken into account by including the variable ϕ , as a phase offset for the experimental data. In order to test the hypothesis of whether the experimental variable is closer to $-\omega$ or $+\omega$ we can create two comparison cases, one for clockwise motion,

$$P_{CW} = e^{-i\omega t} \quad [30]$$

and a similar comparison for counterclockwise motion,

$$P_{CCW} = e^{i\omega t}. \quad [31]$$

Comparison of the hypothetical cases with the experimental case can be accomplished by dividing the experimental motion variable with the reference variables and integrating over a full cycle, giving two cases,

$$A = \int_0^{\frac{2\pi}{\omega}} \frac{P_{exp}}{P_{CW}} dt = \int_0^{\frac{2\pi}{\omega}} K e^{i\omega_{exp} t + i\omega t + i\phi} dt \quad [32]$$

and the counter clockwise test of

$$B = \int_0^{\frac{2\pi}{\omega}} \frac{P_{exp}}{P_{CCW}} dt = \int_0^{\frac{2\pi}{\omega}} K e^{i\omega_{exp} t - i\omega t + i\phi} dt. \quad [33]$$

In order to determine which hypothesis is more valid, we can take the ratio of A and B

$$C = \frac{A}{B} = \left(\frac{e^{i2\pi(\omega_{exp} + \omega)} - 1}{\omega_{exp} + \omega} \right) * \left(\frac{\omega_{exp} - \omega}{e^{i2\pi(\omega_{exp} - \omega)} - 1} \right) = \frac{\omega_{exp} - \omega}{\omega_{exp} + \omega} \quad [34]$$

which allows us to cancel out the factor of ϕ and the amplitude parameter K from the experimental data, allowing the system to become wholly dependant on the comparison frequency, ω and the experimental rotation frequency ω_{exp} . If C is greater than 1 then it indicates that hypothesis A proved a better fit than hypothesis B, and the point moves in a clockwise direction. A value for C less than 1, conversely indicates motion in a counterclockwise direction. While this method is simple, it provides a robust framework for detecting the time based motion of the directionality of points in spite of noise and the reality that K varies throughout the motion of the trajectory.

9. Cellular Segmentation (based on Watershed algorithm)

For each dataset video we can extract all possible information about the structure in each frame:

- the voxels (or cells),
- the location and area of each cell,
- the connectivity of which cells are neighbors,
- the events flagged for when and where T1 events occur,
- the vertex locations throughout the video,
- the vertex connectivity,
- the edges (in 3 formats: straight-line approximations, binary data representing the real edge shape curve, and circular arc fits, reducing each edge to only a length and a curvature DOF).

Below we provide an overview of each sequential step in the process. This is also depicted visually in SI Movie 2.

The first step is removing background noise and gradients from the image with a series of background-subtraction steps. This step may be unnecessary for extremely clean data with exceptionally flat illumination.

The second step is using the watershed algorithm to segment each image into cellular regions.

$$\text{watOut}(\text{frameNumber}) == 1$$

contains a binary skeleton of the foam which is exactly 1 pixel wide (regardless of the foam volume fraction).

$$\text{watOut}(\text{frameNumber}) = \text{localID}$$

corresponds to the image region associated with a voxel, named by that local ID. Properties of the voxel are extracted from each of these subregions.

A. Network Connectivity. The centroid of each cell is computed and we use a region adjacency method(58) to identify which cells are neighbors. This results in a connectivity matrix, describing which cells are adjacent (with “local” ID numbers that describe the cell in a frame).

A simple tracking algorithm (60) is then used to associate the local IDs of each cell in a given frame to a global ID for the cell across all frames in a video, such that each individual voxel has a single ID number in a video. Each cell is saved as a voxel object, with properties that record its centroid location and time series of local IDs over the video. Segmentation errors will sometimes cause the voxel to persist for less than the entire video; typically this is due to tracking assigning the voxel a new global ID at some point in the video. This error may be tuned using fitting parameters in track.m.

The area of each cell is computed and saved.

A connectivity tensor is built (from the set of local connectivity matrices), that associates the connectivity of the global voxel IDs in the entire video.

B. Vertex Identification. Further, we track the location of each vertex in the foam. This is done by using the skeleton from the segmentation step. Because the watershed algorithm always produces a skeleton that is exactly 1 pixel in width in all places, the vertex locations have a particular signature 3x3 pattern. A 3x3 by 50 tensor of all possible vertex patterns is built. The binary skeleton of each frame is convolved with each 3x3 pattern in a loop (this is done in parallel for faster compute time), and matches to the pattern are exact. Error is exceptionally low and typically occurs only in exact locations of T1 events (where 4 edges meet). This is a potential method to perform additional checks on the T1 Identification methods previously described.

The vertexes are then associated with each voxel object in each frame. For each voxel region, the boundary is computationally grown or expanded by a small amount (tunable, but typically only 1 pixel is necessary). Vertexes within the boundary are associated with that voxel. This is repeated for all voxels in each frame.

Vertexes are then tracked throughout the video using the same tracking algorithm as for the centroids, to assign a persistent identity to each junction.

C. Edge Identification. To compute connectivity of the vertexes (to achieve data on edge length distributions), the vertexes associated with each cell are connected counter-clockwise for each voxel object, populating a connectivity matrix for the vertexes. This connectivity matrix is 3-regular, and has an interesting mathematical relationship to the connectivity matrix of cellular centers (which is irregular).

This method is redundant (each edge is counted twice), however the result is exceptionally accurate (no missed edge has yet been identified). Each edge is approximated by a straight line between the voxels — however, we also have the binary edge data (which contains all curvature information as shown in Fig. 17) in the original skeleton. Straight-line edge data is shown in supplementary figure 16. Early algorithms are under development that have used the binary edge and the straight-line abstracted edge that connects the two vertexes to compute curvature for each edge, based on a simple circular-arc fitting method. We then use these circular arcs as edge models to compute the 3 internal angles at each of the vertexes.




Fig. 16. Adjacency plot superposed over original foam image (this is an abnormally wet foam, not one from the main paper datasets). Green lines represent connections between neighboring voxels.

10. Methods for T1 Transition Identification

This voxel connectivity tensor is analyzed for particular local connectivity changes that are unique to a T1 event. A series of initially detected events are cataloged. However, due to known errors in the region adjacency method (to improve speed of the algorithm), a filtering step is necessary to identify which events are truly T1 events versus connectivity events that arise from errors in the segmentation and adjacency methods. This filtering checks that each cell in an event is persistent for a certain minimum time before and after the event, and that the connectivity in the region obeys certain continuities.

A final list of T1 events is produced that contains the location of the event, the frame number in which it occurs, and the global ID numbers of each of the cells that was involved.

11. Data and Code Availability

Data available here (please cite this article if this data is used directly in any way in future work by the community): <https://tinyurl.com/2eey79up>. Code for image segmentation and experimental data feature extraction will be provided upon




Fig. 17. Vertex Finder demonstrates high accuracy in identifying vertex positions (red)

reasonable request.




Fig. 18. T1 visualization on top of Adjacency plot. Green lines represent connections between neighboring voxels, yellow lines represent foam edges, and red-shaded cells are where T1s have occurred within a frame.