Topological metric detects hidden order in disordered media

Dominic J. Skinner,[1] Boya Song,[1] Hannah Jeckel,[2,3] Eric Jelli,[2,3] Knut Drescher,[2,3] and Jörn Dunkel[1]

[1] *Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307, USA*
[2] *Max Planck Institute for Terrestrial Microbiology, 35043 Marburg, Germany*
[3] *Department of Physics, Philipps-Universität Marburg, 35043 Marburg, Germany*

(Dated: December 30, 2021)

Recent advances in microscopy techniques make it possible to study the growth, dynamics, and response of complex biophysical systems at single-cell resolution, from bacterial communities to tissues and organoids. In contrast to ordered crystals, it is less obvious how one can reliably distinguish two amorphous yet structurally different cellular materials. Here, we introduce a topological earth mover's (TEM) distance between disordered structures that compares local graph neighborhoods of the microscopic cell-centroid networks. Leveraging structural information contained in the neighborbood motif distributions, the TEM metric allows an interpretable reconstruction of equilibrium and non-equilibrium phase spaces and embedded pathways from static system snapshots alone. Applied to cell-resolution imaging data, the framework recovers time-ordering without prior knowledge about the underlying dynamics, revealing that fly wing development solves a topological optimal transport problem. Extending our topological analysis to bacterial swarms, we find a universal neighborhood size distribution consistent with a Tracy-Widom law.

Discrete particulate objects, from atoms to cells, compose the majority of physical and living systems. Modern microscopy and simulation techniques enable us to study the elementary building blocks of solids [1, 2], colloidal and granular materials [3–5], bacterial biofilms [6, 7], and tissues [8] with unprecedented resolution over large scales. These experimental and computational advances have highlighted the importance of local spatial organization [9] and disorder [10] for the global behaviors of both equilibrium and non-equilibrium materials, spurring substantial theoretical efforts to link discrete microstructure with macroscale properties. Prime examples include the recent successful characterizations of epithelial cell (EC) layers and other cellular materials through the geometric [11, 12] and topological [13–16] analysis of Voronoi and Delaunay tesselations [17]. In spite of such major progress, high-resolution data continue to pose fundamental conceptual and practical challenges regarding the proper classification of discrete physical and biological structures. Specifically, it is still unclear whether one can recover parametric embeddings, phase space dimensions and time ordering from a topological analysis of static snapshots alone, and whether such analysis can help reveal the governing principles of multicellular development.

To tackle these problems, we introduce here a topological earth mover's (TEM) distance by combining ideas from statistical topology [13–15, 18, 19] and optimal transport theory [20, 21] with non-equilibrium statistical mechanics [22]. The TEM metric compares two discrete material structures by quantifying the statistical differences in the local network topology of their Delaunay triangulations (Fig. 1). Intuitively, computing TEM(A, B) amounts to estimating the smallest number of edge-flips needed to make the local network topology of material A statistically indistinguishable from the local network topology of material B. Physically, this procedure can be interpreted as finding the average lowest-energy path connecting two disordered structures, and we provide an efficient algorithm for realizing this computationally demanding task for systems with $\sim 10^6$ particles [23].

To demonstrate the practical potential of this framework for the analysis of both equilibrium and non-equilibrium systems, we present a broad set of applications: First, we show that the TEM metric successfully distinguishes jammed disordered packings of both monodisperse and polydisperse ellipsoids. Thereafter, we use the TEM framework to reconstruct the non-equilibrium phase diagram of active Brownian particle (ABP) simulations without recourse to time resolved data. Next, by measuring the pairwise TEM distances between unsorted experimental images of a developing fruit fly wing, we are able to reconstruct their temporal ordering and discover that wing development follows an optimal transport geodesic, suggesting a previously unrecognized optimization principle in tissue development. Finally, by extending our topological analysis to single-cell resolution data from bacterial swarming experiments, we observe a universal neighborhood size distribution consistent with a Tracy-Widom law.

To define the TEM distance, we consider the specific example of a two-dimensional (2D) cell layer as shown in Fig. 1(a), although all subsequent definitions generalize to arbitrary point sets in $\mathbb{R}^2$ or $\mathbb{R}^3$. Our starting point is the Delaunay triangulation [17] of the cell centroid positions as shown in Fig. 1(a). In practice, it is often sufficient to take the positions of the EC nuclei as vertices of the Delaunay network [25]. If two random realizations of such networks are generated by the same physical or biological process, they will have different vertex positions and topology, but their local statistical properties (local connectivity patterns, etc.) will be identical provided

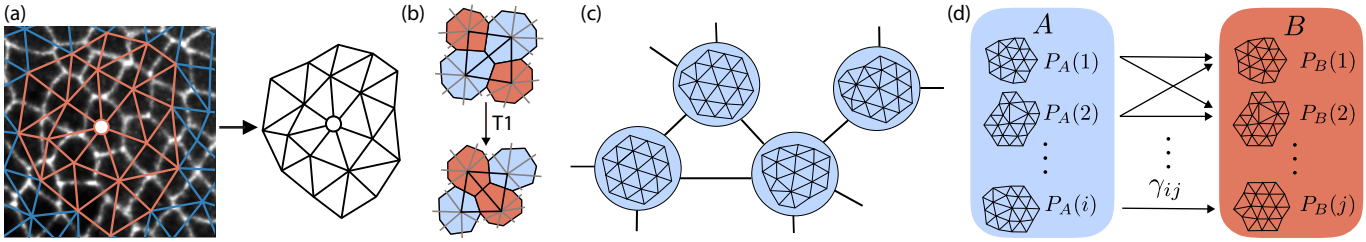

FIG. 1. Key conceptual steps for calculating the topological earth mover's (TEM) distance illustrated for an epithelial cell layer. (a) Experimental image of epithelial cell layer from *Drosophila* embryo (adapted with permission from Ref. [8, 24]) with Delaunay triangulation overlayed (blue, red). The local neighborhood network of radius $r = 2$ (red) is shown for a selected vertex (white circle). (b) A topological T1 transition, corresponding to an exchange of neighbors, is reflected in a change of the local radius-2 neighborhood. Voronoi cells (red, blue) with Delaunay triangulation overlayed (black, grey) are shown. (c) Nodes of the flip graph correspond to networks, and are connected by an edge if the networks are one T1 transition (or flip), away from each other. (d) The map γ_{ij} transports the distribution A to the distribution B.

the networks are sufficiently large. This fact has been exploited previously to define entropic [14] and earth mover's distances between cell complexes [15]. Here, we extend these ideas to define a physically motivated topological metric that measures statistical differences in the local Delaunay triangulations around vertices. Specifically, we define for each vertex a local neighborhood of radius r, which consists of all the vertices that are not more than r edges away from the central vertex (see red subgraph corresponding to $r = 2$ in Fig. 1a). We found that $r = 2$ suffices for many practical applications [23]. Although r can, in principle, be chosen arbitrarily large, TEM computations become expensive for a larger neighborhoods; we therefore focus on the case $r = 2$ from now on. The local neighborhoods of two vertices are of the same topological type i if they are graph-isomorphic. Counting the occurrences of the various neighborhood types i across all vertices yields a probability distribution $P(i) \geq 0$ that characterizes the topo-statistical state of the cell network.

To provide an intuitive physical motivation for the TEM metric, let us recall that the Delaunay network is invariant under infinitesimal perturbations and can change only through a topological T1 transition (Fig. 1b). For EC layers there is an energy barrier to T1 transitions [26], and so the energy cost to transform from one neighborhood type to another is directly related to the number of T1 transitions required. For other packed systems there typically exist similar energetic cost for changing neighbors through T1 transitions. Motivated by this, we can define the energetic distance between two neighborhoods as the minimum number of T1 transitions separating them. This mathematically well-defined metric [23, 27] induces naturally a secondary graph structure, known as the flip graph [28], where nodes correspond to neighborhood types i and are linked with an edge if they are one T1 transition away from each other (Fig. 1c). The minimum path length between two nodes on the flip graph is the smallest number of T1 transitions needed to move between the corresponding neighborhood types. Moreover, the distribution $P(i)$ of neighborhood types in the EC layer can now be viewed as a distribution on the nodes i of the flip graph (blue box in Fig. 1d).

Armed with this intuition, we can now define the TEM distance between the Delaunay triangulations of two materials A and B in a natural manner as the earth mover's or, equivalently, Wasserstein distance [20] between their neighborhood distributions P_A and P_B over the flip graph: If $P_A(i)$ is the probability of neighborhood i occuring in material A, and $P_B(j)$ is the probability of neighborhood j occuring in material B, then a transport map, $\gamma_{ij} \geq 0$, from A to B satisfies $\sum_j \gamma_{ij} = P_A(i)$, $\sum_i \gamma_{ij} = P_B(i)$, see Fig. 1(d). Then, the TEM distance between A and B is

$$\mathrm{TEM}(A, B) = \min_{\gamma} \sum_{ij} \gamma_{ij}\, d(i, j) \qquad (1)$$

where $d(i, j)$ is the distance between the neighborhoods i and j on the flip graph, and the minimum is taken over all possible transport maps $\gamma = (\gamma_{ij})$. We emphasize that, in contrast to widely used entropic distances measures between distributions [14], the definition of TEM uses the physically relevant information encoded in the metric structure $d(i, j)$ of the underlying observable space, which in our case reflects the typical energy cost of a T1 transitions between network motifs. As a consequence, TEM generally outperforms purely entropic Kullback-Leibler/Jensen-Shannon divergences when one needs to distinguish complex structures that are characterized by weakly overlapping distributions; see Ref. [29] and the Supplemental Material [23] for explicit examples.

For large systems, the minimization problem (1) becomes computationally challenging. We combined two algorithmic insights [23] to calculate TEM efficiently for disordered materials with millions of particles. Building on a modification of the Weinberg algorithm [30], our numerical scheme [23] first determines the flip-graph distances $d(i, j)$ of N observed neighborhood motifs in $O(N)$

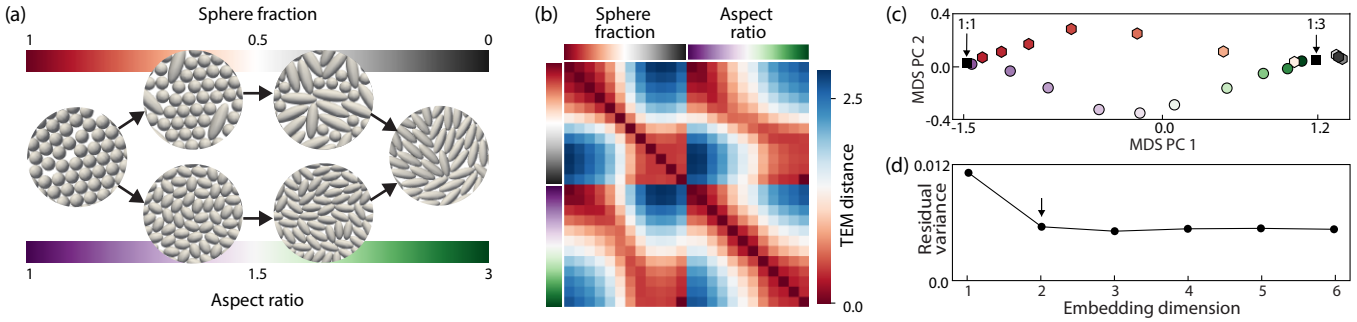

FIG. 2. Poly- and monodisperse packings are distinguished by the TEM distance. (a) Two alternative paths from aspect ratio 1:1 to aspect ratio 1:3 ellipsoid packings. Top is monodisperse with varying aspect ratios, bottom is polydisperse with a mixture of 1:1 and 1:3 aspect ratios. (b) Distance matrix where each pixel represents the distance between two simulated ellipsoid packings, each containing 10,000 ellipsoids. (c) Simulations are embedded in 2D using MDS, recovering two distinct paths, one for monodisperse simulations (hexagons) and one for polydisperse simulations (circles). (d) The residual variance plateaus after embedding dimension 2, correctly identifying the true embedding dimension of the phase space (arrow).

steps. Given $d(i,j)$, the minimization over the transport maps $\{\gamma_{ij}\}$ can be recast as an minimum cost flow problem [20, 23], which is efficiently solved with linear programming [31]. To demonstrate the broad applicability of our TEM framework, we focus in the remainder on applications relevant to current major research areas: colloidal packings, collective far-from-equilibrium dynamics, tissue development, and spatio-temporally heterogeneous multicellular systems.

Recent advances in the fabrication of geometrically complex colloids [3, 4] and confocal imaging techniques [6] have led to a renewed practical and theoretical interest in the characterization of granular [5, 33] and biological materials [6, 7]. Of particular importance in this context are the often fundamentally different behaviors of monodisperse [34] and polydisperse [35] colloidal systems. While the former are much better understood theoretically, the latter are often practically more relevant to natural systems and processes, such as particle segregation seen in industrial agriculture, cereals, or avalanches [36]. To demonstrate the usefulness of the TEM framework for capturing the essential topo-statistical differences between and across mono- and polydisperse systems, we generated jammed disordered packings of 10,000 ellipsoids using an event-driven packing code [37]. Specifically, we were interested in distinguishing two different pathways for transitioning from a monodisperse packing of spheres (ellipsoids with aspect ratio 1:1) to a monodisperse packing of ellipsoids with aspect ratio 1:3 (Fig. 2a). The first 'monodisperse' transition path was realized by simulating 12 monodisperse packings of ellipsoids with aspect ratios varying from 1:1 to 1:3 (bottom path in Fig. 2a). The second 'polydisperse' transition path was realized by simulating 12 different binary mixtures of 1:1 and 1:3 ellipsoids (top path in Fig. 2a). Computing the TEM distances between all 24×24 pairs of simulations produces the symmetric TEM distance matrix shown in Fig. 2(b). Given this matrix, it is natural to seek a faithful low-dimensional embedding in Euclidean space $\mathbb{R}^d$ that approximately preserves the TEM distance structure. To construct the embedding we choose Multi-Dimensional Scaling (MDS), a generalized principal component (PC) analysis based on the TEM distance [38]. Since each pathway corresponds to a 1D manifold (as only one parameter is varied in each case), the phase space can be embedded in $\mathbb{R}^2$; indeed the $\mathbb{R}^2$ embedding clearly distinguishes the two different pathways (Fig. 2c). To find the dimensionality of the phase space, we calculate the residual variance which plateaus at the relevant dimension [23, 39] and correctly identifies the ellipsoid embedding as 2D (Fig. 2d). We show in the Supplementary Information [23] that the same approach can be used to infer the non-equilibrium phase space of active Brownian particle (ABP) simulations [40] from instantaneous system configurations. More broadly, these examples illustrate how the TEM metric can discover phase spaces from configurational snapshots alone.

In the remainder, we show that the topological analysis of data from two recent experiments [8, 32] can reveal previously unrecognized biophysical optimization principles and universal statistical signatures. We begin by considering shuffled images (Fig. 3a-e) of developing fruit fly embryo wings [8, 24]. Using the Delaunay triangulation of the cell centroids, hierarchical clustering [41] of the TEM distance matrix of the shuffled images (Fig. 3f) reveals three developmental main phases (Fig. 3g). The resulting MDS embedding is essentially 1D, with the first principal component corresponding to time (Fig. 3h), and thus restores the temporal order of the data (Fig. 3i). This shows how the TEM framework can be used to infer temporal ordering from ensemble measurements [42]. More importantly, however, the TEM analysis reveals the developmental trajectory of the fly wing follows a topological geodesic, a continuous curve that minimizes the total length with respect to the TEM distance. Whilst earth mover's geodesics are in general

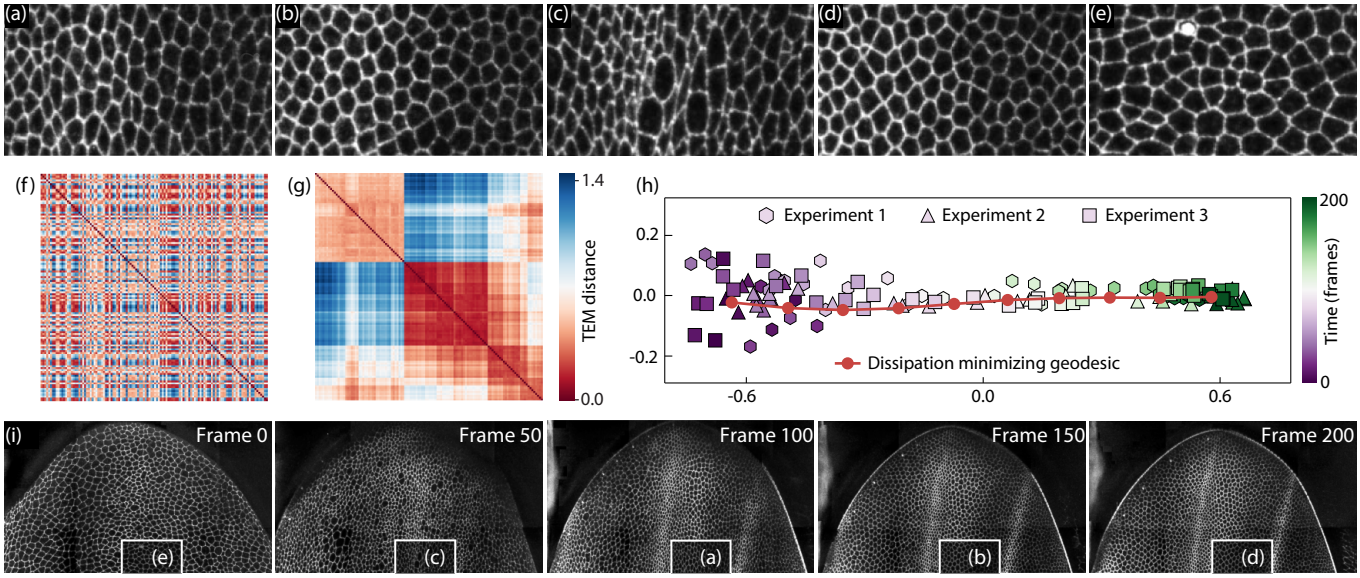

FIG. 3. Developing *Drosophila* embryos solve topological optimal transport problem, with MDS embedding recovering temporal order. (a-e) Enlarged images showing epithelial cells at unknown times. Three experiments and 40 time frames per experiment were used. (f) Matrix of TEM distances between unsorted experimental images has no apparent structure. (g) TEM distance matrix sorted by hierarchical clustering shows approximately three phases. (h) 2D MDS embedding recovers the temporal order as the principal component. Included in the embedding are intermediate stages of a dissipation-minimizing [23] geodesic between average start and end states, which the data falls on. (i) The correct temporal ordering of the image time series is recovered. The white boxes show the source of images (a-e). Data adapted with permission from [8, 24].

not unique, a unique path can be found by additionally minimizing transport dissipation [21, 23]. The data fluctuates closely around this minimum-dissipation geodesic (red curve in Fig. 3h), meaning that fly wing development approximately solves a dissipation-constrained topological optimal transport problem.

Finally, we provide a more detailed characterization of the neighborhood motif distributions in 2D non-equilibrium systems. To this end, we analyze recent bacterial swarming experiments [32] using machine learning [23] to identify individual cells (Fig. 4a,b). By determining the motif size distributions for snapshots taken at different space-time locations in a growing swarm, we find that both mean and variance vary systematically with space, time, and cell density (Fig. 4c-e). Strikingly, after rescaling to zero mean and unit variance, the combined motif size distribution closely matches [23] a universal Tracy-Widom (TW) distribution (Fig. 4f). TW distributions were recently reported for growing fluctuating fronts [43], dynamics of self-assembly [44], active particle dynamics [45, 46], and phase transitions between strongly and weakly coupling regimes [47]. We also find TW motif size distributions in the ABP and fly wing data when subsampling from the liquid-like phase [23], suggesting that TW distributions play a central role in the topo-statistics of non-equilbrium systems.

To conclude, the TEM metric framework will be broadly applicable, from single-cell RNA-sequencing [48], cryo-electron microscopy [49] and organoid characterization [50], to structural transitions in living [6, 11] and nonliving [3] matter. In particular, it enables a direct comparison of the topological statistical properties of a wide range of fundamentally different systems, the only requirement being that transitions between basic motifs (Delaunay neighborhood structures, DNA strings, etc.) can be mapped onto a joint flip-graph structure.

This work was supported by a MathWorks Fellowship (D.J.S.), a James S. McDonnell Foundation Complex Systems Scholar Award (J.D.), and the Robert E. Collins Distinguished Scholar Fund (J.D.). We thank Martin Abt for help with cell segmentation and the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing HPC resources.

[1] G. Binnig and H. Rohrer, Rev. Mod. Phys. **59**, 615 (1987).
[2] Y. Sugimoto, P. Pou, M. Abe, P. Jelinek, R. Pérez, S. Morita, and Ó. Custance, Nature **446**, 64 (2007).
[3] S. C. Glotzer and M. J. Solomon, Nature Materials **6**, 557 (2007).
[4] V. Soni, E. Bililign, S. Magkiriadou, S. Sacanna, D. Bartolo, M. Shelley, and W. Irvine, Nat. Phys. **15**, 1188 (2019).
[5] S. Nauer, L. Böttcher, and M. A. Porter, J. Complex Networks , doi:10.1093/comnet/cnz037 (2019).
[6] R. Hartmann, P. Singh, P. Pearce, R. Mok, B. Song, F. Díaz-Pascual, J. Dunkel, and K. Drescher, Nat. Phys.

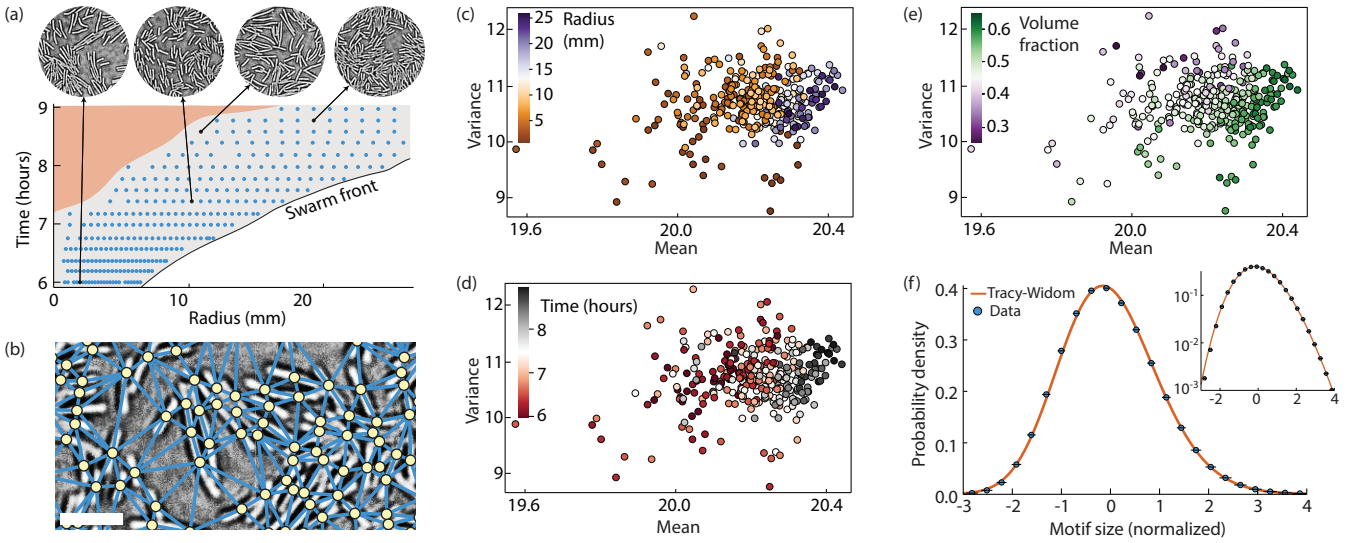

FIG. 4. Motif sizes in heterogenous bacterial swarms exhibit universal non-equilibrium statistics. (a) An azimuthally symmetric bacterial swarm expands as motile cells grow and divide [32]. We analyzed configurational snaphots at different space-time positions in the swarming monolayer phase (blue dots; select examples at top); the central multilayered biofilm phase (red) was excluded. (b) Aided by a machine learning algorithm, cell centroids (yellow circles) and their Delaunay tesselation (blue lines) were determined. Scale bar: 10μm. (c-e) The distribution of motif sizes, defined as the number of vertices per $r = 2$ neighborhood, is heterogenous in space and time. The mean motif size increases, on average, with radial distance and time (c,d). The variance decreases with area filling fraction (e), reflecting more ordered cell packings in dense swarming regions. (f) After normalizing to zero mean and unit variance, the combined histogram over all space-time snaphots (circles) follows a universal Tracy-Widom distribution (line). Horizontal error bars represent the standard deviation within a bin (inset: log scale plot).

15, 251 (2019).

[7] D. Dell'Arciprete, M. Blow, A. Brown, F. Farrell, J. Lintuvuori, A. McVey, D. Marenduzzo, and W. Poon, Nat. Commun. **9**, 4190 (2018).

[8] R. Etournay, M. Popović, M. Merkel, A. Nandi, C. Blasse, B. Aigouy, H. Brandl, G. Myers, G. Salbreux, F. Jülicher, and S. Eaton, eLife **4**, e07090 (2015).

[9] P. Pearce, B. Song, D. J. Skinner, R. Mok, R. Hartmann, P. K. Singh, H. Jeckel, J. S. Oishi, K. Drescher, and J. Dunkel, Phys. Rev. Lett. **123**, 258101 (2019).

[10] C. P. Goodrich, A. J. Liu, and S. R. Nagel, Nature Physics **10**, 578 (2014).

[11] D. Bi, X. Yang, M. C. Marchetti, and M. L. Manning, Phys. Rev. X **6**, 021011 (2016).

[12] D. Bi, J. H. Lopez, J. M. Schwarz, and M. L. Manning, Nature Physics **11**, 1074 (2015).

[13] E. A. Lazar, J. K. Mason, R. D. MacPherson, and D. J. Srolovitz, Phys. Rev. Lett. **109**, 095505 (2012).

[14] J. K. Mason, E. A. Lazar, R. D. MacPherson, and D. J. Srolovitz, Phys. Rev. E **86**, 051128 (2012).

[15] B. Schweinhart, J. K. Mason, and R. D. MacPherson, Phys. Rev. E **93**, 062111 (2016).

[16] D. Wenzel, S. Praetorius, and A. Voigt, J. Chem. Phys. **150**, 164108 (2019).

[17] F. Aurenhammer, R. Klein, and D.-T. Lee, *Voronoi Diagrams and Delaunay Triangulations*, 1st ed. (World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2013).

[18] H. Ronellenfitsch, J. Lasser, D. Daly, and E. Katifori, PLoS Comput. Biol. **11**, e1004680 (2015).

[19] J. W. Rocks, A. J. Liu, and E. Katifori, Phys. Rev. Research **2**, 033234 (2020).

[20] M. Essid and J. Solomon, SIAM J. Sci. Comput. **40**, A1961 (2018).

[21] J. Solomon, R. Rustamov, L. Guibas, and A. Butscher, arXiv:1603.06927.

[22] P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro, and Z. Nussinov, Scientific Reports **2**, 329 (2012).

[23] Supplementary Information.

[24] R. Etournay, M. Merkel, M. Popović, H. Brandl, N. Dye, B. Aigouy, G. Salbreux, S. Eaton, and F. Jülicher, eLife **5**, e14334 (2016).

[25] R. van Drongelen, T. Vazquez-Faci, T. A. Huijben, M. van der Zee, and T. Idema, J. Theor. Biol. **454**, 182 (2018).

[26] D. Bi, J. H. Lopez, J. M. Schwarz, and M. L. Manning, Soft Matter **10**, 1885 (2014).

[27] C. L. Lawson, Discrete Math. **3**, 365 (1972).

[28] P. Bose and F. Hurtado, Comput. Geometry **42**, 60 (2009).

[29] J. Solomon, arXiv:1801.07745.

[30] L. Weinberg, IEEE Trans. Circuit Theory **13**, 142 (1966).

[31] I. Dunning, J. Huchette, and M. Lubin, SIAM Review **59**, 295 (2017).

[32] H. Jeckel, E. Jelli, R. Hartmann, P. K. Singh, R. Mok, J. F. Totz, L. Vidakovic, B. Eckhardt, J. Dunkel, and K. Drescher, Proc. Nat. Acad. Sci. **116**, 1489 (2019).

[33] A. Donev, I. Cisse, D. Sachs, E. A. Variano, F. H. Stillinger, R. Connelly, S. Torquato, and P. M. Chaikin, Science **303**, 990 (2004).

[34] S. Torquato and F. H. Stillinger, Rev. Mod. Phys. **82**, 2633 (2010).

[35] I. S. Aranson and L. S. Tsimring, Rev. Mod. Phys. **78**, 641 (2006).

[36] A. Kudrolli, Rep. Prog. Phys. **67**, 209 (2004).
[37] A. Donev, S. Torquato, and F. H. Stillinger, J. Comput. Phys. **202**, 765 (2005).
[38] I. Borg and P. J. F. Groenen, "Modern multidimensional scaling: Theory and applications," (Springer, 2005) Chap. 12, pp. 201–268, 2nd ed.
[39] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, Science **290**, 2319 (2000).
[40] Y. Fily, S. Henkes, and M. C. Marchetti, Soft Matter **10**, 2132 (2014).
[41] J. H. Ward, J. Am. Stat. Assoc. **58**, 236 (1963).
[42] P. Pearce, F. G. Woodhouse, A. Forrow, A. Kelly, H. Kusumaatmaja, and J. Dunkel, Nat. Commun. **10**, 5368 (2019).
[43] K. A. Takeuchi, M. Sano, T. Sasamoto, and H. Spohn, Scientific Reports **1**, 34 (2011).
[44] G. Makey, S. Galioglu, R. Ghaffari, E. D. Engin, G. Yıldırım, Ö. Yavuz, O. Bektaş, Ü. S. Nizam, Ö. Akbulut, Ö. Şahin, *et al.*, Nat. Phys. **16**, 795 (2020).
[45] Z. Chen, J. de Gier, I. Hiki, and T. Sasamoto, Phys. Rev. Lett. **120**, 240601 (2018).
[46] C. A. Tracy and H. Widom, Commun. Math. Phys. **290**, 129 (2009).
[47] S. N. Majumdar and G. Schehr, J. Stat. Mech.: Theory Exp. **2014**, P01012 (2014).
[48] A. K. Shalek, R. Satija, X. Adiconis, R. S. Gertner, J. T. Gaublomme, R. Raychowdhury, S. Schwartz, N. Yosef, C. Malboeuf, D. Lu, *et al.*, Nature **498**, 236 (2013).
[49] N. Fischer, A. L. Konevega, W. Wintermeyer, M. V. Rodnina, and H. Stark, Nature **466**, 329 (2010).
[50] J. F. Dekkers, M. Alieva, L. M. Wellens, H. C. R. Ariese, P. R. Jamieson, A. M. Vonk, G. D. Amatngalim, H. Hu, K. C. Oost, Snippert, *et al.*, Nature Protocols **14**, 1756 (2019).

Supplemental Material: Topological metric detects hidden order in disordered media

Dominic J. Skinner,[1] Boya Song,[1], Hannah Jeckel[2,3], Eric Jelli[2,3], Knut Drescher[2,3], and Jörn Dunkel[1]

[1] *Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307, USA*
[2] *Max Planck Institute for Terrestrial Microbiology, 35043 Marburg, Germany*
[3] *Department of Physics, Philipps-Universität Marburg, 35043 Marburg, Germany*

(Dated: December 30, 2021)

Voronoi diagrams and Delaunay triangulations

Given a set of points X in $\mathbb{R}^n$, the Voronoi tessellation partitions $\mathbb{R}^n$ into regions known as Voronoi cells. A point y is in the Voronoi cell associated with $x \in X$ if y is closer to x than to any other $z \in X \setminus \{x\}$. The Delaunay triangulation is a graph with vertices at the points in X, with two points sharing an edge if their corresponding Voronoi cells share a face. For more properties refer to [S1].

Interior and exterior points

For simulations with periodic boundary conditions the Deluanay graph extends periodically and there are no boundary cases. For simulations or experiments that are not periodic, there exist exterior points on the edge which may have quite different properties from points in the bulk. We do not wish to include edge effects and so we only take local networks of radius r for points at least r edges away from exterior points. Exterior points are identified by calculating the alpha shape, or concave hull [S2]. In short, a point is an exterior point if a circle (or sphere in $\mathbb{R}^3$) of radius α can intersect that point without enclosing any other points. For every point x_i, there is a largest circle which intersects x_i, but does not include any other points; call its radius α_i. For the non-periodic *Drosophila* example, $\alpha = 2\times$(median α_i) identifies cells on the edge without incorrectly identifying interior points (Fig. S1).

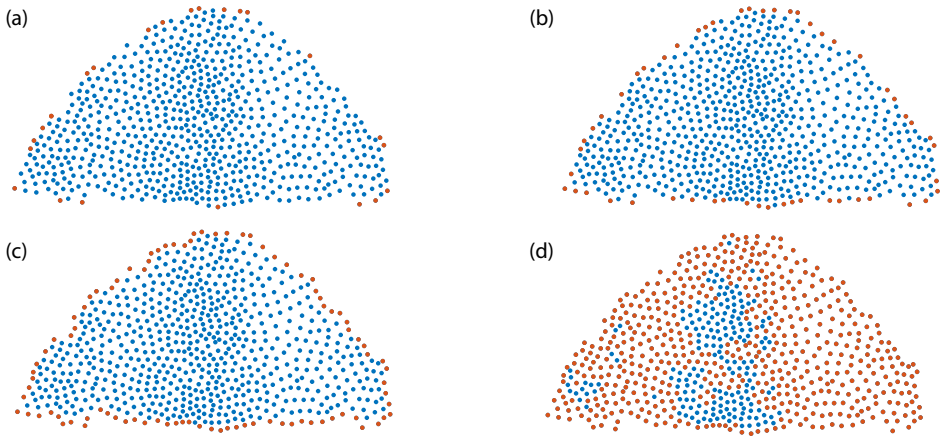

FIG. S1. Taking the concave hull identifies exterior points for segmented cell data from a *Drosophila* embryo [S3]. Cells centroids are shown in red for exterior points, and blue for interior. (a) Exterior cells are identified by the convex hull, which underidentifies exterior points. (b) Edge cells are identified using $\alpha = 10\times$(median α_i), underidentifying exterior points. (c) $\alpha = 2\times$(median α_i) does a reasonable job at identifying exterior points. (d) $\alpha =$(median α_i) overestimates exterior points.

Storing and comparing networks

Many thousands of topologically distinct networks were observed, requiring fast methods to store and compare them. The approach taken here was to replace each network with a vector of integers that uniquely represents its topological type. Once represented as a vector, the networks can be stored as a dictionary with an $O(\ln N)$ cost to read for N topologically distinct graphs. Calculating the probability distribution for N networks is then $O(N \ln N)$,

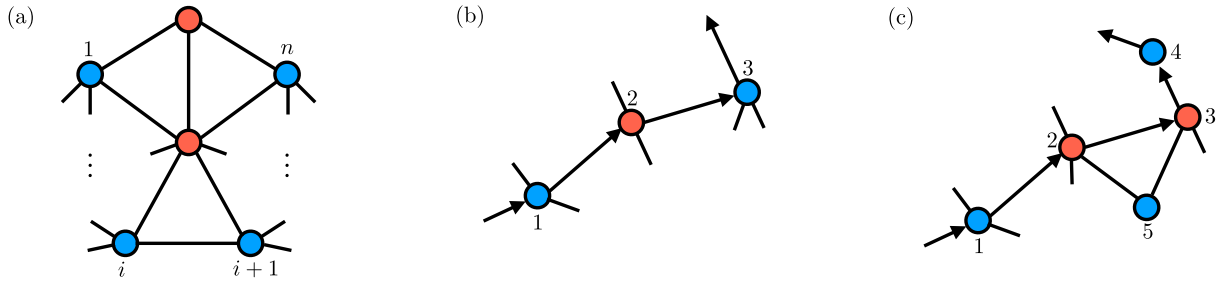

FIG. S2. **Claim:** The local network of radius r formed from a Delaunay triangulation is triply connected.
Proof: First show the local network of radius $r = 1$ is triply connected. Suppose two vertices are removed, if the central vertex is kept, as everything is connected to the central vertex, the graph is still connected. Suppose now that the central vertex and one other vertex are removed, as in (a). Label the remaining vertices by their anti-clockwise ordering about the central vertex for a particular embedding. Then as the local network is a triangulation, the central vertex and vertices i, $i+1$ form a triangle, so vertex i is connected to $i + 1$, hence the network remains connected.
Now consider a local network of arbitrary radius, and a path between two vertices, A and B. If no vertices on this path are removed A and B are still connected. If one vertex is removed, vertex 2 in (b), the path enters the vertex from 1 and leaves toward 3. But 1 and 3 are in the local network of radius 1 around vertex 2, hence are connected, and so an alternative path from A to B can be found.
The final case is when vertex 3 is also deleted, so an alternative path to vertex 4 must be found, as in (c). Here, note that as the network is a triangulation, vertices 2 and 3 are connected to a common vertex, vertex 5 which is in the local network of radius 1 for both. So 1 and 5 are connected as the local network of radius 1 is connected around 2, but 5 and 4 are also connected as the local network of radius 1 is connected around 3. Therefore an alternative path from A to B can be found, hence the network remains connected.

rather than the $O(N^2)$ cost that would be required if one used an algorithm that could only compare two graphs at a time.

To encode the topology we use a modified Weinberg algorithm. The Weinberg algorithm uniquely encodes the topology of a triply connected planar graph, where triply connected means at least three vertices need to be removed to disconnect the graph [S4, S5]. The local network is triply connected, see Fig. S2.

In short, Weinberg's algorithm canonically labels Eulerian circuits, and from all possible Eulerian circuit picks the lexographically first labeling. This labeling is taken as the vector; every isomorphic graph has the same vector and if two graphs have the same vector they are isomorphic. The procedure for finding the Eulerian circuit from a given oriented starting edge is detailed in algorithm 1, and the canonical labeling for a particular starting edge is detailed in algorithm 2. The total algorithm is described in algorithm 3.

Algorithm 1 Trémaux's algorithm for finding an Eulerian circuit in a directed graph with a particular edge to be traversed first [S4].

1: **Input:** A directed graph and a chosen edge.
2: The first vertex is the source of the chosen edge, the first step is to the destination of the edge. No edge is traversed twice and future steps are made according to the following rules:
3: If a new vertex is reached, exit this vertex with the outgoing edge to the right of the edge that you entered from.
4: If a previously visited vertex is reached, exit, if possible, towards the vertex that you were previously at.
5: If a previously visited vertex is reached, and it is not possible to exit towards the vertex that you were previously at, exit to the nearest available outgoing edge to the right.
6: If there are no edges available, the algorithm terminates, and a Eulerian circuit has been found.

Algorithm 2 Canonical labeling for a graph using a specific oriented edge to start. A worked example is shown in Fig. S3.

1: **Input:** an undirected graph and an edge for which a direction is chosen.
2: Replace every edge of the graph with two directed edges oriented in opposite directions.
3: Take the Eulerian circuit starting with the chosen edge according to algorithm 1.
4: Label the starting vertex as 1, and as the circuit is traversed, label every new vertex reached with consecutive integers.
5: The labeling vector for this starting edge is the ordered record of vertices that are seen as the circuit is traversed (so if a vertex is crossed n times it appears in the vector n times).

Algorithm 3 Algorithm for finding a graph's Weinberg vector (modified from [S4])

1: **Input:** A local network (an undirected graph with central vertex).
2: **for** edges connected to central vertex **do**
3: Find canonical labeling as described in algorithm 2 with the edge oriented outwards.
4: **end for**
5: Do the above for the mirrored embedding of the network as well. Collect all resulting labelings.
6: All of the labelings are vectors of integers with length $2|E|$. Lexographically sort these vectors and take the first sorted vector. This first vector is the Weinberg vector for the network.

Weinberg's algorithm as described in [S4] differers from algorithm 3 in that it traverses every edge in both directions rather than just the edges originating at the central vertex and so recognizes the same graph with different embeddings as isomorphic. The local Delaunay network is a near triangulation meaning all of the faces are triangular except the one at infinity [S6], so there is only one embedding that we will observe (together with the mirrored embedding). We also only define isomorphism between egocentric networks to mean the networks are isomorphic and they have the same central vertex (it is possible for an egocentric network to have two possible candidates for the central vertex, although these are quite rare in practice). Therefore to check if two egocentric networks are isomorphic we need only consider the labelings that start by moving away from the central vertex. If they are isomorphic they share the same labelings. The resulting vector will not be the same as the vector calculated by trying all edges, but still works as a topological identifier when compared only to other vectors calculated in the same manner.

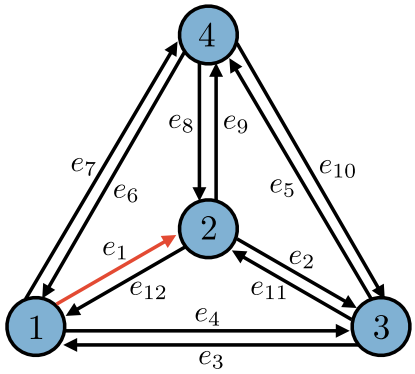

FIG. S3. Calculation of a labeling vector from a specific starting edge for an example local network of radius 1. The edge labeled e_1 is the starting edge, and the first step is from vertex 1 to vertex 2. Next we exit to the right along edge e_2 to vertex 3. We have reached a new vertex so we exit to the right along edge e_3. Vertex 1 has been visited before so we take edge e_4 back to vertex 3. Vertex 3 has been visited before, but we have already used edge e_3, so take the nearest available edge to the right, edge e_5, onto vertex 4. We continue as before, using algorithm 2, until we run out of edges. The edges are labeled in the order in which they are used. The canonical labeling is the record of which vertices are visited, which is (1,2,3,1,3,4,1,4,2,4,3,2,1).

The algorithm in [S4] finds the Weinberg vector in $O(E^2)$ time, as finding the path and labeling it (which can be done simultaneously) takes $O(E)$ time, and this must be done for $2|E|$ edges. By only taking the edges that start at the central vertex, which is typically $O(6)$, the calculation grows like $O(E)$ as the size of the local network grows. For local networks of radius $r = 2$ this typically allows the algorithm to run 10-20 times faster.

Calculating the flip graph

Given N observed networks, we wish to calculate which are connected by an edge in the flip graph, from which the minimum path length between two networks gives a measure of distance between them. This measure of distance is similar to the Levenshtein distance between two strings, although in that case distances are calculated, as needed, using dynamic programming [S7]. Here, path lengths on the flip graph give distances, but this requires checking if up to $N(N-1)/2$ edges exist. Instead we are able to calculate the flip graph in $O(NE^2)$ time using algorithm 4. The central insight is that after a T1 transition or flip, path distances to the central vertex either do not increase or do not decrease, as proved in Fig. S4. After flipping, the new graph need not be a local network, but due to the claim in Fig. S4, for one of the local networks, flipping means the other local network is a subgraph of the flipped graph

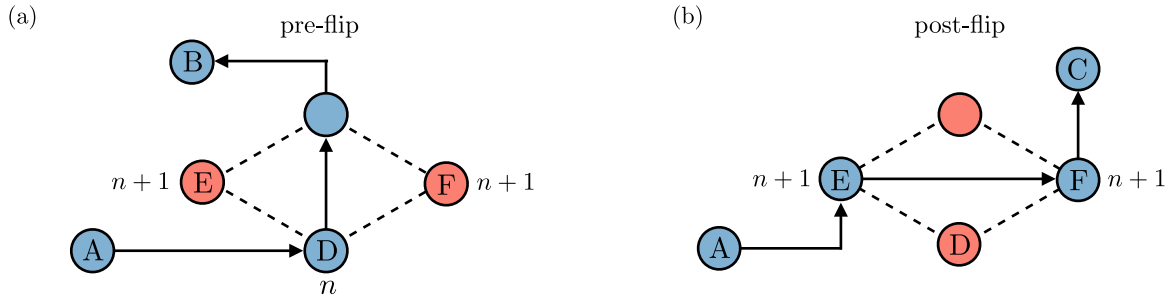

FIG. S4. **Claim:** A topological transition, or flip, either does not increase all minimum path lengths from a specific vertex, or it does not decrease all minimum path lengths. That is to say, given a vertex, there cannot exist a vertex that gets closer after flipping and a vertex that gets further away. **Proof:** Suppose A is our chosen vertex, B is a vertex which gets further away after flipping, and C is a vertex which gets closer after flipping. Therefore the pre-flip minimum path from A to B goes through the edge that will be flipped, and there are no other paths from A to B of the same length. In particular, if n is the distance between D and A, then $n+1$ must be the path distance between E, F, and A, else flipping would not increase the distance. However, this means that no post-flip minimum path from A to C can go through the edge connecting E to F as they are both the same path length from A. This means that the flip does not affect the path length from A to C, a contradiction which proves the claim.

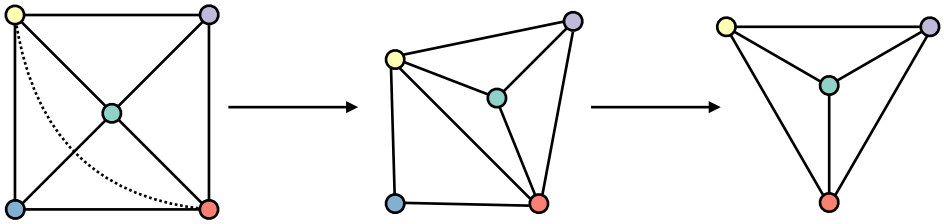

FIG. S5. Procedure for calculating the flip graph. (a) A local network of radius $r = 1$, with a potential flip identified (dotted line). (b) The flip is performed, but the resulting graph is not a local network. However, all path lengths from the central vertex have not decreased, meaning the new local network of radius 1 is a subgraph of this graph. (c) Taking the local network of radius 1 around the central vertex (in green), gives a local network 1 flip away from the original network.

(see Fig. S5). Since this requires calculating a Weinberg vector for potentially each edge of a network, the total cost is $O(NE^2)$.

The flip graph is connected, but in practice this calculation often yields a disconnected graph. This occurs when the path between two states goes through states that were not observed in the N observed networks (but do theoretically exist). If this occurs for a few isolated states that make up a negligible proportion of the total, then the largest connected component can be taken and these isolated states can be ignored. If the number of disconnected states is large, then the N observed networks can be augmented by additional networks observed in a Poisson-Voronoi process or similar [S5], and a larger flip graph can be calculated. Networks are added until the connected component of the new flip graph contains all (or almost all) of the original N observed networks.

Algorithm 4 Calculating the flip graph for N networks

Input: N topologically distinct local networks each with central vertex known.
for 1:N networks **do**
  for edges in network **do**
    Check if edge can be feasibly flipped [S6]. If it can check to see if the distances to the central vertex do not decrease (From claim in Fig. S4 finding at least one distance that increases or decreases suffices, and this can be done in the neighborhood of the flip).
    If the distances do not decrease, find the new local network around the central vertex, which will be a subgraph of the flipped network.
    Connect the vertices corresponding to the original local network and the new local network in the flip graph.
  end for
end for

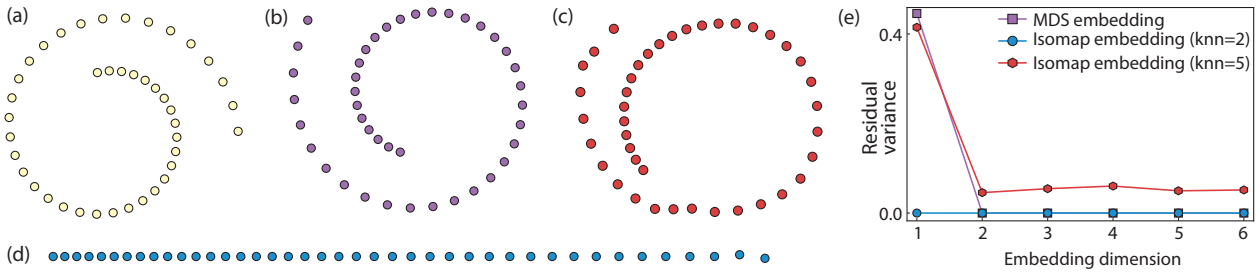

FIG. S6. Using Isomap to calculate residual variances recovers the correct manifold dimension, but an MDS embedding preserves the structure. (a) Points from a 2D spiral lie on a 1D manifold but have a non-trivial embedding in 2D space. (b) 2D MDS embedding (with Euclidean distances) of the data from (a) exactly preserves the structure. (c) 2D Isomap embedding (knn = 5) mostly preserves the structure. (d) 2D Isomap embedding (knn = 2) is effectively 1D. (e) The MDS embedding has no residual variance for dimension ≥ 2 suggesting that the data is 2D, whereas the Isomap embedding with knn = 2 has almost no residual variance for a 1D embedding, correctly identifying that the data lies on a 1D manifold. Isomap with knn = 5 suggests the data lies on a 2D manifold.

Calculating the TEM distance

The topological earth mover's (TEM) distance is defined to be the earth mover's distance between two probability distributions on the flip graph,

$$\mathrm{TEM}(A, B) = \min_{\gamma} \sum_{i,j} \gamma_{ij}\, d(i, j), \tag{S1a}$$

where the sum is taken over all pairs of networks i, j, $d(i, j)$ is the minimum path length on the flip graph between networks i and j, and γ is a map between distributions satisfying

$$\gamma_{ij} \geq 0, \qquad \sum_j \gamma_{ij} = P_A(i), \qquad \sum_i \gamma_{ij} = P_B(j). \tag{S1b}$$

Rather than optimize over all possible maps γ, the problem can be rephrased as a minimum cost flow problem over the flip graph [S8]. This is done by first converting to a minimum cost flow problem over the complete graph on N vertices, where the weight, or cost, on the edge between i and j is $d(i, j)$, and each vertex is a source or sink with strength $P_A(i) - P_B(i)$. This is then equivalent to solving on the flip graph, because the cost of sending mass directly from i to j is the same cost as sending it through the minimum length path between i and j.

The minimum cost flow is converted into the standard formulation, by taking two additional nodes, one a source, one a sink, and connecting every existing source to the new source by an edge with a capacity of the existing source strength and connecting every existing sink to the new sink by an edge with demand of the existing sink strength. With the exeption of the new source and sink, all other sources and sinks are then set to strength 0. To reduce the stiffness of the problem, each capacity was multiplied by $1+10^{-5}$ and each demand was multiplied by $1-10^{-5}$. Under this relaxation of the problem, the algorithm always converged, and the additional error was found to be $\sim 0.005\%$. For significantly larger or stiffer problems, approximations to optimal transport, such as entropic regularization are possible [S8], but were not required here.

Residual variance

The residual variance is defined as $1 - R^2(\hat{D}, D_U)$, where $\hat{D}$ is the (Euclidean) distance matrix in the embedded space, D_U is an unembedded distance matrix, and R^2 is the linear correlation coefficient [S9]. To calculate the residual variance, we do not take the TEM distance matrix, instead we take D_U as the Isomap distance matrix derived from the TEM distance matrix [S9]. In short, this means replacing the TEM distance between two points with their distance along the manifold; local distances are preserved and global distances become a sum of local distances along the path between two points [S9]. Local here means the k nearest neighbors of a point (knn), where k is a parameter to be chosen. The dimension of the manifold on which the points lie is the dimension for which the residual variance becomes negligible or does not decrease for higher dimensional embeddings [S9].

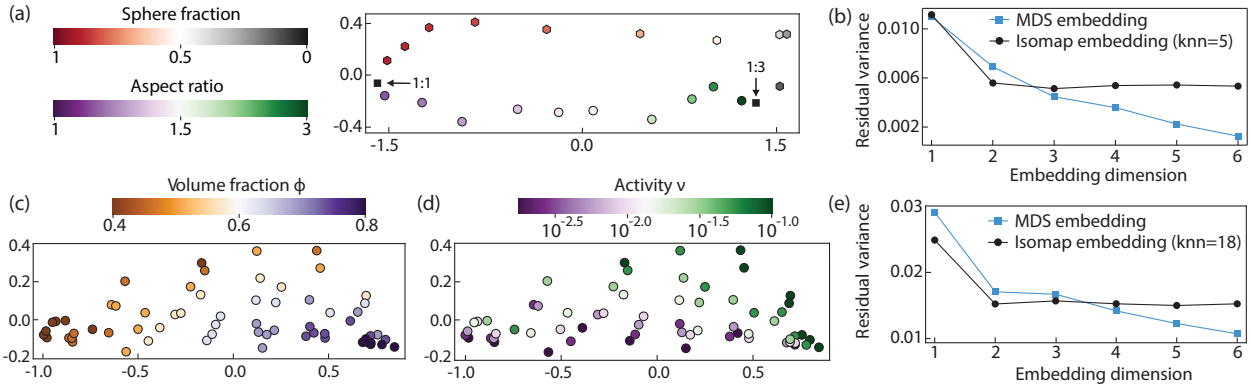

FIG. S7. Calculating the residual variance with Isomap finds the dimension of the phase space, with the Isomap embedding comparable to the MDS embedding. (a) Isomap embedding of the ellipsoid packings with knn = 5 recovers the two distinct paths from 1:1 to 1:3 ellipsoid packings even more clearly than the MDS embedding (Fig. 2 main text). (b) The residual variance from the Isomap embedding plateaus after two dimensions indicating a 2D phase space, whereas the MDS residual variance does not clearly indicate any dimension. (c-d) Isomap embedding of the ABP simulations with knn = 18 colored by volume fraction (c), and activity (d). Both volume fraction and activity are recovered as coordinates, but do not correlate as strongly with the principal components of the embedding as the MDS embedding (Fig. 3 main text). (e) Both in the MDS and Isomap embedding there is somewhat of a plateau after two dimension, but only Isomap definitively shows that the underlying space is 2D.

To understand why Isomap produces the most accurate estimate of dimension, consider the example of a 2D spiral; a 1D manifold with non-trivial embedding in 2D Euclidean space (Fig. S6a). The 2D MDS embedding maintains the spiral (Fig. S6b), as does a 2D Isomap embedding with a large number of neighbors (Fig. S6c). In contrast, a 2D Isomap embedding using 2 neighbors "unrolls" the shape to get a straight line (Fig. S6d). Therefore, Isomap can correctly identify the manifold as 1D, whereas MDS incorrectly identifies the manifold as 2D (Fig. S6e). That said, the MDS embedding preserves the non-trivial structure of the manifold in 2D; for this reason we stick with MDS embeddings for visualization purposes.

For the ellipsoid packing and ABP examples, using the Isomap distance matrix to calculate the residual variance correctly recovers the dimension of the subspaces, unlike MDS (Fig. S7b,e). The number of neighbors used was chosen to be large enough to make the Isomap embedding consistent with the MDS embedding, whilst being small enough to "unroll" the manifold (Fig. S7a,c,d).

Active Brownian particles

Simulations of 2D active Brownian particles (ABPs) were performed with 2,000 particles following the method described in [S10], which we briefly outline here. Periodic boundary conditions were used for a box of size $L \times L$. Let $\boldsymbol{r}_i$ and θ_i describe the center position and the orientation of the i-th particle respectively. The over-damped dynamics of each particle is governed by the following equations,

$$\frac{\mathrm{d}\boldsymbol{r}_i}{\mathrm{d}t} = v\hat{\boldsymbol{n}}_i + \mu \sum_{j \neq i} \boldsymbol{F}_{ij}, \tag{S2}$$

$$\frac{\mathrm{d}\theta_i}{\mathrm{d}t} = \eta_i(t), \tag{S3}$$

where $\hat{\boldsymbol{n}}_i = [\cos\theta_i, \sin\theta_i]$ describes the orientation of the i-th particle, v is the self-propulsion speed, and μ is the mobility. $\boldsymbol{F}_{ij}$ is a pairwise soft repulsive force such that $\boldsymbol{F}_{ij} = \boldsymbol{0}$ when the particles i and j are not overlapping, and $\boldsymbol{F}_{ij} = k(a_i + a_j - r_{ij})\hat{\boldsymbol{r}}_{ij}$ with $r_{ij} = ||\boldsymbol{r}_i - \boldsymbol{r}_j||$ and $\hat{\boldsymbol{r}}_{ij} = (\boldsymbol{r}_i - \boldsymbol{r}_j)/r_{ij}$ when the particles overlap. Eq. (S3) is a stochastic differential equation with Gaussian white noise, $\eta_i(t)$, satisfying $\langle \eta_i(t)\eta_j(t')\rangle = 2v_r\delta_{ij}\delta(t-t')$, where v_r is the the rotational diffusion rate. The radius of i-th particle a_i is drawn from a uniform distribution between $0.8a$ and $1.2a$. The domain size L is computed from the volume fraction ϕ and the radii of the particles $L = \sqrt{\sum_i \pi a_i^2/\phi}$. Choosing mean particle radius a as the unit of length and the elastic time scale $\tau = (\mu k)^{-1}$ as the unit of time, the parameter space is reduced to the effective self-propulsion speed $\tilde{v} = v/(a\mu k)$, the packing fraction ϕ and the effective

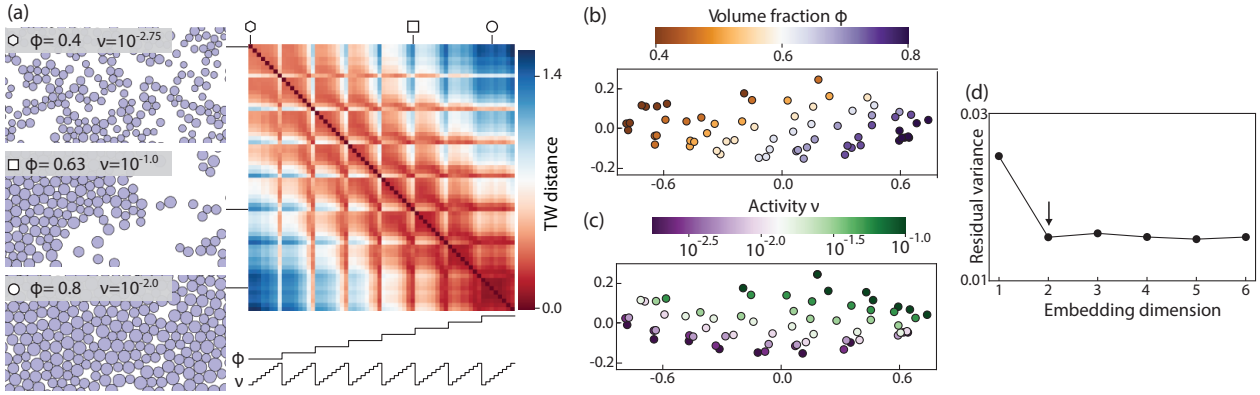

FIG. S8. Recovery of phase space dimension and coordinates for active Brownian particles. (a) Snapshots of simulations for different parameter values showing liquid like, phase separated, and glass like states, and their corresponding entries in the distance matrix. For every simulated parameter set, 5 runs with 2000 particles were combined to obtain an average distribution. (b-d) 2D MDS embedding, (b-c) colored according to volume fraction and activity respectively, showing the original phase space is recovered. (d) The residual variance plateaus after embedding dimension 2, correctly identifying the true embedding dimension of the phase space (shown with arrow).

rotational diffusion $\tilde{v}_r = v_r/(\mu k)$. For all simulations, we fix $\tilde{v}_r = 5 \times 10^{-4}$ and simulate values in $\tilde{v} \in [1.8 \times 10^{-3}, 0.1]$ and $\phi \in [0.4, 0.8]$.

A custom, parallelized code employing graphics processing units (GPUs) was implemented to perform the simulations, following [S10]. We use a standard explicit Euler scheme to numerically integrate the dimensionless form of Eq. (S2) and (S3) from $t = 0$ to $t = 5000\tau$ with a time step $\Delta t = 0.01\tau$. Only the snapshot of $t = 5000\tau$ is used to calculate the TEM distances, no dynamic information is used.

Following [S10], we neglect translational noise, although activity becomes equivalent to translational noise in the limit where the orientational correlation time becomes much smaller than the mean free time between collisions. For a given volume fraction, this limit will be realized as $v \to 0$, so reducing v makes the system closer to a thermal system [S10]. That said, phase separation for active particle systems is a distinctly non-equilibrium phenomenon, making much of our simulated phase space far from equilibrium [S10].

Three representative partial simulation snapshots can be seen in Fig. S8(a). We computed the Delaunay tessellations of the ABP-centroid positions from each snapshot and then the associated $r = 2$ motif distributions for each simulation (ϕ, v). The TEM distance matrix for all simulation pairs is depicted in Fig. S8(a). Notably, the first two principal components of the associated MDS embedding recover the phase space spanned by volume fraction and activity parameter (Fig. S8b,c). In particular, the second principal component correlates closely with activity, demonstrating that the TEM metric detects the transition to far-from-equilibrium dynamics (large v), which is recovered by the embedding without need for time-resolved data (Fig. S8c).

Comparison with Jensen-Shannon

The Wasserstein or earth mover's distance is only one of many possible metrics that could be taken between two distributions. Another possible distance is the Jensen-Shannon (JS) distance which has also been used to distinguish cellular structures [S11]. It is defined by

$$JS(A,B)^2 = \frac{1}{\ln 2} \sum_i \frac{1}{2} P_i^A \log \left[\frac{P_i^A}{\frac{1}{2}(P_i^A + P_i^B)} \right] + \frac{1}{2} P_i^B \log \left[\frac{P_i^B}{\frac{1}{2}(P_i^A + P_i^B)} \right], \tag{S4}$$

where the sum is taken over all networks i, and P_i^A is the probability of observing network i in distribution A, similarly for P_i^B. The JS distance is an entropic distance between distributions, based on the idea of mutual information. It does not use any notion of distance between networks, only using their isomorphism classification. This has the drawback that while it can distinguish distinct distributions, it cannot tell to what degree they are different, for example all non-overlapping distributions are JS distance 1 away from eachother regardless of their particular forms [S12].

For monodisperse packings with varying aspect ratio (Fig. S9a), both the TEM distance, and the JS distance are consistent across different simulations with the same parameters (Fig. S9b,c). Unlike the JS distance, the TEM

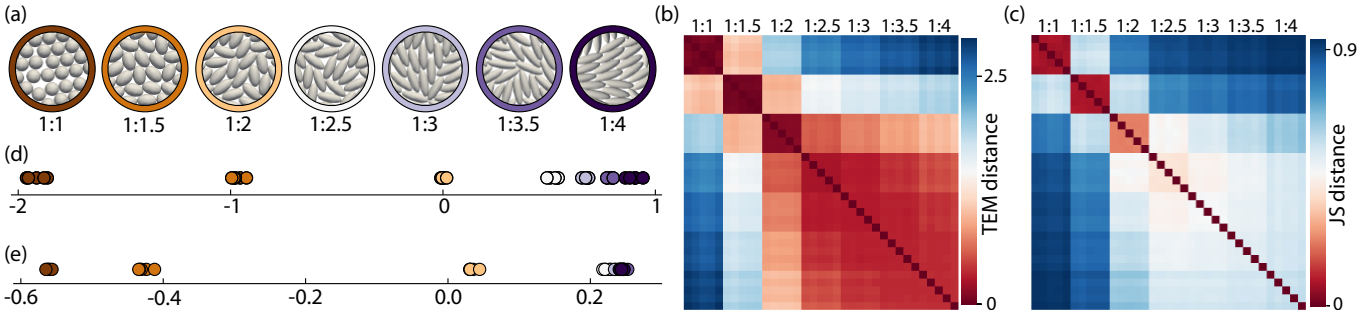

FIG. S9. Comparison of TEM and JS distances ability to distinguish simulations of jammed disordered packings of ellipsoids for different aspect ratios. (a) Snapshots of ellipsoid packings for various aspect ratios. (b-c) Distance matrices for the TEM and JS distances respectively. Each pixel represents the distance between two simulations, with five simulations of 10,000 ellipsoids performed for each aspect ratio. (d-e) 1D embedding of the simulations using MDS for TEM and JS respectively. Color coding is according aspect ratio as in (a).

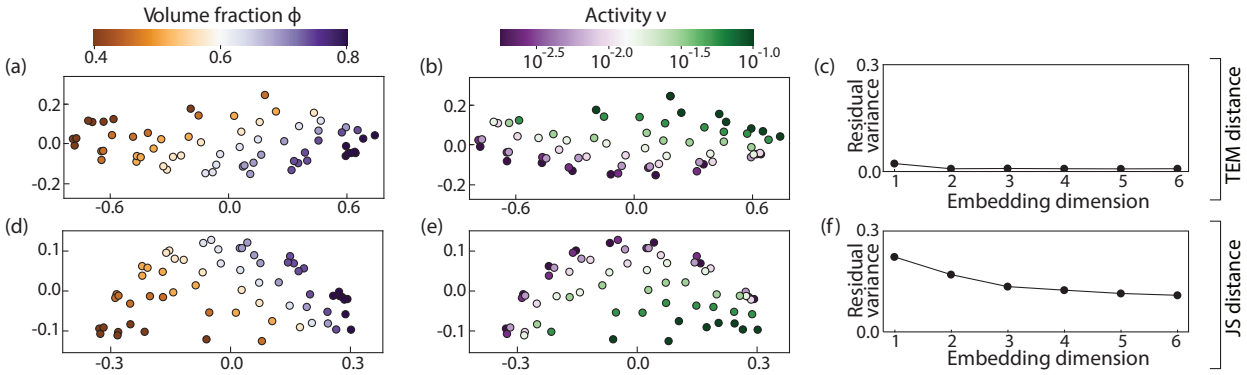

FIG. S10. While both the JS and TEM distances recover the underlying phase space in a 2D MDS embedding of ABP simulations, only the TEM distance matrix can be comfortably embedded in a low dimensional space. (a-b) 2D MDS embedding of ABP simulations using the TEM distance, colored according to volume fraction and activity respectively. (c) over 98% of the variance is recovered in a 2D MDS embedding of the TEM distance matrix. (d-e) 2D MDS embedding of ABP simulations with the JS distance. (f) even a 6D MDS embedding only accounts for around 90% of the variance of the JS matrix, with no clear preference for an embedding dimension.

distance recognizes different simulations of 1:3-4 ellipsoids as very similar to each other. The distance matrix was embedded in a 1D space, the true dimension of the manifold on which the data lies, using MDS. For the TEM distance this recovers the correct ordering of aspect ratios, whereas the JS distance is unable to separate some of the larger aspect ratio packings, Fig. S9d,e.

For dense enough phase spaces, both the JS and TEM distances should be able to predict which points are neighbors, even if the JS distance cannot tell how far distant points are. For this reason it is not surprising that both distances, when embedded in 2D using MDS, recover the phase space for the ABP simulations (Fig. S10a,b,d,e). However, while the residual variance for the TEM distance clearly indicates that the phase space is 2D (Fig. S10c), and a 2D embedding recovers over 98% of the variance, the JS distance shows no clear preference for any dimension, and even a 6D embedding recovers only 90% of the variance (Fig. S10f).

The inability of the JS distance to tell how closely related networks are is a significant drawback when limited data is available. In this regime, the true distribution is not sampled well; while this would reduce the accuracy for all distances, a distance that has no concept of similarity between networks will particularly struggle. To see this, imagine the extreme case where for each distribution we take only the one network that occurs most frequently. The TEM distance may still provide information about the distributions, but the JS distance will give all distances as 0 or 1. To test this intuition, we took jammed disordered packings of 10,000 ellipsoids for aspect ratios 1:1-2.5 (Fig. S11a), and subsampled them, retaining either 1000, 100, or 20 local networks. The TEM and JS distances were calculated (Fig. S11b,d), as were their MDS embeddings, (Fig. S11c,e). The correct MDS embedding is recovered for the TEM distance even when only 20 local networks from each simulation are available, and this is consistent across simulations. For the JS distance, the MDS embedding is incorrect even for 100 local networks. While the correct

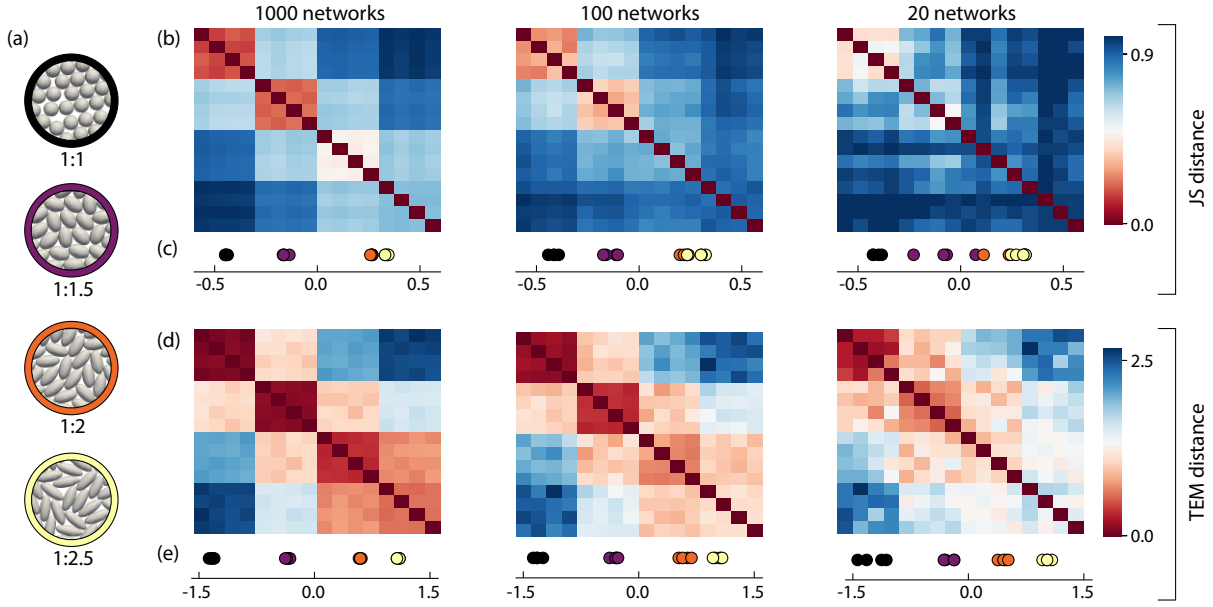

FIG. S11. TEM distance can correctly distinguish ellipsoid packings from as little as 20 local networks, unlike the JS distance. From a large jammed disorded ellipsoid packing simulation, a small number (1000, 100, and 20) of local networks were sampled. This was done for aspect ratios 1:1-2.5 shown in (a), with 5 simulations for each aspect ratio. (b) Distance matrix for the JS distance. Taking 1000 networks it is apparent that there are 4 physically distinct materials, whereas for 20 networks this is not apparent. (c) 1D MDS embedding of the distance matrix in (b) with color coding according to (a). The correct MDS embedding is only recovered when 1000 networks are sampled. (d) Distance matrix for the TEM distance, even for only 20 sampled networks, the 4 distinct parameter regimes are visible. (e) 1D MDS embedding of the distance matrix in (d) with color coding according to (a). The correct MDS embedding is recovered even when only 20 local networks are sampled.

ordering is recovered for 1000 local networks, there is hardly any separation between the 1:2 and 1:2.5 ellipsoids in the MDS embedding, unlike the 4 clear clusters that are apparent for all TEM MDS embeddings. Therefore, in the case when limited data is available, the TEM distance outperforms the JS distance.

Choice of radius r

To capture the local ordering around a point we take the local network of radius r. The larger r is, the more local information we capture, but also many more distinct networks are observed with corresponding increase in computational cost. Taking the ABP simulations as an example, 320 simulations of 2000 particles were performed, so 640,000 networks were computed in total. For $r = 1$, $O(20)$ distinct networks were observed, for $r = 2$, $O(30,000)$ were observed and for $r = 3$, $O(500,000)$ were observed. This means that computing optimal transport exactly for $r \geq 3$ becomes computationally infeasible, but could be solved approximately using entropic regularization [S8]. However, since $r = 2$ contains sufficient information to recover the phase space, the question becomes whether $r = 1$ may be sufficient as well. Although taking $r = 1$ works to some extent, it is not sufficient to recover the 2D phase space for the ABP simulations, Fig. S12. Using the MDS embedding for $r = 1$, over 99.6% of the variance is in the first principal component (compared to 97.5% for $r = 2$), suggesting, erroneously, that the manifold is 1D, despite the true phase space lying on a 2D manifold.

To understand why taking $r = 1$ is insufficient, we calculated the flip graph for all ABP simulations, and calculated the frequency at which each local network was observed (Fig. S13). Neglecting a negligible fraction of the total networks observed ($< 0.1\%$), the flip graph is simply a 1D, or path, graph and the networks tell us only how many neighbors each Voronoi cell has. Euler's theorem tells us that the average number of neighbors will be 6, and as demonstrated in Fig. S13, the number of neighbors remains close to 6. Therefore, the distribution is approximately 1D; if the fraction of 5 sided shapes is p_5, then the fraction of 7 sided shapes is $p_7 \approx p_5$, and the fraction of 6 sided shapes is $p_6 \approx 1 - 2p_5$, meaning the whole distribution is approximately described by a single parameter. While in reality there are further degrees of freedom, the fact that the distribution is almost 1D explains why $r = 1$ will struggle to reconstruct a 2D or higher phase space.

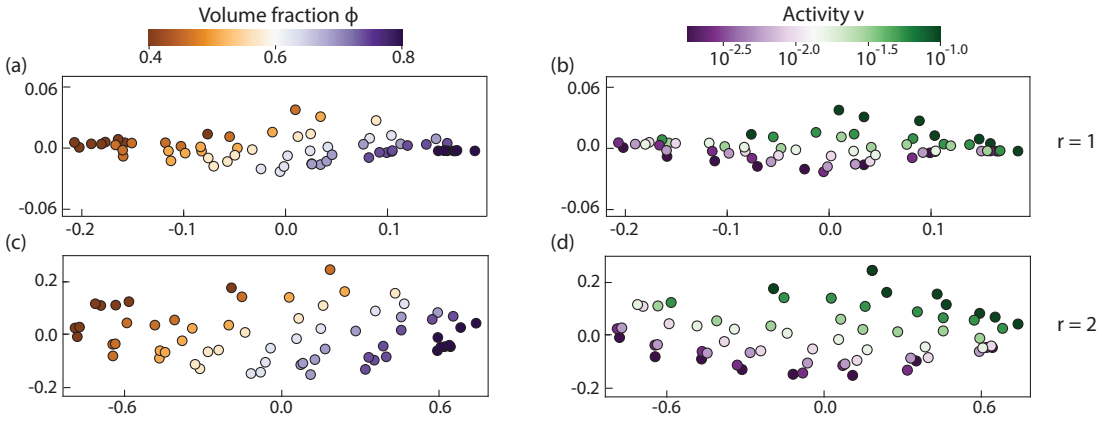

FIG. S12. Taking the local network of radius $r = 2$ is sufficient to capture the 2D nature of the phase space, whereas $r = 1$ is insufficient. (a-d) MDS embedding of ABP simulations for various values of volume fraction and activity. (a-b) Taking local networks of radius $r = 1$ recovers the volume fraction as a principal component, (a), but fails to recover the activity, except for intermediate values of the volume fraction, (b). (c-d) Taking local networks of radius $r = 2$ recovers a 2D phase space with one principal component corresponding roughly to volume fraction, (c), and the other to activity, (d).

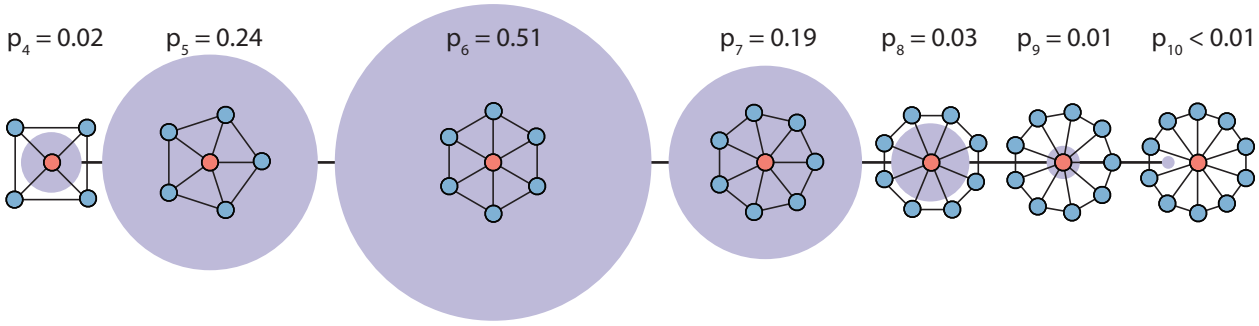

FIG. S13. The flip graph for $r = 1$ is essentially 1D and contains only the number of sides of each Voronoi cell. All ABP simulations were combined and the flip graph calculated for local networks of radius $r = 1$. Each node of the flip graph is drawn in purple with the area of the node representing the proportion of local networks found in this state. The nodes that were not drawn represent less than 0.1% of the total neworks seen. Each node has the local network that it represents overlayed.

Bacterial swarm

The bacterial strain used for the swarming assay in this study was a $\Delta comI$ derivative of the ancestral *B. Subtilis* strain NCBI3610, obtained from Daniel Kearns' lab at Indiana University. Cell cultures were grown at 37°C in Luria-Bertani (LB) liquid medium for 18h before incoluation. Swarming plates were prepared by filling $12ml$ of LB medium containing 0.5% bacto agar into a 90mm petri dish prior to a drying period of 20 minutes. The plates were then immediately inoculated with a small drop of bacterial culture of volume $< 100nl$, which was placed in the center of the plate. Swarms were kept at 37°C in an enclosure containing a water reservoir and imaged for 12h using the automated microscopy routine and setup described in Ref. [S13]. At each space-time location an image of 1024×1024 pixels is taken corresponding to an area of $170\mu\text{m} \times 170\mu\text{m}$.

To segment the cell centroids, we used the cell segmentation tool Stardist together with a U-Net backbone [S14, S15]. For training data, we anually annotated 30 images containing 26,228 cells in total. The following table contains the hyperparameters, if not otherwise stated the default parameters were used.

parameter	value
backbone	U-Net
n_rays	32
unet_n_depth	3
batch_size	4
patch_size	(256,256)

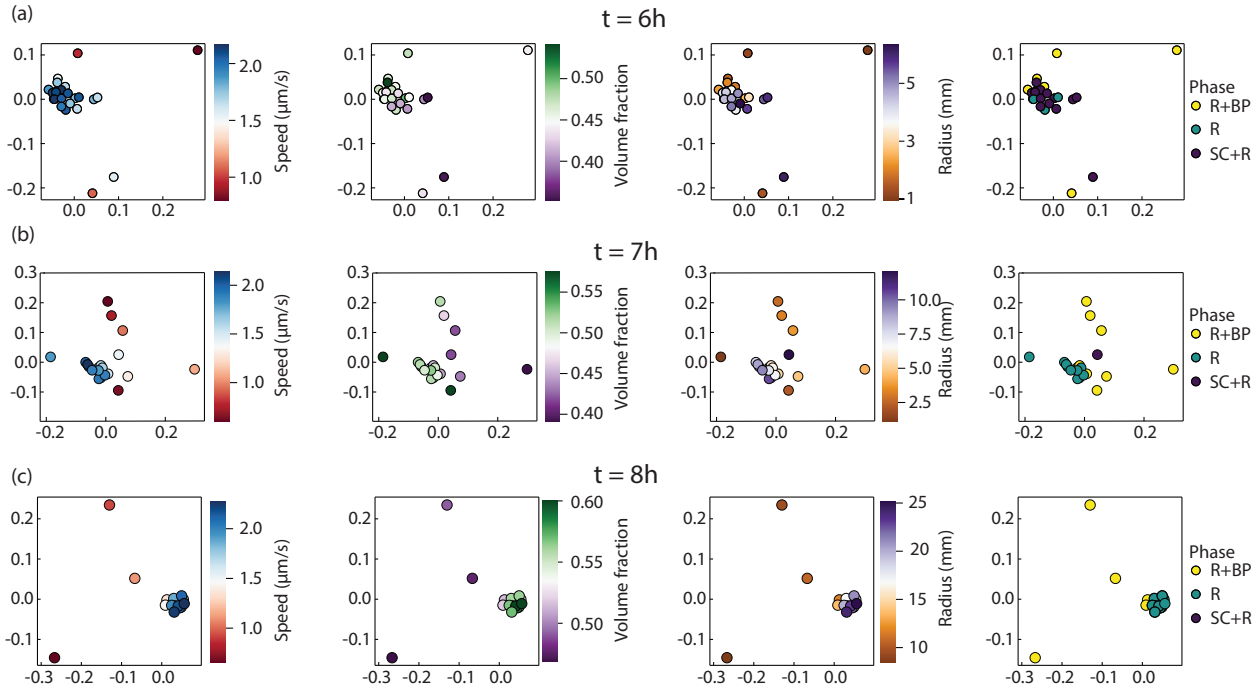

FIG. S14. MDS embedding of spatial snapshots at a fixed time find hereogeneous spatial structure. For times $t = 6, 7, 8$h, all spatial snapshots were taken, pairwise distance matrix computed, and a MDS embedding was performed. The embeddings are colored by speed, volume fraction, radius and phase. (a) At early times ($t = 6$h), many snapshots contain fast moving cells and this region of the embedding approximately recovers the spatial ordering. The remaining snapshots contain slower moving cells and are either at the core of the swarm in the biofilm precursor region, or at the edge. (b) At mid times ($t = 7$h), the spatial ordering is not recovered, but either speed or volume fraction does a reasonable job at explaining the first principal component. The phases are separated in this embedding, except some R+BP points which are close to the R cluster. (c) At late times ($t = 8$h), either speed, volume fraction, or radius could explain the first principal component. The phases are also separated.

There are 320 experimental snapshots across time and space which, due to the quadratic growth of computing the distance matrix, makes the full MDS embedding expensive (at least 16× the cost of the ABP embedding). The embedding by motif size parameters has no such quadratic cost, and was discussed in the main text. Here we consider instead embedding all snapshots for a fixed time, and hence investigate the spatial structure at a particular time. We performed this embedding for three time points, each containing 15-20 snapshots, Fig. S14. We first note that the swarm is heterogeneous, and the snapshots are separated by $O(500\mu\text{m})$ or hundreds of bacteria lengths, meaning we do not necessarily expect the properties to change smoothly with radius. While there is often a relationship between the embedding and the radius, the primary principal component does not always correspond to the radius (except for $t = 8$h), Fig. S14. Volume fraction and speed are also strongly linked to the embedding and can explain certain principal components. Finally, we find that the embeddings can cluster phases found by a data-driven clustering of snapshot parameters [S13]. For the region of the swarm we are analyzing, there are three identified phases: rafting + biofilm precursor (R+BP), rafting (R), and single cell + rafting (SC+R) [S13]. We find that the phases tend to separate in the embedding, Fig. S14. In addition to phase classification, we can now use the MDS embedding to see if a snapshot is truly representative of its phase or if it is close to transitioning into a different phase.

Geodesics in W_1 space

The energetic barrier required to perform a flip in epithelial cell layer suggests that the system might evolve in such a way to minimize the total number of these flips required. To investigate, we seek to find a path between the average start and end experimental distributions that exactly solves this minimization problem. Recall to find the earth mover's or 1^{st} Wasserstein (W_1) distance, we take the start and end distributions, say P_0, and P_1, and solve the for the optimal transport plan moving mass γ_{ij} from i to j. A natural way to interpolate this process would be to

instead move mass $t\gamma_{ij}$, from i to j, for $t \in [0,1]$, which results in the intermediate distributions $P_t = (1-t)P_0 + tP_1$. This path is a geodesic in the sense that $\sum_{t_k} d(P_{t_k}, P_{t_{k+1}}) = d(P_0, P_1)$, for $0 = t_0 \leq \cdots \leq t_n = 1$, but P_t is not the only interpolating distribution to have this property, as geodesics are not unique under the W_1 metric [S16]. This distribution could be interpreted as phase separated growth, a fraction $(1-t)$ of the system is in the initial distribution, the remainder in the final, and the size of the phase in the final distribution grows until the whole system is in that distribution. However, this has a certain unphysical aspect of mass instantaneously appearing in the final state as the path is traversed [S16]. In contrast, it is well known that when computing the second Wasserstein (W_2) distance on a continuous manifold $\mathcal{M}$, one finds a naturally interpolating path between initial and final distributions. Moreover, the distance can be computed as

$$\begin{aligned} W_2^2(\rho_0, \rho_1) = \min_{\rho, v} \int_{t=0}^{t=1} \int_{\mathcal{M}} \rho(x,t) ||v(x,t)||^2 \, \mathrm{d}x \, \mathrm{d}t \\ s.t.\ \rho(x,0) = \rho_0,\ \rho(x,1) = \rho_1, \text{ and} \\ \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0. \end{aligned} \tag{S5}$$

for a probability density ρ over $\mathcal{M}$ advected by a velocity field, v [S16]. The term in the integrand to be minimized is the natural analogue of dissipation for this transport problem, and the intermediate path $\rho(x,t)$ is the natural interpolation between start and end distributions [S16].

In order to compute a W_1 geodesic on a graph, we must find a transport plan $\gamma_{ij}(t)$, that minimizes

$$W_1 = \sum_{ij} \int_t \gamma_{ij}(t) \mathrm{d}t \tag{S6}$$

together with the conservation equation

$$\frac{d}{dt} P_t(i) = \sum_j \gamma_{ji}(t) - \sum_j \gamma_{ij}(t), \tag{S7}$$

with initial and final conditions $P_{t=0} = P_0$ and $P_{t=1} = P_1$; $\gamma_{ij}(t)\mathrm{d}t$ is the mass transported from i to j in time $\mathrm{d}t$. As previously discussed, there are many such $\gamma_{ij}(t)$ that satisfy this equation. Here, we take all transport plans $\gamma_{ij}(t)$ that minimize the W_1 cost, and out of these, find the plan that minimizes the transport dissipation. To write down the graph-based equivalent of the transport dissipation, we can interpret the mass transport $\gamma_{ij}(t) = (\text{density}) \times (\text{velocity})$, so that dissipation would become,

$$R = \sum_{i,j} \int_t \frac{\gamma_{ij}(t)^2}{P_t(i,j)} \mathrm{d}t, \tag{S8}$$

where the sum is assumed to be taken over non-zero γ_{ij}, and $P_t(i,j)$ denotes the probability density over edge (i,j). To define an edge based density (previously density was defined on vertices), we follow [S16], taking the harmonic mean of the densities at the corresponding vertices, $P_t(i,j) = 2P_t(i)P_t(j)/(P_t(i) + P_t(j))$. To solve numerically, we discretize in time, finding $T+1$ states, $P^{(0)}, P^{(1)}, \ldots, P^{(T)}$, and T transport plans $\gamma^{(1)}, \ldots \gamma^{(T)}$, making the transport dissipation [S16],

$$R = T \sum_{l=1}^{l=T} \sum_{i,j} (\gamma_{ij}^{(l)})^2 \left(\frac{P^{(l)}(j) + P^{(l-1)}(i)}{2P^{(l)}(j)P^{(l-1)}(i)} \right), \tag{S9}$$

where the sum over i,j is assumed to only include $\gamma_{ij}^{(l)} > 0$ terms so that $P^{(l)}(j)$, $P^{(l-1)}(i)$ are always non-zero when they appear in the sum. Note that the advective nature of transport is reflected in the harmonic mean by taking the density at the source of the transport at time $l-1$ and the density at the destination of transport at time l. The full discretized problem for finding the dissipation regularized geodesic can be stated as

$$\begin{aligned} \min_{P^{(l)}, \gamma^{(l)}} R(P^{(l)}, \gamma^{(l)}) \ \ s.t.\ \ P^{(0)} = P_0, P^{(T)} = P_1, \\ \sum_j \gamma_{ji}^{(l)} - \sum_j \gamma_{ij}^{(l)} = P^{(l)}(i) - P^{(l-1)}(i), \\ \sum_{l=1}^{T} \sum_{ij} \gamma_{ij}^{(l)} = W_1(P_0, P_1), \end{aligned} \tag{S10}$$

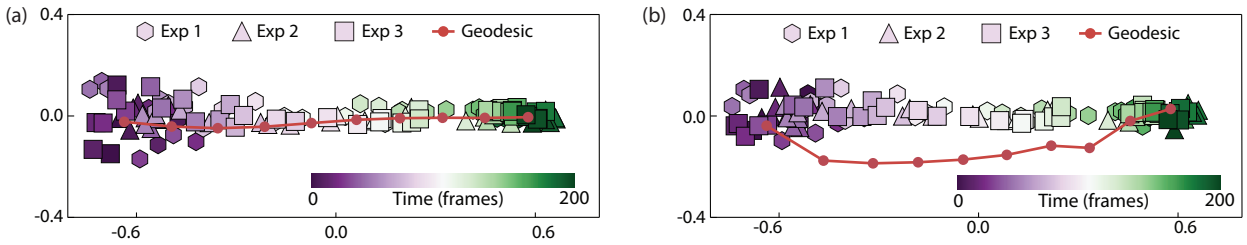

FIG. S15. Comparison of dissipation regularized geodesic and phase separated growth geodesic. (a) MDS embedding of data and dissipation regularized geodesic (also Fig. 3h of main text). (b) MDS embedding of data and the phase separated growth geodesic.

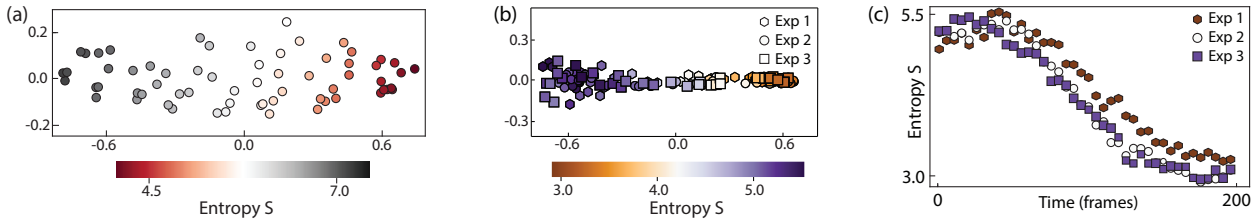

FIG. S16. Topological entropy as an analysis tool for embeddings. (a) MDS embedding of ABP data colored by topological entropy shows that the topological entropy corresponds to the first principal component, and increases with decreasing density as more states become available. (b) MDS embedding of *Drosophila* data colored by topological entropy. Topological entropy corresponds to the first principal component. (c) Topological entropy decreases with time for the *Drosophila* embryos as the cell layer becomes more regular.

where the last constraint enforces that the path is still optimal with respect to the W_1 metric. The integer parameter T determines the level of discritization, with the continuous limit being recovered for $T \to \infty$ [S16]. This problem can be solved by second order conic optimization programs, and whilst significantly slower than solving the linear min-cost flow for W_1, it is still practical for the systems considered here.

We see that the dissipation regularized geodesics calculated by this method closely match the actual path the data takes, unlike the phase separated interpolation, Fig. S15. Fluctuations around the geodesic mean that the system will not exactly minimize dissipation, but nevertheless follows the minimizing path closely.

Topological entropy

In practice, if the underlying parameters are unknown, interpretation of the principal components directly from the data is desirable. In the next section, we introduce tools to analyze the motif size distribution, which yields further interpretations of the embedding. Here, we note that from the observed motif distributions $P(i)$, we can compute the topological Shannon entropy [S17], $S = -\sum_i P(i) \ln P(i)$. The interpretation of entropy is that it is a measure of how many states are available to the system. The entropy S often explains the first MDS principal component, for instance in the ABP embedding, Fig. S16. In the *Drosophila* example, the entropy corresponds to the first principal component, notably decreasing with time, Fig. S16, consistent with known results about the increasing regularity of the developing wing [S3].

Tracy-Widom distribution

We seek to understand the structure of the data by understanding the probability distribution of motifs. This space is hard to analyze directly, so we consider a projection onto $\mathbb{N}$, by taking the distribution of motif sizes, defining the size of a motif to be the number of vertices that it has. In the main text, we claimed that for the swarm data, the distributions differed in mean and variance, but all lay on a universal distribution, namely the Tracy-Widom (TW) distribution. Here we provide stronger numerical evidence that the distribution is indeed TW by analyzing the network formed by a Poisson point process, known as a Poisson-Voronoi (PV) tesselation [S18], which is equivalent

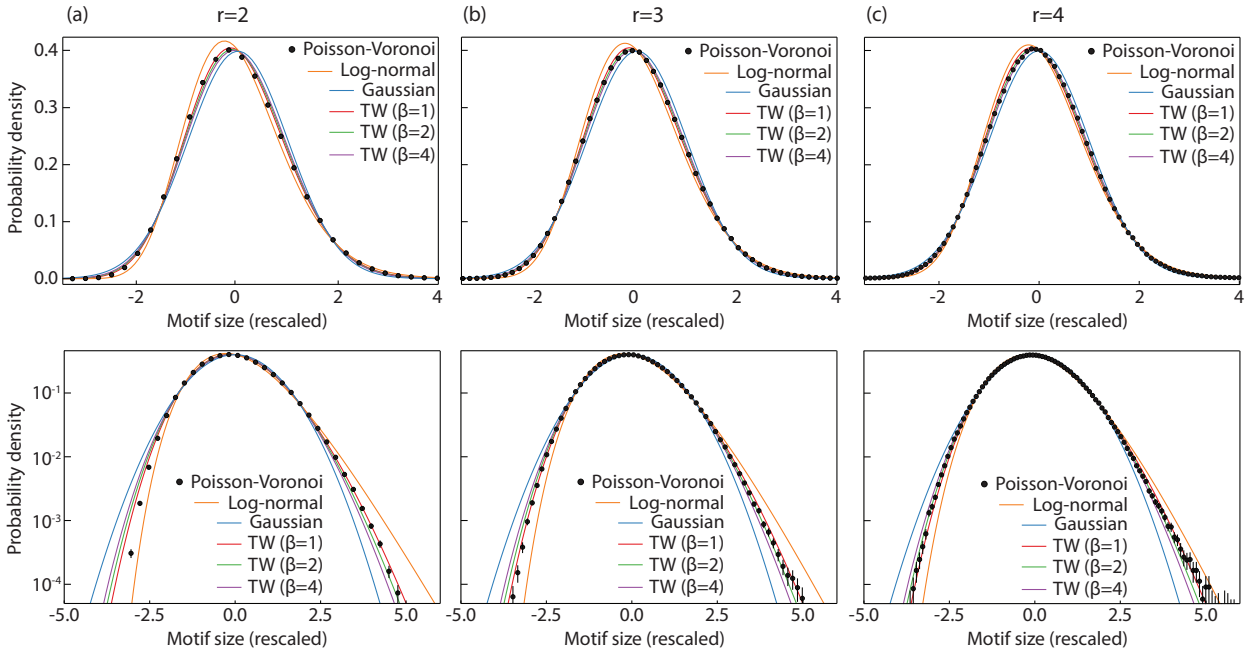

FIG. S17. Rescaled motif size distribution for various motif radii, compared against known distributions. (a) For $r = 2$, the distribution is close to TW, but differs in the left tail due to geometric constraints on the minimum motif sizes. (b) For $r = 3$, the distribution approaches TW on the left tail, but does not yet lie exactly on it. (c) For $r = 4$, the distribution closely matches TW in the regions shown. Error bars are ± 2 standard deviation, from sampling error.

to a non-interacting gas, taking 2×10^6 points in a periodic domain. To analyze the motif size distribution, we took the observed distribution for fixed r, and rescaled so that it has mean 0, variance 1. We compare this against various known distributions including a Gaussian, Log-normal, and TW for all β parameters, all rescaled to have unit variance and zero mean. The distribution were chosen for their universal characteristics, and hence plausible appearance in the motif size distribution. Gaussian distributions are universal due to the central limit theorem, and for positive variables a Log-normal distribution has universal properties [S19]. We discuss in the main text the universal character of TW distributions. The parameters for the Log-normal were chosen by maximum likelihood estimation before rescaling, the TW and Gaussian distributions were only rescaled, not fitted. We see for $r = 2$, that the distribution lies enough close to TW to be plausible, but differs slightly in the left tail and due to the discrete nature of motif sizes, does not finely resolve the distribution, Fig. S17a. However, by taking ever larger motif sizes, $r = 3, 4$, the observed distribution gets significantly closer to TW, matching in both the $O(1)$ regions and in the tails until at least probability density 10^{-4}, Fig. S17b-c. In particular, the distribution for larger motif sizes lies on TW, and not a Gaussian or Log-normal, even though Log-normal (for our parameters) and TW are very similar distributions. TW with $\beta = 2$ appears to be the best fitting of the TW distributions, although they only differ slightly when rescaled, so further numerical tests would be needed to confirm this. In the main text, and the rest of this document we use $\beta = 2$ for the TW distribution. We note that this distribution is universal in the sense that the motif size distributions for PV and for the swarm data do not share the same mean or variance, and the underlying processes are different. The mean $r = 2$ motif size of a PV process is 20.7 and variance 14.8, which is significantly greater than the corresponding quantities measured from the swarm snapshots. Further, PV can be thought of as a snapshot of an non-interacting ideal gas, whereas the bacteria in the swarm have strong excluded volume interactions and density plays an important role. Nevertheless, the same distribution is found.

Analysis of motif size distribution

Applying the motif size analysis to the ABP data, we find a range of possible distributions, Fig. S18d, not all of which appear to be TW. We can approximately characterize these distributions by taking their mean, variance, and third moment. We find that the mean and variance correspond to the first principal component of the MDS embedding, and the third moment corresponds to the second principal component, Fig. S18(a-c). In an alternative to

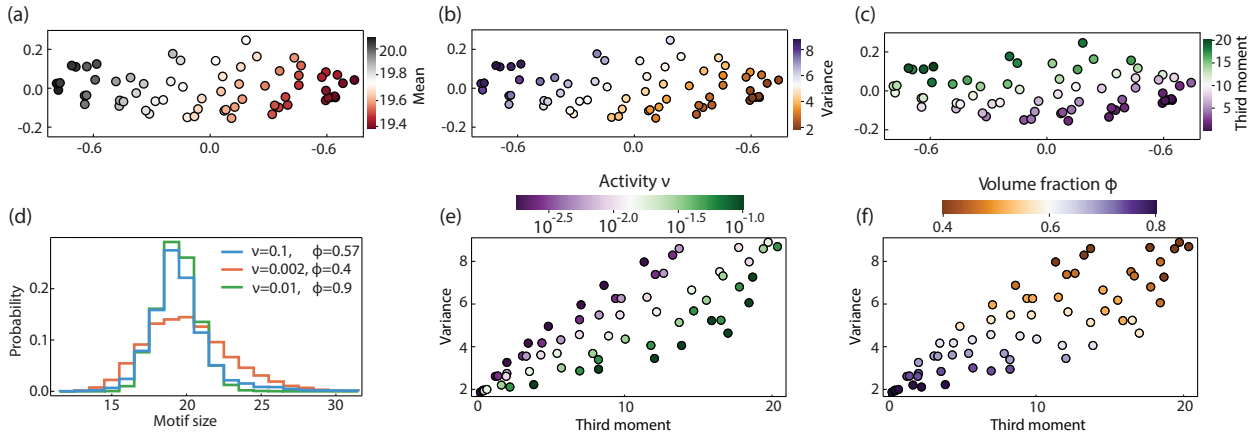

FIG. S18. Motif size distributions characterize ABP simulations. (a-c) Coloring the MDS embeddings by mean, variance and third moment respectively, finds that the mean and variance correspond to the first principal component, and the third moment corresponds to the second principal component. (d) Example motif size distributions for various values of volume fraction and activity. (e-f) Simulations embedded in the variance-third moment plane colored by activity and volume fraction respectively. The embedding, like the MDS embedding, recovers the parameter space.

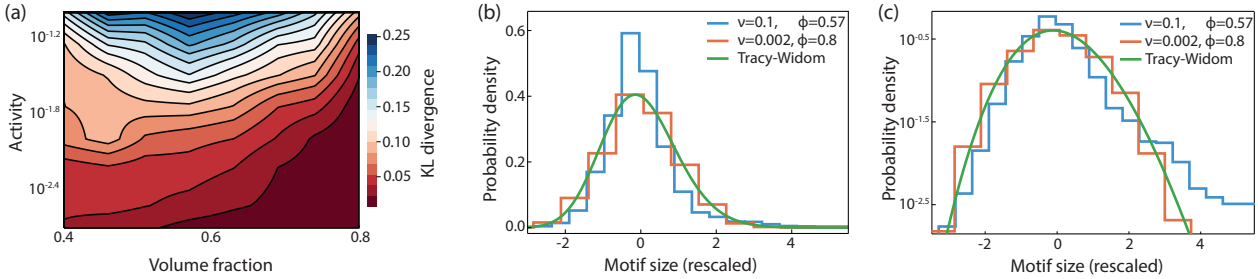

FIG. S19. Motif size distributions for ABP simulations are approximately Tracy-Widom in the liquid-like region of phase space. (a) KL divergence between rescaled motif size distributions and TW distributions shows regions of phase space where the TW approximation holds. (b) Representative example histograms of distributions which are far from TW, and distributions that are close to TW, together with the TW distribution. (c) Same as (b) with log-scale.

the TEM-MDS embedding, one can plot the simulations in the variance-third moment plane, finding that this again recovers the activity/volume fraction phase space, Fig. S18(e-f). This is not preferable to the metric embedding, since there is no notion of distance, although here one need not compute $O(N^2)$ components of a distance matrix, providing an alternative when N is large. To understand these embeddings, we see that reducing the volume fraction for fixed activity causes more states to become available as particles interact less. Hence the motif size variance increases and so does the third moment (as does the topological entropy). Interestingly, when the activity is increased, the variance does not change much, but the third moment increases significantly. In the most extreme case of $\phi = 0.8$, going from the lowest activity to the highest, the third moment increases by a factor of 16, whereas the variance increases by a factor of 1.2, and the mean does not chance. Hence activity changes the distribution by increasing the probability of large motif fluctuations, which directly correspond to the giant number fluctuations of the phase separated regime [S10].

The distributions seen in the ABP simulations appear to take a variety of forms, although the Tracy-Widom (TW) distribution appears, Fig. S19. We evaluate where in the parameter space the motif distribution appears most TW like, by comparing the distribution rescaled to zero mean and variance 1, and the TW distribution similarly scaled. We use the Kullback-Leibler (KL) divergence between two distributions, defined as $KL(P||Q) = \sum_n P(n)\log(P(n)/Q(n))$, representing the amount of information lost when the data distribution (P) is approximated by the TW distribution (Q). To compare continuous and discrete distributions in calculating the KL divergence, we consider the discritized TW pdf to be the mass lying within a bin as defined by the histogram of the discrete data. We find that in the most active regimes the observed distribution is furthest from TW. Perhaps this is because in the phase separated regimes, the dense phase is a strongly coupled glass like phase, with a narrow distribution of motif sizes, but the gas-like

phases can take a wide variety of motif sizes, leading to unusual distribution tails. The region where the distribution is closest to TW, is the region where the system is most liquid like.

For the *Drosophila* data, we find that the motif size distribution is well approximated by TW in early times, whilst the EC layer is irregular and the system is most liquid-like. However, at later times when the final state of a more regular lattice appears, the distribution is no longer close to TW, Fig. S20.

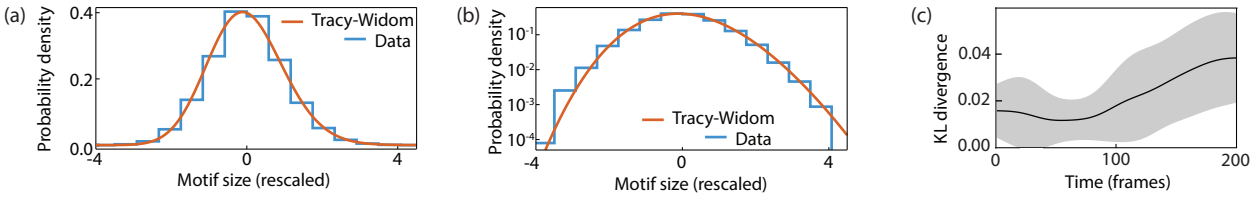

FIG. S20. Motif size distribution for the wings of *Drosophila* embryos is well approximated by TW at early times, but takes a different form at later times. (a) We combinine the first 100 frames into a single distribution, and rescale to have zero mean and variance 1. The resulting histogram is well approximated by the TW distribution. (b) same as (a) with log scale. (c) The KL divergence between the observed distribution and TW remains low until around 100 frames where it increases. Black line is local average, grey region shows $\pm 2\times$(standard deviation).

For the bacterial swarm data, while the motif distributions have varying mean and variance, the distributions appeared to have a common form when rescaled to have 0 mean and variance 1, Fig. S21a. We rescale every observed motif by the mean and variance of the distribution at that snapshot and combine in a single histogram. Since all distributions have similar mean, and are originally defined on the integers, the combined histogram still shows discrete peaks, Fig. S21b. We therefore choose histogram bins so that each contains exactly one discrete spike, and plot them as points so that the height represents the probability density, the x location represents the center of mass, with horizontal bars of $\pm$(within bin standard deviation) representing the width of the peak. We compare this against various known distributions, Fig. S21(c-d), as before with the PV tesselation. We once again find that TW fits the data well, although since we are using $r = 2$ here, they disagree slightly at the left tail.

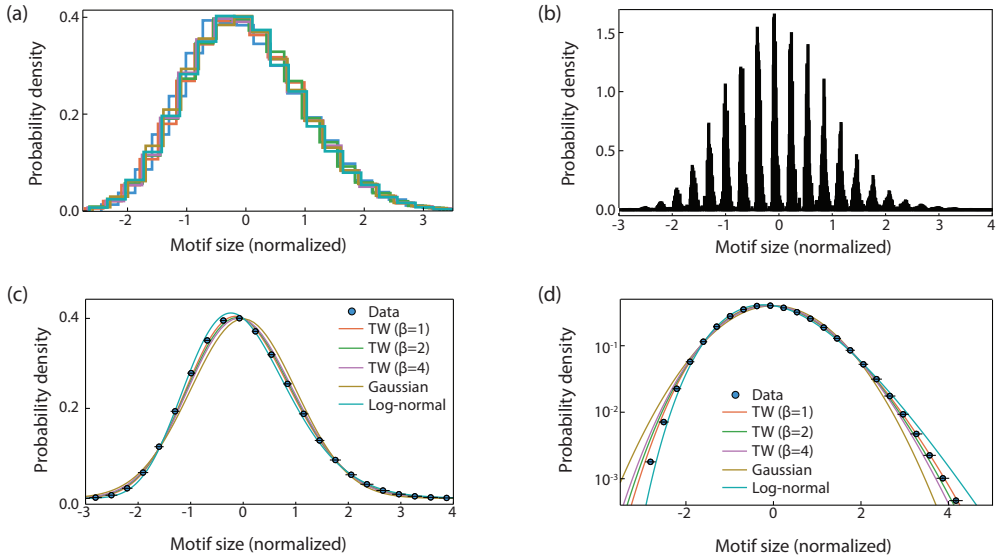

FIG. S21. Observed motif size distribution against plausible known distributions. (a) Randomly chosen snapshot distributions rescaled lie on the same curve. (b) Overall motif distribution with each snapshot data rescaled to have 0 mean and variance 1. Peaks are visible due to the integer sizes of the motifs, as well as a smooth envelope. (c-d) Histogram with bins containing a single peak, compared against known distributions with normal (c) and log scale (d).

[S1] F. Aurenhammer, R. Klein, and D.-T. Lee, *Voronoi Diagrams and Delaunay Triangulations*, 1st ed. (World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2013).
[S2] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel, IEEE Trans. Information Theory **29**, 551 (1983).
[S3] R. Etournay, M. Popović, M. Merkel, A. Nandi, C. Blasse, B. Aigouy, H. Brandl, G. Myers, G. Salbreux, F. Jülicher, and S. Eaton, eLife **4**, e07090 (2015).
[S4] L. Weinberg, IEEE Trans. Circuit Theory **13**, 142 (1966).
[S5] E. A. Lazar, J. K. Mason, R. D. MacPherson, and D. J. Srolovitz, Phys. Rev. Lett. **109**, 095505 (2012).
[S6] P. Bose and F. Hurtado, Comput. Geometry **42**, 60 (2009).
[S7] R. A. Wagner and R. Lowrance, J. ACM **22**, 177 (1975).
[S8] M. Essid and J. Solomon, SIAM J. Sci. Comput. **40**, A1961 (2018).
[S9] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, Science **290**, 2319 (2000).
[S10] Y. Fily, S. Henkes, and M. C. Marchetti, Soft Matter **10**, 2132 (2014).
[S11] J. K. Mason, E. A. Lazar, R. D. MacPherson, and D. J. Srolovitz, Phys. Rev. E **86**, 051128 (2012).
[S12] J. Solomon, arXiv:1801.07745.
[S13] H. Jeckel, E. Jelli, R. Hartmann, P. K. Singh, R. Mok, J. F. Totz, L. Vidakovic, B. Eckhardt, J. Dunkel, and K. Drescher, Proc. Nat. Acad. Sci. **116**, 1489 (2019).
[S14] U. Schmidt, M. Weigert, C. Broaddus, and G. Myers, in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, edited by A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, and G. Fichtinger (Springer International Publishing, Cham, 2018) pp. 265–273.
[S15] O. Ronneberger, P. Fischer, and T. Brox, in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lecture Notes in Computer Science, Vol. 9351 (Springer, 2015) arXiv:1505.04597.
[S16] J. Solomon, R. Rustamov, L. Guibas, and A. Butscher, arXiv:1603.06927.
[S17] T. J. Yoon, M. Y. Ha, E. A. Lazar, W. B. Lee, and Y.-W. Lee, Phys. Rev. E **100**, 012118 (2019).
[S18] E. A. Lazar, J. K. Mason, R. D. MacPherson, and D. J. Srolovitz, Phys. Rev. E **88**, 063309 (2013).
[S19] E. Limpert, W. A. Stahel, and M. Abbt, BioScience **51**, 341 (2001).