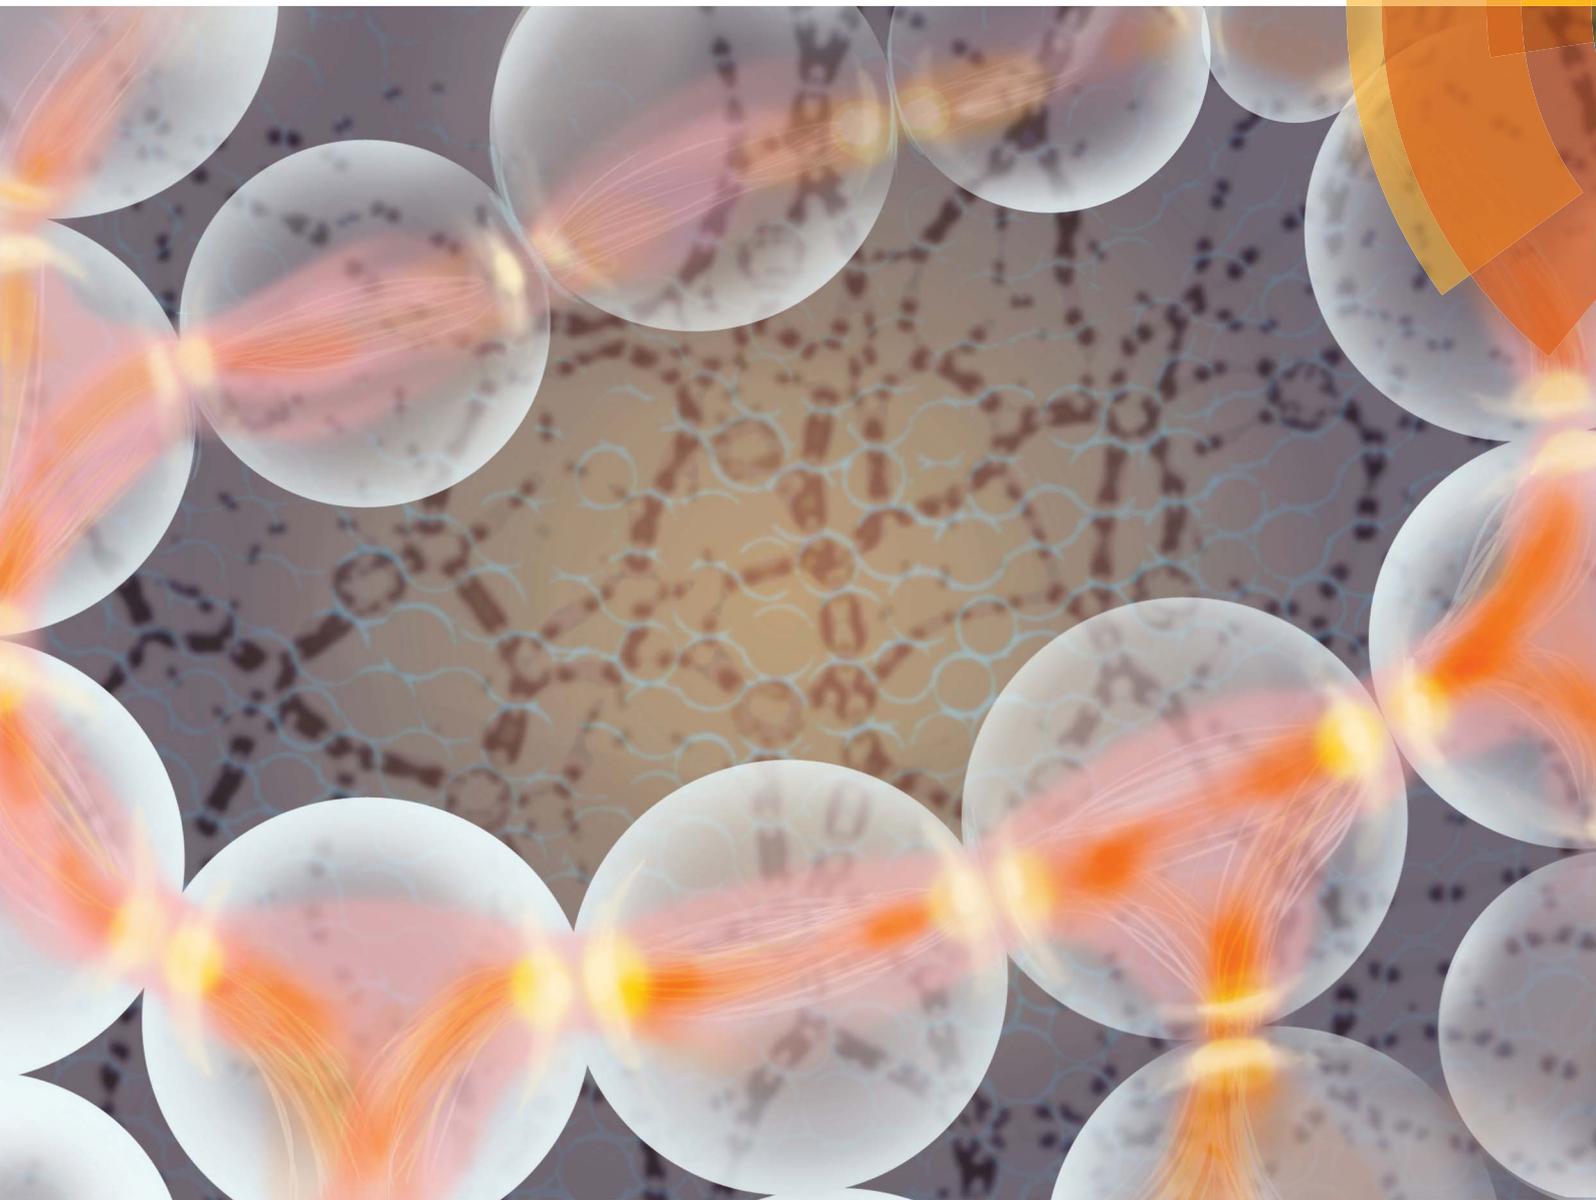


Soft Matter

www.softmatter.org



ISSN 1744-683X



PAPER

Danielle S. Bassett *et al.*

Extraction of force-chain network architecture in granular materials using community detection

Cite this: *Soft Matter*, 2015, 11, 2731

Extraction of force-chain network architecture in granular materials using community detection

 Danielle S. Bassett,^{*ab} Eli T. Owens,^{cd} Mason A. Porter,^{ef} M. Lisa Manning^g
and Karen E. Daniels^c

Force chains form heterogeneous physical structures that can constrain the mechanical stability and acoustic transmission of granular media. However, despite their relevance for predicting bulk properties of materials, there is no agreement on a quantitative description of force chains. Consequently, it is difficult to compare the force-chain structures in different materials or experimental conditions. To address this challenge, we treat granular materials as spatially-embedded networks in which the nodes (particles) are connected by weighted edges that represent contact forces. We use techniques from community detection, which is a type of clustering, to find sets of closely connected particles. By using a *geographical null model* that is constrained by the particles' contact network, we extract chain-like structures that are reminiscent of force chains. We propose three diagnostics to measure these chain-like structures, and we demonstrate the utility of these diagnostics for identifying and characterizing classes of force-chain network architectures in various materials. To illustrate our methods, we describe how force-chain architecture depends on pressure for two very different types of packings: (1) ones derived from laboratory experiments and (2) ones derived from idealized, numerically-generated frictionless packings. By resolving individual force chains, we quantify statistical properties of force-chain shape and strength, which are potentially crucial diagnostics of bulk properties (including material stability). These methods facilitate quantitative comparisons between different particulate systems, regardless of whether they are measured experimentally or numerically.

Received 17th August 2014
Accepted 6th February 2015

DOI: 10.1039/c4sm01821d

www.rsc.org/softmatter

1. Introduction

Particulate matter comes in many forms: it ranges from frictionless emulsions and frictional granular materials to bonded composites and biological cells. A long-known hallmark of particulate systems is the heterogeneous distribution of inter-particle forces within them^{1–6}—a phenomenon that has come to be known as *force chains*. The term “force chain” arises from the appearance—particularly in two-dimensional (2D) systems—of chains of particles that transfer force from one particle to another along the chain. In Fig. 1, for example, the eye readily focuses on the backbone that arises from the set of bright particles. These particles experience forces that are larger than average, and one can see such forces by using photoelastic

particles and placing them between crossed polarizers. In both two and three dimensions, it is increasingly common to measure inter-particle forces in a variety of systems – including frictional grains,^{6,7} frictionless grains,⁸ and frictionless emulsions.^{9–11} This opens the door for achieving a better understanding of force chains in a variety of systems.

It is known that force chains are important for resisting shear⁵ and directing sound propagation,^{13–16} but little is known about which properties of these chains are universal to particulate systems and which are sensitive to details (such as friction, adhesion, boundary conditions, and body forces like gravity). Improved methods for automatically identifying force chains and quantifying their specific properties (*e.g.* size and shape) would yield a deeper understanding of how they impact mechanical properties of particulate systems.

Although the presence of force chains is a generic feature of particulate systems, the term “force chain” is often used colloquially, and the field still lacks a quantitative definition of the term. Recently, several techniques have been proposed that aim to identify which subset of particles form a “force-chain network.” Peters *et al.*¹⁷ calculated force chains in low friction by demanding two requirements: particles must occur in a “quasi-linear” arrangement, and they must contain a concentrated stress. They used an algorithm that requires a choice of a

^aDepartment of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA. E-mail: dsb@seas.upenn.edu

^bDepartment of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA

^cDepartment of Physics, North Carolina State University, Raleigh, NC 27695, USA

^dDepartment of Physics, Presbyterian College, Clinton, SC 29325, USA

^eOxford Centre for Industrial and Applied Mathematics, Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

^fCABDyN Complexity Centre, University of Oxford, Oxford, OX1 1HP, UK

^gDepartment of Physics, Syracuse University, Syracuse, NY 13244, USA

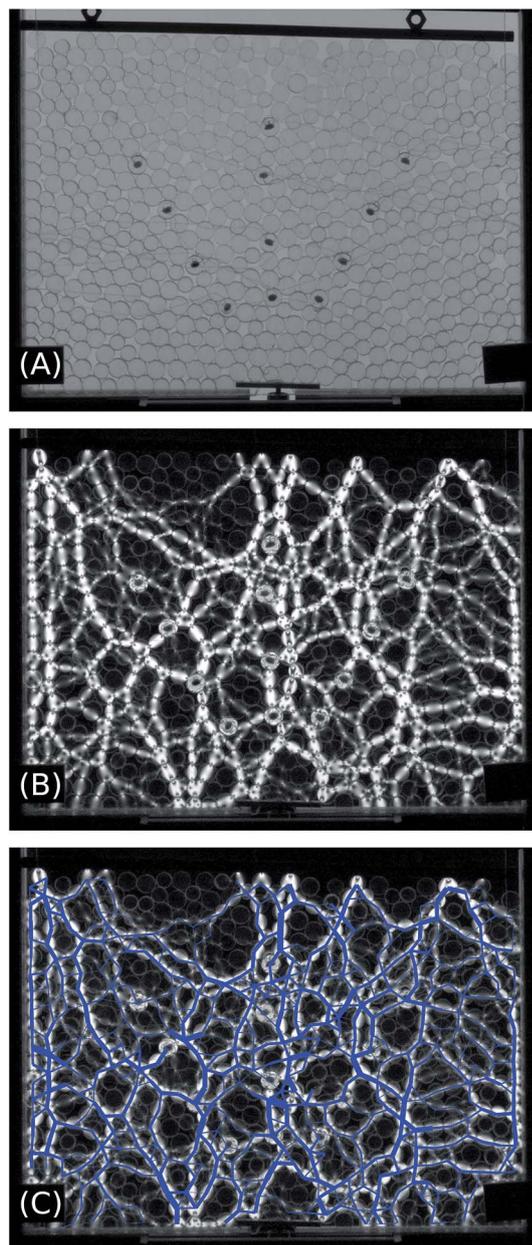


Fig. 1 We represent granular materials as spatially-embedded networks¹² whose nodes (particles) are connected by weighted edges that represent contact forces. (A) Image of a 2D vertical aggregate of photoelastic disks that are confined in a single layer by a pressure of $P \approx 6.7 \times 10^{-4}E$. Several particles are embedded with a piezoelectric sensor, for which the wires are visible. (B) The internal stress pattern in the photoelastic particles manifests as a network of force chains. (C) The blue line segments show the edges of a weighted graph, which we determine from image processing and overlay on the image from panel (B). An edge between two particles (*i.e.*, nodes) exists if the two particles are in physical contact with each other; the forces between particles give the weights of the edges. Note: the orientation of the packings is larger horizontal coordinates to the right, and larger vertical coordinates upward.

threshold value for each of the two conditions, and their notion does not allow force chains to branch. Using discrete-element simulations with a flat-punch geometry, Peters *et al.*¹⁷ obtained force chains that contain approximately half of the particles in a

system. Additionally, they reported that the lengths of these chains satisfy an approximately exponential distribution. Arevalo *et al.*¹⁸ examined polygonal structures that underly a force-chain network, which they determined by selecting all inter-particle forces above some threshold. In their numerical simulations, they observed that triangular structures dominate the network near a critical packing fraction. More recently, topological tools such as computational homology have been used to address the question of how to determine the structure of force chains.^{19–21} For example, using numerical simulations, Kondic *et al.*¹⁹ were able to distinguish between frictional and frictionless packings *via* a topological invariant known as the 0th Betti number.²² They demonstrated that the 0th Betti number changes with packing density, and they used a force threshold to isolate strong forces from weak forces.

The above investigations have been informative, but they also share a common viewpoint that the strongest inter-particle forces form the backbone of a particulate system. However, even a linear chain of strong particles would not be stable against buckling without the participation of the particles that lie alongside them.²³ Therefore, it is necessary to develop techniques that do not include a minimum threshold force to be able to consider a particle to be part of a network.

The purpose of the present paper is to use an approach based on network science²⁴ to develop an example of an appropriate method. Network science provides a powerful set of tools to represent complex systems by focusing not only on the components of such systems, but also on the interactions among those components. A network representation is particularly appropriate for particulate systems, where particles (network nodes) are connected to one another by contact forces (network edges).

In Fig. 1, we show an example network, which we constructed from laboratory experiments. Network approaches in granular materials have already had several successes in describing the dynamics of granular materials. For example, prior studies have investigated spatial patterns in the breaking of edges in a (binary) contact network under shear^{25–28} and the influence of network topology on acoustic propagation.²⁹ The latter study established the importance of using weighted contact networks to take into account the strength of inter-particle forces. Several other papers have also recently contributed to this line of inquiry using a variety of different network-based approaches.^{30–33}

In the present paper, we use network representations to build a set of practical methods to (1) automatically extract force chains from force networks (without the use of thresholding) and (2) define a set of three scalar quantities that one can use to characterize and classify different particulate materials. In Section II, we present two commonly-used granular systems (which are rather different from each other) in which it is possible to measure inter-particle forces. These two case studies—one in a “wet” laboratory and the other based on a standard computational model for jammed packings—provide a basis for describing the utility of our new technique. In Section III, we describe our method to automatically extract the force-chain structure from a force network.

Our method uses community detection, which is a form of clustering for networks.^{34,35} In contrast to previous uses of community detection in particulate systems,^{29,36} our method incorporates a *geographical null model* to account for spatial embeddedness. This geographical null model makes it possible to identify clusters of particles that are more densely interconnected *via* strong contact forces than expected given their geographical proximity. We then calculate a *gap factor* that allows one to optimize the resolution at which the detected communities are maximally branched. One can use the gap factor to quantify similarity of a community's geometry to that expected in a force chain. The gap factor is larger when communities are more branch-like and thus more similar to the expected geometry of a force chain; it is smaller when communities are either more compact or more string-like, and thus less similar to the expected geometry of a force chain. In addition to the gap factor, we use two other diagnostics—size and network force—to help characterize particulate structures. (We define the three diagnostics in Section III B.) In Section IV, we then apply our methodology to force networks that we measure in two very different case studies. We demonstrate the utility of our diagnostics for (1) quantifying changes in force-chain structure as a function of the confining pressure applied to a granular material and (2) facilitating the comparison and classification of force-chain network architectures across different media. We conclude and discuss several implications of our work in Section V.

II. Methods

To better understand which features of force chains are universal and to demonstrate that our methodology for identifying force chains is insightful for particulate systems in general, we compare and contrast results from two very different case studies: granular experiments and frictionless simulations. In both situations, the particles are confined in two dimensions, and they each consist of packings of bidisperse disks that interact with each other in a Hertzian-like manner. Both cases are jammed under constant pressure and use approximately the same number of particles (600). The key differences between the two systems are (1) the presence of friction and gravitational pressure in the experiments and (2) the use of periodic boundary conditions in the simulations. Additional differences include a different fabric tensor due to different initial conditions and the fact that laboratory experiments have a slight fine-scale polydispersity for particles of the “same” size in addition to the coarse-scale bidispersity. For each of the experiments and simulations, we measure force-contact networks for (approximately) the same seven different values of confining pressure.

A. Granular experiments

We perform experiments on a vertical 2D granular system of bidisperse disks that are confined between two sheets of Plexiglas. The particles are 6.35 mm thick, their diameters are $d_1 = 9$ mm and $d_2 = 11$ mm (which yields a diameter ratio of approximately 1.22), and they are cut from Vishay PSM-4 photoelastic material to

provide measurements of the internal forces. These particles have an elastic modulus of $E = 4$ MPa. We produce new configurations by rearranging the particles by hand. We increase the pressure by placing additional brass weights on the top surface of the packing. The values of pressure, which we report in units of the elastic modulus E (recall that the configuration is two-dimensional), are $2.7 \times 10^{-4}E$, $4.1 \times 10^{-4}E$, $6.7 \times 10^{-4}E$, $1.1 \times 10^{-3}E$, $2.2 \times 10^{-3}E$, $3.8 \times 10^{-3}E$, and $5.9 \times 10^{-3}E$. Particles are constrained by vertical walls to prevent deformation due to these pressures from occurring in the direction perpendicular to the loading direction. Such constraints on deformation can influence the shape of force chains, which tend to form in the direction of the major principal stress.^{6,26,37} See ref. 16, 29 and 38 for additional details about the experiments.

For each of 21 particle configurations and the seven values of pressure, we compute particle positions and forces using two high-resolution pictures of the system. We use one image, which we take without the polarizers, to determine the particle positions and contacts. (See³⁹ for a description of the technique.) We take particles to be in contact if the force between them is measurable by our photoelastic calculations. Using a second image that we take with the polarizers, we then determine the particle contact forces by solving the inverse photoelastic problem.³⁹

B. Frictionless simulations

We perform numerical simulations of a 50 : 50 ratio of bidisperse frictionless disks. We use a diameter ratio of 1.22, and particles interact *via* a Hertzian potential in a box with periodic boundary conditions in both directions and zero gravity.^{40–42} This model has been well-studied and it is significantly different from our experimental system with respect to friction, gravity, and the fixed boundaries. We generate mechanically-stable packings *via* a standard conjugate-gradient method.⁴⁰ We then perform simulations for a fixed packing fraction and volume, and we analyze 20 mechanically-stable packings at each packing fraction ϕ . We choose the seven values of the packing fraction so that the mean pressure \bar{p} at that packing fraction⁴³ matches the ones in the experiments: $[\phi, \bar{p}] = [0.8499, 3 \times 10^{-4}E]$, $[0.8521, 4 \times 10^{-4}E]$, $[0.8560, 7 \times 10^{-4}E]$, $[0.8621, 11 \times 10^{-4}E]$, $[0.8760, 22 \times 10^{-4}E]$, $[0.8927, 38 \times 10^{-4}E]$, and $[0.9106, 59 \times 10^{-4}E]$, where the modulus E is defined as the energy scale for the Hertzian interaction ε divided by the mean Voronoi area of a particle in the packing. The smallest value of ϕ provides a data point for a jammed packing that is less dense than what is accessible in our experiments.

III. Force-chain extraction and characterization

A. Force-chain extraction

To locate the force chains, we begin by recording which particles are in contact with (and exert force on) which other particles. In network language, we represent each particle as a node, and we represent each inter-particle contact as an edge whose weight is given by the magnitude of the force at that contact. We

thereby construct a force-weighted contact network \mathbf{W} from a list of all inter-particle forces. If particle i and j are in contact, then $W_{ij} = f_{ij}/\text{mean}(f)$, where f_{ij} is the normal force between them. If two particles are not in contact, then $W_{ij} = 0$. In addition, we let $W_{ii} = 0$. We also construct an unweighted (*i.e.*, binary) matrix \mathbf{B} whose elements are

$$B_{ij} = \begin{cases} 1, & W_{ij} \neq 0, \\ 0, & W_{ij} = 0. \end{cases}$$

The matrix \mathbf{B} is often called an “adjacency matrix”,²⁴ and the matrix \mathbf{W} is often called a “weight matrix.”

To obtain force chains from \mathbf{W} , we want to determine sets of particles for which strong inter-particle forces occur amidst densely connected sets of particles. We can obtain a solution to this problem *via* “community detection”,^{34,35,44} in which we seek sets of densely connected nodes called “modules” or “communities.” A popular way to identify communities in a network is by maximizing a quality function known as *modularity* with respect to the assignment of particles to sets called “communities.” Modularity Q is defined as

$$Q = \sum_{ij} [W_{ij} - \gamma P_{ij}] \delta(c_i, c_j), \quad (1)$$

where node i is assigned to community c_i , node j is assigned to community c_j , the Kronecker delta $\delta(c_i, c_j) = 1$ if and only if $c_i = c_j$, the quantity γ is a resolution parameter, and P_{ij} is the expected weight of the edge that connects node i and node j under a specified null model.

One can use the maximum value of modularity to quantify the quality of a partition of a force network into sets of particles that are more densely interconnected by strong forces than expected under a given null model. The resolution parameter γ provides a means of probing the organization of inter-particle forces across a range of spatial resolutions. To provide some intuition, we note that a perfectly hexagonal packing with non-uniform forces should still possess a single community for small values of γ and should consist of a collection of single-particle (*i.e.*, singleton) communities for large values of γ . At intermediate values of γ , we expect maximizing modularity to yield a roughly homogeneous assignment of particles into communities of some size (*i.e.*, number of particles) between 1 and the total number of particles. (The exact size depends on the value of γ .) The strongly inhomogeneous community assignments that we observe in the laboratory and numerical packings (see Section IV) are a direct consequence of the disorder in the packings.

An important choice in maximizing modularity optimization is the null model P_{ij} .^{45,46} The most common null model for modularity optimization is the Newman–Girvan (NG) null model^{34,35,47,48}

$$P_{ij}^{\text{NG}} = \frac{k_i k_j}{2m}, \quad (2)$$

where $k_i = \sum_j W_{ij}$ is the strength (*i.e.*, weighted degree) of node i and $m = \frac{1}{2} \sum_{ij} W_{ij}$. The NG null model is most appropriate for

networks in which a connection between any pair of nodes is possible. Importantly, many networks include (explicit or implicit) spatial constraints that exert a strong influence on which edges are present.¹² For particulate systems, numerous edges are simply physically impossible, so it is important to improve upon the NG null model for such applications. We use the term *geographical constraints* to describe the explicit spatial constraints in such systems. These constraints exert a significant effect on network structure, so it is important to take them into account when choosing a null model. For granular materials (and other particulate systems), each particle can only be in contact (*i.e.*, $W_{ij} \neq 0$) with its immediate neighbors. We therefore use a null model, which we call the *geographical null model*, to account for this constraint. The geographical null model is

$$P_{ij} = \rho B_{ij}, \quad (3)$$

where ρ is the mean edge weight in a network and \mathbf{B} is the binary adjacency matrix of the network (such a null model was used previously for applications in neuroscience⁴⁶). Recall that the adjacency matrix encapsulates the presence or absence of contact between each pair of particles. For a granular material, $\rho = \bar{f} := \langle f_{ij} \rangle$ is the mean inter-particle force. Because we have normalized the edge weights in the force network ($W_{ij} = f_{ij}/\text{mean}(f)$), we note that in our case $\rho = 1$.

Maximizing Q yields a so-called “hard partition” of a network into communities in which the total edge weight inside of modules is as large as possible relative to the chosen null model. A hard partition assigns each node to exactly one community. (An alternative is a “soft partition”,⁴⁹ which allows each node in a network to be associated with multiple communities.) For the geographical null model in (3), maximizing Q assigns the particles into communities that have inter-module particle forces that are larger than the mean force. Such communities represent the force chains in a granular system.

Because maximizing Q is NP-hard,⁵⁰ the success of the maximization is subject to the limitations of the employed computational methods. In the present paper, we use a Louvain-like locally greedy algorithm.⁵¹ Additionally, given the numerous near-degeneracies in the modularity landscape that tends to inflict networks that are constructed from empirical data⁵² (*i.e.*, many different partitions often yield comparably large values of Q), we report community-detection results that are ensemble averages over 20 optimizations.

B. Properties of force chains

We characterize communities using several diagnostics: size, network force, and a *gap factor* (a novel notion that we introduce in the present paper). The *size* s_c of a community c is simply the number of particles in that community. The systemic size s is given by the mean of s_c over all communities. The modularity Q is composed of sums of magnitudes of bond forces and therefore also has units of force. Therefore, we use the term *network force* to indicate the contribution of a community c to modularity. The network force of a community is given by the formula

$$\sigma_c = \sum_{ij \in C} [W_{ij} - \gamma \rho B_{ij}], \quad (4)$$

where C is the set of nodes in community c . The systemic network force σ is the mean of σ_c over all communities. Communities that correspond to force chains composed of densely packed particles with large forces between them have large values of network force, whereas communities that correspond to force chains composed of sparsely packed particles with small forces between them have small values of network force.

To identify the presence of gaps and the extent of branching in the geometry of a force-chain network, we calculate the Pearson correlation between “physical distance” (which we measure using the standard Euclidean metric) and “hop distance” (which is often called the “topological distance” and counts distance measured only along network edges), and we examine how the correlation depends on the force-chain topology. Communities with compact or linear-chain characteristics (see the main diagonal in Fig. 2) occur when there is a perfect correlation between hop distance and physical distance.

The perfect correlation arises because particles that are one hop away from each other (*i.e.*, they are adjacent to each other in the binary contact network) are also 1 particle-distance away. Small linear chains occur close to the origin, where both hop distance and physical distance are small, whereas long linear chains occur in the upper right quadrant of Fig. 2. Note that we use the term “compact” in the spirit of the mathematical sense of the term, although our exact meaning is somewhat different: “compact” communities of particles are at the opposite extreme as linear chains. Small compact blobs contain particles that are 1 hop-distance away from one another and 1 particle-distance away from one another, and they therefore occur near the origin of Fig. 2. Larger compact blobs contain particles that are several hop-distances away (and an equal amount of particle-distances away), and they therefore occur in the upper right quadrant of Fig. 2.

In contrast to compact blobs and linear chains, force chains with gaps, branches, and rings have a larger hop distance than

physical distance (see the upper triangle in Fig. 2). Particles that are close together in space do not necessarily have strong forces to bind them; the presence of more complicated shapes decreases the correlation between the physical and hop distance. (Note for particulate systems that the lower triangle is unphysical, as it would require forces between particles that are not in contact with one another to achieve a larger physical distance than hop distance).

To identify communities that are composed of branched structures, we measure the amount of correlation between the hop distance and the physical distance in the set of all node pairs in a given community. To compute the hop distance of a community, we define the community contact network \mathbf{B}^c . Its elements B_{ij}^c are entries of the matrix \mathbf{B} for which the corresponding nodes have both been assigned to the same community c . We then calculate the path lengths between possible pairs of nodes in a community c using the hop distance on the matrix \mathbf{B}^c . The resulting matrix of pairwise distances is \mathbf{L}_t . To compute physical distance, we calculate the Euclidean distances between all possible pairs of nodes in a community. The resulting matrix of pairwise distances is \mathbf{L}_p .

In defining the community gap factor, we choose to weight each community by its size. We thus weight large communities more heavily than small communities in linear proportion to the number of particles that they contain. In this case, the *gap factor* g_c of a community c measures the presence of gaps and the extent of branching in a community. We calculate it using with the formula

$$g_c = 1 - \frac{r_c s_c}{s_{\max}}, \quad (5)$$

where r_c is the value of the Pearson correlation between the upper triangle of \mathbf{L}_t and the upper triangle of \mathbf{L}_p , and s_{\max} is the size of the largest community. (Note that we exclude the diagonal elements of the matrices \mathbf{L}_p and \mathbf{L}_t). To provide further illustration of this quantity, a set of communities colored by their respective gap-factors is shown in Fig. 3C.

We define the (weighted) systemic gap factor as

$$g = 1 - \frac{1}{n} \sum_c \frac{r_c s_c}{s_{\max}}, \quad (6)$$

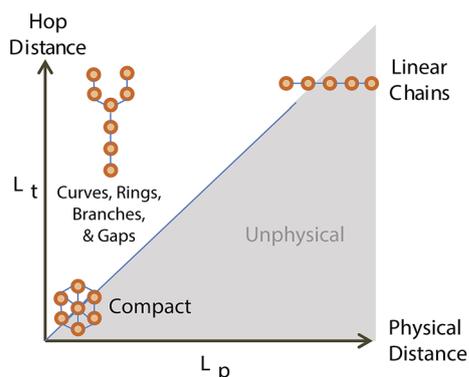


Fig. 2 Schematic illustration of the *gap factor*, which we measure via the Pearson correlation between the hop distance L_t and physical distance L_p in granular force networks. Network communities with gaps, branches, and rings can have a larger hop distance than physical distance. They thus reside in the upper triangle.

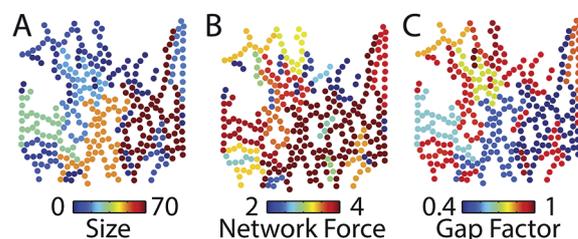


Fig. 3 Sample community-detection results for a single granular experiment at $4.1 \times 10^{-4} E$. We color communities according to (A) community size s_c , (B) network force σ_c , and (C) community gap factor g_c . In these calculations, we use $\gamma_{\text{opt}} = 0.9$, which we choose to maximize the systemic gap factor g across all pressures (rather than for a single pressure). This example illustrates that the three network diagnostics can reveal very different spatial distributions in the data and thereby supports examining all three diagnostics.

where the quantity n is the total number of communities (including singletons) and we again note that large communities contribute more than small communities to the averaging in eqn (6). An alternative choice for the systemic gap factor is to weight all communities uniformly to calculate an alternative systemic gap factor

$$g_{\text{uniform}} = 1 - \frac{1}{n} \sum_c r_c. \quad (7)$$

As we discuss in more detail in the appendix, branched communities typically include more nodes than linear communities, so the weighted gap factor g tends to be larger for communities with more branching. By contrast, g_{uniform} tends to be larger for more linear communities. For the remainder of the main manuscript, we focus on the weighted gap factor g because it varies far less across packings. Further studies of g_{uniform} would also be interesting.

IV. Characterizing force chains

We now characterize the force-chain networks of the experiments and simulations that we described in Section II using the methodology that we described in Section III. Our community-detection procedure consists of two stages. First, we maximize modularity (eqn (1)) with the geographical null model (eqn (3)) for different values of the resolution parameter γ : from $\gamma = 0.1$ to $\gamma = 2.1$ in increments of $\Delta\gamma = 0.2$. We then choose a resolution that approximately maximizes the systemic gap factor g [see eqn (6)]. This ensures that we extract communities of particles that have strong and dense force-weighted contacts with one another (*i.e.*, our network partition has a large value of the modularity Q), and that they are spatially sparse (*i.e.*, they tend to have a small topological-physical distance correlation r_c). The communities that we obtain tend to take the form of chain-like structures that are reminiscent of force chains; this provides visual support that our technique is successful. Our communities are much better than what one obtains using the NG null model. In previous work using the NG null model,²⁹ we observed that these latter communities always tend to be compact in form. We know, however, that communities with other qualitative features are very common. (See the schematic in Fig. 2.)

The resolution parameter γ sets the spatial resolution of the communities^{34,53,54} [see eqn (1)]. By tuning γ , one can either examine large communities (using small values of γ) or small communities (using large values of γ). In Fig. 4A, we show an example computation using data from experiments. (Note that we do not show single-particle communities). Observe that small values of γ produce communities that are dominated by compact structures (small g), whereas large values of γ produce communities that are dominated by linear structures (small g as well). This suggests that we can identify an optimal value γ_{opt} that maximizes the systemic gap factor g . This also corresponds to the choice for which the detected communities are most similar to the force-chain structure that we observe visually (see Fig. 4B). For this particular packing, we identify $\gamma_{\text{opt}} = 0.7$ as the

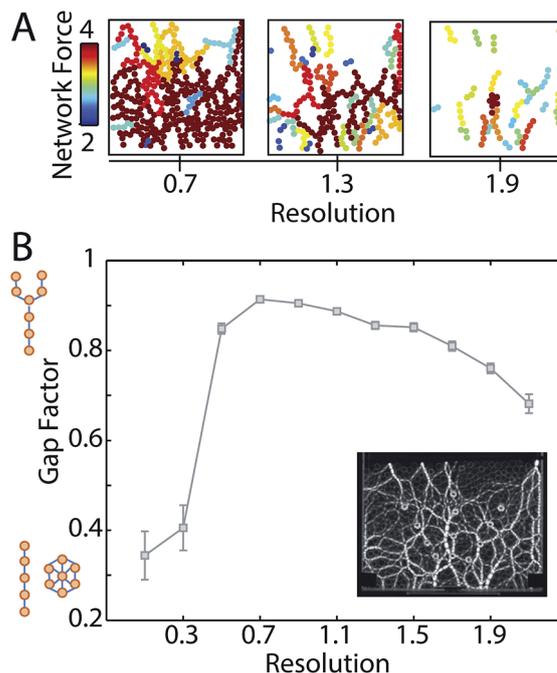


Fig. 4 Sample community-detection results for a single granular experiment at $4.1 \times 10^{-4}E$ illustrate multiresolution structure in chain-like communities. (A) Community structure as a function of the resolution parameter γ . Color indicates the logarithm (base 10) of the network force σ_c of community c . (We don't show single-particle communities.) (B) The systemic gap factor g [see eqn (6)] as a function of the resolution parameter γ exhibits a maximum at $\gamma = 0.7$ (for $\gamma \in \{0.1, 0.3, \dots, 2.1\}$). The error bars indicate the standard deviation of the mean over the laboratory packings. In the inset, we show an image of the 2D vertical packing of photoelastic disks that we use for panels (A) and (B).

best choice in our examination of $\gamma \in \{0.1, 0.3, \dots, 2.1\}$. As one can see in Fig. 4A, we observe at all values of γ that the detected communities vary in their size, network force, and gap factor. We also note that single-particle communities are common for all values of γ near 1 (and become increasingly common as γ increases after it exceeds 1).

Community detection *via* modularity maximization with a resolution parameter chosen to yield a maximal systemic gap factor provides an automated approach for detecting force chains in granular media. Once identified, we then calculate the size s , network force σ , and gap factor g to describe key features of a force-chain network. In the following subsections, we quantitatively examine how the distributions of s , σ , and g vary as a function of pressure for both granular experiments and frictionless simulations.

A. Granular experiments

Our experiments are at seven different pressures, which range from $2.7 \times 10^{-4}E$ to $5.9 \times 10^{-3}E$. As we illustrate in Fig. 5A, we observe that the communities become more compact (so they are closer to the diagonal line in Fig. 2) as pressure increases. At a small value of the resolution parameter ($\gamma = 0.7$), the detected communities metamorphose from branched structures to

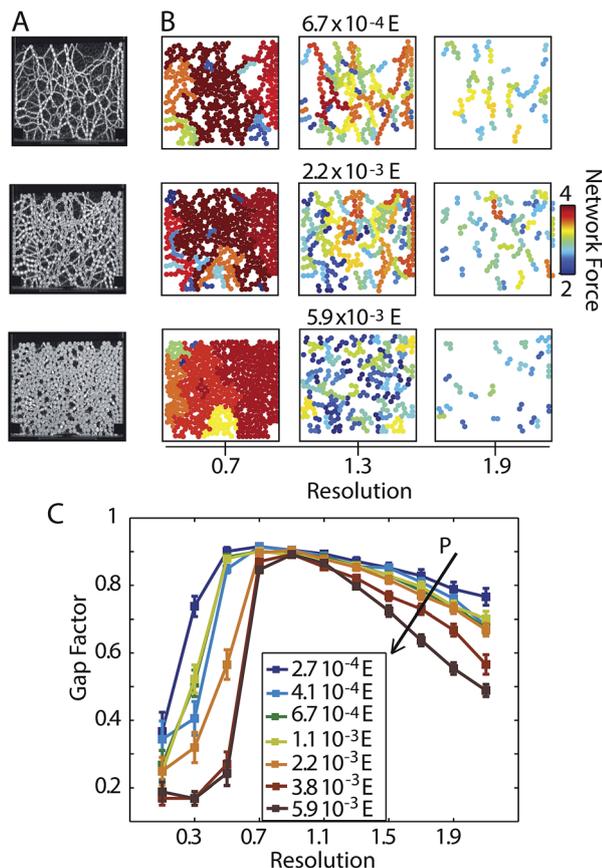


Fig. 5 (A) Images of experimental 2D vertical packings of photoelastic disks. These images reveal the manifestation of the internal stress pattern in a set of photoelastic particles as a network of force chains. (B) Community structure as a function of the resolution parameter γ for the following pressures: (top) $6.7 \times 10^{-4} E$, (middle) $2.2 \times 10^{-3} E$, and (bottom) $5.9 \times 10^{-3} E$. Color indicates the logarithm (base 10) of the network force σ_c of community c . (C) Gap factor as a function of both the resolution parameter γ and the pressure. The error bars indicate a standard deviation of the mean over the laboratory packings. The arrow emphasizes increasing pressure.

compact domains as pressure increases. At a large value of the resolution parameter ($\gamma = 1.9$), the detected communities tend to shrink from many-particle chains to 2-particle chains as pressure increases. This provides a way to quantify our earlier observations that force-chain networks at higher pressures are more homogeneous (observations at small γ) and less chain-like (observations at large γ) than they are at lower pressures.

At all pressures (and despite the aforementioned differences), we find a maximum in the systemic gap factor g as a function of the resolution parameter γ (see Fig. 5B). The shape of the curve of the gap factor *versus* the resolution-parameter has a more pronounced peak at higher pressures than at lower pressures: larger slopes descend from and (especially) lead up to the values of γ near the maximum. At all pressures, small values of γ select compact communities and large values of γ select two-particle communities. In between, we observe a value of γ at which the most branch-like structures appear; we refer to this value that approximately maximizes the gap factor as the

“optimal” value. This optimal value changes slightly as a function of pressure; for example, $\gamma_{\text{opt}} = 0.9$ at $5.9 \times 10^{-3} E$ and $\gamma_{\text{opt}} = 0.7$ for $2.7 \times 10^{-4} E$. High-pressure packings also exhibit a much smaller systemic gap factor than low-pressure packings at both small and large values of the resolution parameter. This observation is consistent with both the increased compactness of the communities and the increasingly homogeneous nature of the force-chain structure as one increases the pressure.

The resolution that maximizes the gap factor identifies structures in a force network that are most reminiscent of the force chains that are apparent by eye; in other words, it identifies branching communities. Therefore, to extract force chains from force networks across different packings and pressures, we examine community structure for a range of resolution parameters and identify the resolution-parameter value that approximately maximizes the gap factor. We refer to the communities detected at γ_{opt} as the “force chains” in our calculations, and we characterize their properties in terms of their size, network force, and gap factor. For an illustrative example of a single packing whose communities are color-coded by either size, network force, or gap factor (see Fig. 3).

As we illustrate in Fig. 6, we observe that the size and network force of the force chains have approximately exponential distributions for all seven pressures. This identifies that the majority of communities are relatively small and weak, but a few communities are relatively large and strong. By contrast, the gap-factor distribution is skewed: most communities have a large gap factor, and only a few communities have a small gap factor. In high-pressure packings, communities tend to be less broadly distributed with respect to both size and network force; this is consistent with the homogenization of force chains as pressure increases.

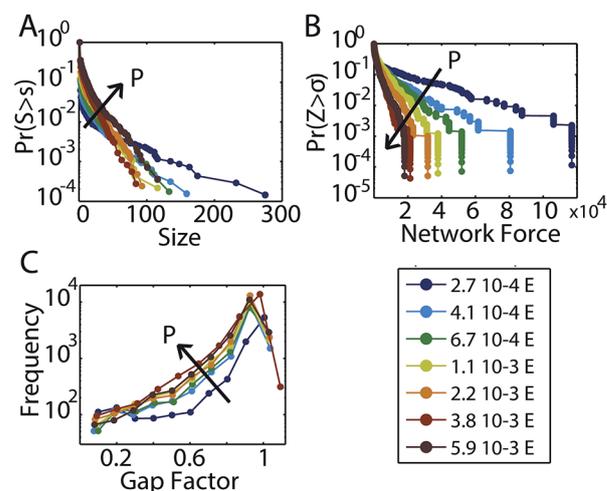


Fig. 6 Cumulative probability distributions of (A) community size s_c and (B) network force σ_c for all communities. (C) Histogram of the gap factor g_c for all communities. In these calculations, we use $\gamma_{\text{opt}} = 0.9$, which we choose to maximize the systemic gap factor g across all pressures (rather than for a single pressure). The arrows emphasize increasing pressure.

B. Frictionless simulations

To illustrate quantitative differences in force-chain structure between laboratory and numerical packings, we use our method to identify force chains in frictionless simulations. In addition to their lack of friction, the numerical packings also differ from our experiments in that they use periodic boundary conditions, have zero gravity, have zero fine-scale polydispersity, and have a different fabric tensor due to their different initial conditions. (A commonality between the two types of packings is that they are both bidisperse). As in Section IV A, we find communities *via* modularity maximization at different values of the resolution parameter for different packing fractions, which we select to match the seven pressures from the experiments.

As we show in Fig. 7A, the simulated packings change qualitatively as a function of the resolution parameter. For relatively small values of the resolution parameter ($\gamma = 0.7$), the granular force network has collapsed into a single compact

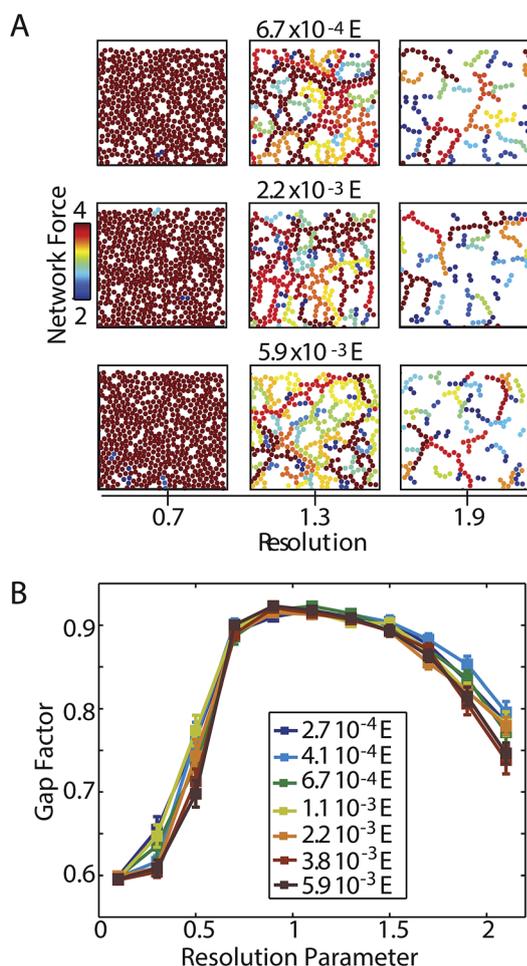


Fig. 7 Multiresolution structure of chain-like communities in numerical packings in frictionless simulations as a function of pressure. (A) Community structure as a function of the resolution parameter γ for the following pressures: (top) $6.7 \times 10^{-4} E$, (middle) $2.2 \times 10^{-3} E$, and (bottom) $5.9 \times 10^{-3} E$. Color indicates the logarithm (base 10) of the network force σ_c of community c . (B) Gap factor as a function of both the resolution parameter γ and the pressure. The error bars indicate the standard deviation of the mean over the numerical packings.

community. In contrast, for relatively large values of the resolution parameter ($\gamma = 1.3$), we identify many small branch-like communities that are reminiscent of force chains. Unlike in the granular experiments, we do not observe a strong qualitative change in community structure as a function of pressure (regardless of the value of γ). Compare Fig. 5B and 7B.

In order to select the most branched community structure, we again identify a value γ_{opt} that is associated with the maximum value of the systemic gap factor g . As with the experimental packing, there is a clearly identifiable maximum (see Fig. 7B), which occurs at $\gamma = 0.9$ for $\gamma \in \{0.1, 0.3, \dots, 2.1\}$. We also observe that the shape of $g(\gamma)$ is more consistent across pressures for the simulations than it is for the experiments.

As in the laboratory packings, we refer to the communities at γ_{opt} as our “force chains,” and we characterize their properties in terms of size, network force, and gap factor (see Fig. 8). Across the set of pressures from $2.7 \times 10^{-4} E$ to $5.9 \times 10^{-3} E$, we observe that chain size and network force have approximately exponential distributions. This indicates that the majority of communities are relatively small and weak, but a few communities are relatively large and strong. The gap-factor distribution has a left skew, which indicates that most communities have a large gap factor and only a few communities have a small gap factor.

In contrast to the laboratory experiments, the shape of the distributions for size, network force, and gap factor in frictionless simulations do not change dramatically as a function of pressure. Nevertheless, the frictionless simulations do exhibit a systematic difference in the mean size and mean network force for communities as a function of pressure [see Fig. 8B].

C. Comparison between laboratory and numerical force chains

Using our methodology of extracting force chains, it is possible to quantitatively compare and contrast two particulate systems

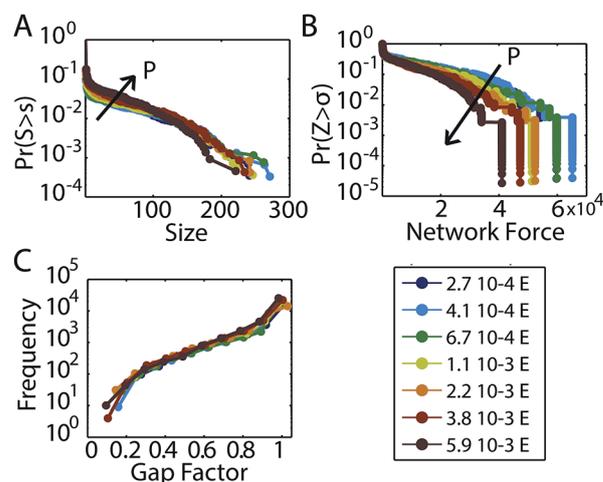


Fig. 8 Size and network force of chain-like communities in numerical packings as a function of pressure. Cumulative distributions of (A) community size s_c and (B) network force σ_c for the resolution-parameter value ($\gamma = 0.9$) that maximizes the gap factor in the seven different pressures separately. (C) Histogram of the gap factor g_c for the communities at $\gamma = 0.9$. The arrows emphasize increasing pressure.

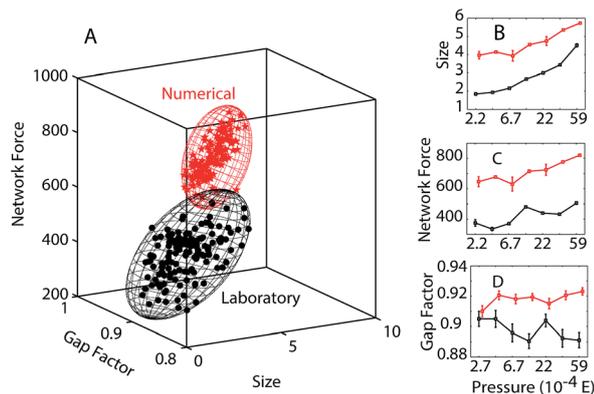


Fig. 9 Comparison between the laboratory and numerical packings. (A) Scatter plot of $[s, \sigma, g]$ for each of the 21 runs and seven different pressures. (B, C, D) we average the same diagnostics over all equal-pressure runs for (B) size, (C) network force, and (D) gap factor. All of these calculations use a resolution-parameter value of $\gamma_{\text{opt}} = 0.9$. The error bars indicate the standard deviation of the mean over the 21 runs.

(such as the aforementioned granular experiments and frictionless simulations). In Fig. 9, we summarize the three main diagnostics (size, network, force, and gap factor) that we calculate for each of our two case studies. We anticipate that these diagnostics might provide a helpful means of characterizing how well a given set of simulations—*e.g.*, with anisotropic cell shapes^{55,56} or different models of friction^{57,58}—captures experimental features. It also provides a systematic way to compare the properties of different particulate systems or to monitor the temporal evolution of a particulate system.

The first key difference between our results for the experimental and numerical packings is that the granular experiments have force chains with smaller mean size and network force (compare Fig. 9A and B). This result is consistent with the more homogeneous nature of force chains in the frictionless packings than in the laboratory packings. Laboratory packings exhibit a few large and strong force chains (see Fig. 6A and B), but the majority of a packing tends to be dominated by singleton communities (see Fig. 6A). By contrast, the frictionless packings have community-size and network-force distributions that are far less skewed (see Fig. 8). Additionally, the majority of such a packing is dominated by the force chains (see Fig. 8A); there are few singletons. These differences in the homogeneity of the packings and in the distributions of size and network force leads on average to larger and stronger force chains in the frictionless numerical packings than in the laboratory ones.

We also observe that the gap factor is smaller on average in the laboratory packings than in the frictionless packings. This result is also consistent with the homogenization of structure in frictionless packings, but it might also be influenced by differences between the two types of packings from the presence *versus* absence of gravity. The frictionless packings are gravity-free, whereas the laboratory packings are influenced by gravity. In contrast to the frictionless packings, the laboratory packings therefore often exhibit linear vertical chains, which decrease the systemic gap factor. (See our discussion in Appendix 2).

Finally, we observe that the size and network force of chains increases with pressure in laboratory packings, whereas the systemic gap factor decreases with pressure for such packings. These observations indicate the presence of a changing length scale, as is to be expected.^{59,60} As pressure increases, the force structure becomes more homogeneous, leading to extended sections of material with densely packed particles that exert strong forces on one another. However, the force structure becomes less branch-like, leading to a decrease in the systemic gap factor. The frictionless numerical packings also exhibit increases in size and network force of chains with pressure; however, unlike in the laboratory packings, the frictionless packings do not exhibit a decrease in systemic gap factor with pressure. This quantitative difference between the two types of packings yields a technique for measuring the visual differences between Fig. 5A and 7A, in which it is clear that the geometry of the force chains is affected less by pressure in the simulated setting than in the laboratory setting.

Fig. 9B also demonstrates that the size of communities that resemble force chains decreases as pressure approaches 0, which provides some elucidation for a longstanding question in the field of disordered matter. In particulate systems at 0 temperature, the point at which pressure $P = 0$ is the jamming transition, and much recent work has been devoted to attempting to understand the nature of this transition.⁶¹ There is now strong evidence that there is a *growing* length scale,^{59,60} which corresponds to the size of regions that are mechanically unstable, near the transition. As we discuss in Section V, it is natural to postulate that our force-chain communities are negatively correlated with weak regions in a particulate packing. Therefore, we expect force chains to shrink as mechanically unstable regions grow.

V. Conclusions and discussion

In this paper, we treated granular materials as spatially-embedded networks in which nodes (particles) are connected by edges whose weights we determine from contact forces. We developed and applied a network-based clustering method, in which we detect tightly-connected “communities” of particles *via* modularity maximization with a geographical null model, to extract chain-like structures that are reminiscent of force chains in both numerical and laboratory packings. From these chain-like structures, we calculated three key diagnostics (size, network force, and gap factor), and we illustrated their utility for identifying and characterizing force-chain network architectures across two types of packings (laboratory and numerical) and across seven different pressures. To characterize force chains, we identified an “optimal” resolution-parameter value that approximately maximizes the gap factor in each of these scenarios. An alternative approach would be to examine the structure of force chains that are identified at a set of resolution parameters that one optimizes separately by considering other parameters (*e.g.*, size, network force, *etc.*). This is an interesting direction for future study.

After we identified an optimal value for the resolution parameter, we systematically evaluated and compared

properties of force-chain communities. A common feature that we observed in both the (frictional) laboratory and (frictionless) numerical packings is that the distribution of force-chain sizes is consistent with an exponential distribution. In the laboratory packings of 2D granular materials, we found that high-pressure force networks exhibit compact rather than branching communities, which is consistent with the notion that increasing pressure causes a breakdown in the long-range heterogeneous structure in a material. In contrast, the geometry of the force chains in the frictionless numerical packings appear to be affected less by pressure. The force chains in the numerical packings also tend to be both larger and stronger than their counterparts in laboratory packings. Together, these results support the conclusion that the force-chain structure, topology, and geometry is different in the two systems.

Methodological insights

An important conclusion of our paper is that choosing an appropriate null model is critical for extracting information from community detection in networks. We have developed and applied a geographical null model for the detection of network communities, which are strongly reminiscent of force chains, in particulate materials. The power of this null model lies in its ability to fix a network's topology (which is given by a binary contact network) while scrambling its geometry (which is given by a force-weighted contact network). The geographical null model thereby includes more of the fundamental physics of particulate systems than the Newman-Girvan null model that is commonly used in modularity maximization. An interesting future direction would be to incorporate different physical constraints and principles (*e.g.*, force balance) and to examine the different results that one obtains by including different combinations of relevant physical ideas. A key question in community detection is what is the minimal set of physical ideas—and, more generally, the minimal amount of problem-specific information—that one can include in a null model to get answers that are more insightful than using a one-size-fits-all null model like the Newman-Girvan model.

Practical utility

We expect that our methodology will provide a framework for understanding which features of force chains are universal *versus* which are governed by detailed particle-particle or particle-environment interactions. For example, one can use our methodology to provide a quantification of the similarity of force chains between a simulation and a given experiment. It can also provide similar quantifications between two experiments (or two simulations) with different types of particles or boundary conditions. It is also a viable tool to help predict differences in macroscopic behavior based on subtle differences in the mean or distributions of force-chain diagnostics.

Our methodology also provides additional information that is not available *via* traditional methods for identifying force chains. For example, we know which particles are strongly connected within communities, and we can also observe the spatial arrangement of the various communities that comprise a force network. This allows one to ask whether the arrangement of communities helps to govern the linear and nonlinear response of a disordered packing. For example, it is possible that large, strong communities indicate a region of relatively high mechanical stability or that boundaries between these communities indicate areas of weakness. In the future, these types of investigations should be helpful for obtaining a better understanding of the onset of flow or failure in particulate systems.

Appendix 1: Effect of community weighting on force-chain extraction

In this appendix, we describe an alternative set of measurements that we make using the *uniformly-weighted systemic gap factor* g_{uniform} , which we defined in eqn (7).

Gap factor versus resolution parameter

In Fig. 10, we observe that the curve of the maximum of g_{uniform} versus the resolution parameter γ varies significantly with pressure in the (frictional) laboratory packings but not in the (frictionless) numerical packings. In the laboratory packings,

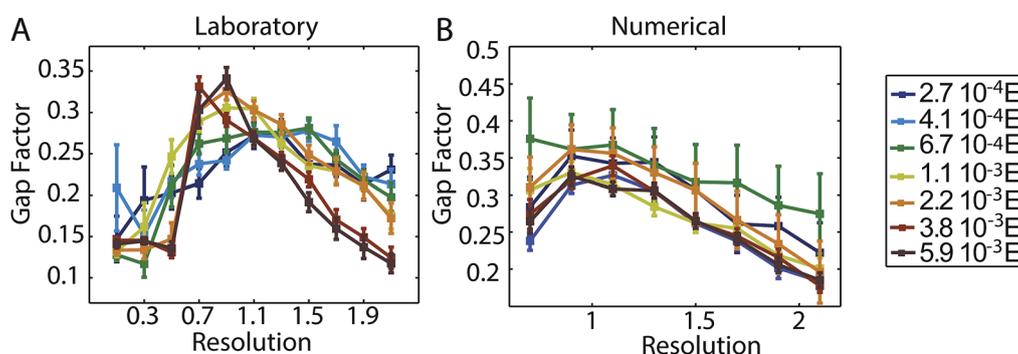


Fig. 10 (A) In the laboratory packings, the systemic uniformly-weighted gap factor g_{uniform} exhibits a maximum as a function of the resolution parameter γ between roughly $\gamma = 0.9$ and $\gamma = 1.3$ (for $\gamma \in \{0.1, 0.3, \dots, 2.1\}$) that is different for different pressures. The error bars indicate the standard deviation of the mean over the laboratory packings. (B) In the numerical simulations, g_{uniform} exhibits a maximum (which varies little with pressure) as a function of the resolution parameter γ at $\gamma = 1.1$. The error bars indicate the standard deviation of the mean over the numerical packings.

we observe a maximum of g_{uniform} at $\gamma = 0.9$ (for $\gamma \in \{0.1, 0.3, \dots, 2.1\}$) in high-pressure packings ($5.9 \times 10^{-3}E$) and at $\gamma = 1.5$ for low-pressure packings ($2.7 \times 10^{-4}E$). In the numerical packings, we observe a maximum of g_{uniform} at $\gamma = 1.1$ for all pressures. In comparison to our observations in the main text from employing the size-weighted systemic gap factor g , we find that the optimal value of γ is larger when we instead employ g_{uniform} (compare Fig. 10 to Fig. 5 and 7). We also observe that the curves of the systemic gap factor *versus* resolution parameter exhibit larger variation for the uniformly-weighted gap factor than for the size-weighted gap factor.

Optimal value of the resolution parameter

The large variation in the maximum of g_{uniform} over packings and pressures makes it difficult to choose an optimal resolution-parameter value. We choose to take $\gamma_{\text{opt}} = 1.1$ because (1) it corresponds to the maximum of g_{uniform} in the numerical packings and (2) it corresponds to the mean of the maximum of g_{uniform} in the laboratory packings. To facilitate the comparison of optimal values of γ from the two weighting schemes, we denote γ_{opt} for g as $\hat{\gamma}$ and we denote γ_{opt} for g_{uniform} as $\hat{\gamma}_{\text{uniform}}$. Note that $\hat{\gamma}_{\text{uniform}} = 1.1$ differs from (and is larger than) $\hat{\gamma} = 0.9$.

Force-chain structure at the optimal value of the resolution parameter

The force chains that we identify for the optimal value for the uniformly-weighted gap factor (at $\hat{\gamma}_{\text{uniform}} = 1.1$) differ from those that we identified in the main text for the optimal value of the size-weighted gap factor (at $\hat{\gamma} = 0.9$). We show our comparison in Fig. 11. For both laboratory and numerical packings, the force chains that we identify at $\gamma = 0.9$ are larger and more branched than the ones that we identify at $\gamma = 1.1$ (which are smaller and more linear). Indeed, the communities that we identify at $\gamma = 1.1$ have more singletons than the communities that we identify at $\gamma = 0.9$. These results follow from the difference in the two weighting schemes for calculating a systemic gap factor. The size-weighted systemic gap factor g weights larger communities more heavily than smaller ones, and the larger communities tend to be the more branched communities that we identify at smaller values of the resolution parameter (*e.g.*, at $\gamma = 0.9$). In contrast, the uniformly-weighted systemic gap factor g_{uniform} gives equal weight to small and large communities, and it therefore uncovers the linear communities that are evident at larger values of the resolution parameter (*e.g.*, at $\gamma = 1.1$). We can therefore use the size-weighted gap factor to identify larger, more branched force chains and the uniformly-weighted gap factor to identify smaller, more linear force chains.

Appendix 2: Methodological considerations

Robustness of community structure to errors in the estimation of contact forces

In our frictional laboratory experiments, we estimate that errors in the force measurements could be as large as $\pm 30\%$ of the

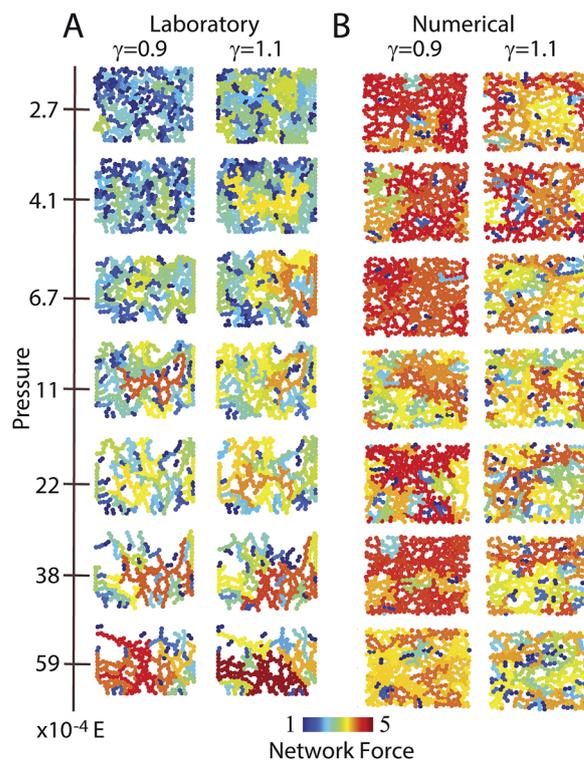


Fig. 11 In both (A) (frictional) laboratory and (B) (frictionless) numerical packings, we identify larger and more branched force chains at the optimal resolution determined by (*left*; $\gamma = 0.9$) the size-weighted gap factor g , and we identify smaller and less branched force chains at the optimal resolution determined by (*right*; $\gamma = 1.1$) the uniformly-weighted gap factor g_{uniform} . These observations are consistent across all pressure values, but they are especially evident at high pressures in the laboratory packings and are least evident at low pressures in the laboratory packings. In the numerical packings, we observe little variation for different values of pressure. In both panels, we highlight the network force for the force chains that we identify in example packings.

contact force f_{ij} ; the high variability arises from the nonlinear fitting process. (Recall that we take particles to be in contact if the force between them is measurable by our photoelastic calculations. We then determine the particle contact forces by solving the inverse photoelastic problem using images taken with polarizers³⁹). Somewhat surprisingly, we find that the errors in the force estimates are independent of both the local force and the global pressure. To ensure that our results are qualitatively robust to these variations, we construct 20 simulated force networks for each experimental network (21 packings, seven pressures) by adding Gaussian noise with width $f_{ij}/3$ to each contact. For each of these simulated networks, we reevaluate the estimated community structure (from which we infer the force chains).

To determine whether the estimates are robust to this amount of noise, we compare the community structure of the actual force networks with those of the simulated force networks using the z-score of the Rand coefficient.⁶² For comparing two partitions α and β , we calculate the Rand z-score in terms of the network's total number M of pairs of nodes, the number M_α of pairs that are in the same community in partition

α , the number M_β of pairs that are in the same community in partition β , and the number $w_{\alpha\beta}$ of pairs that are assigned to the same community both in partition α and in partition β . The z-score of the Rand coefficient comparing these two partitions is

$$z_{\alpha\beta} = \frac{1}{\sigma_{w_{\alpha\beta}}} \left(w_{\alpha\beta} - \frac{M_\alpha M_\beta}{M} \right), \quad (8)$$

where $\sigma_{w_{\alpha\beta}}$ is the standard deviation of $w_{\alpha\beta}$ (as in⁶²). Let the *mean partition similarity* z denote the mean value of $z_{\alpha\beta}$ over all possible partition pairs for $\alpha \neq \beta$.

We observe that the assignment of particles to communities (and consequently to force chains) in the experimental laboratory force networks is, on average, statistically similar (the z-scores for similarity are larger than 18) to the assignment of particles to communities in the simulated networks constructed by adding Gaussian noise with width $f_{ij}/3$ to each contact (see Fig. 12A). These results indicate that our estimates of network

communities is relatively robust to the empirical measurement error in the inter-particle contact forces.

To further clarify the robustness of our algorithm to variations in inter-particle forces on the order of $\pm 30\%$ of the contact forces, we perform similar calculations for the numerical packings. For each numerical packing, we construct 20 simulated force networks by adding Gaussian noise with width $f_{ij}/3$ to each contact. For each of these simulated networks, we reevaluate the estimated community structure (from which we infer the force chains). We observe that the assignment of particles to communities (and consequently to force chains) in the numerical force networks is, on average, statistically similar (the z-scores for similarity are larger than 9) to the assignment of particles to communities in the simulated networks constructed by adding Gaussian noise with width $f_{ij}/3$ to each contact (see Fig. 12B). These results indicate that our estimates of network communities are relatively robust to measurement error in the inter-particle contact forces in the form of $\pm 30\%$ of the contact force.

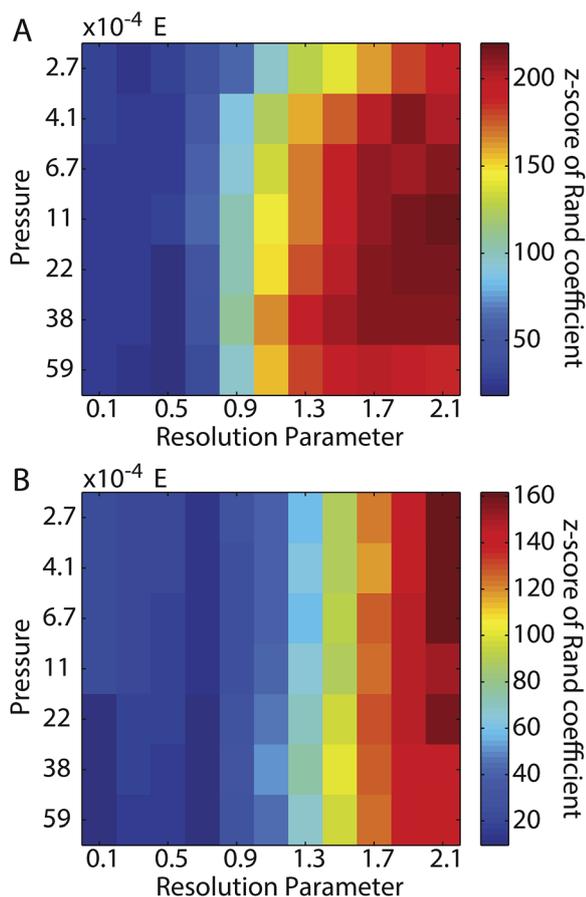


Fig. 12 In both (A) laboratory and (B) numerical packings, we calculate the similarity between the assignment of particles to communities for the actual packings as compared to the packings in which we have added Gaussian noise with width $f_{ij}/3$ to each contact. We estimate partition similarity using the z-score of the Rand coefficient; we compute the z-score for each pair of actual-noisy networks, for each pressure, for each packing, and for each value of the resolution parameter. In this figure, we show the mean z-score of the Rand coefficient over the 21 packings and over the 20 noisy instantiations of the actual force networks.

Effects of gravity on community structure in laboratory packings

In the laboratory packings, gravity can play an important role in the heterogeneity of observed force chains. To determine the impact of gravity on network diagnostics, we calculate the mean vertical position of particles in each network community for each packing, pressure, and optimization of the modularity quality function for the resolution-parameter value ($\gamma = 0.9$) that maximizes the gap factor. For each pressure value separately, we collate these four diagnostics (size, network force, gap factor, and mean vertical position) for every community that we identify over packings and modularity optimizations. To decrease the potential for false positives, we then identify the unique community sizes, and average network force, gap factor, and mean vertical position over all communities of that size. This process serves as a data-reduction procedure. It decreases the degrees of freedom in subsequent statistical testing, and it thereby decreases the potential for false positives. We then ask whether the network diagnostics (size, network force, or gap factor) are significantly correlated with the mean vertical position of the communities. We use a Spearman rank correlation to increase our robustness to outliers in the data. For low values of pressure we observe a significant correlation between each of the three community-level network diagnostics (size, gap factor, and strength) and the mean vertical position of the particles in that community (see Fig. 13).

At pressures of $2.7 \times 10^{-4} E$ and $4.1 \times 10^{-4} E$, larger communities with larger network force tend to be identified at a lower vertical position of the packing; this is consistent with the compaction effects of gravity. At pressures of $4.1 \times 10^{-4} E$, communities with large gap factors tend to be observed at high vertical positions in the packings (where communities tend to be small and less compact). For values of pressure between $22 \times 10^{-4} E$ and $59 \times 10^{-4} E$, we observe no relationship between mean vertical position of particles in a community and community size, network force, or gap factor. These results

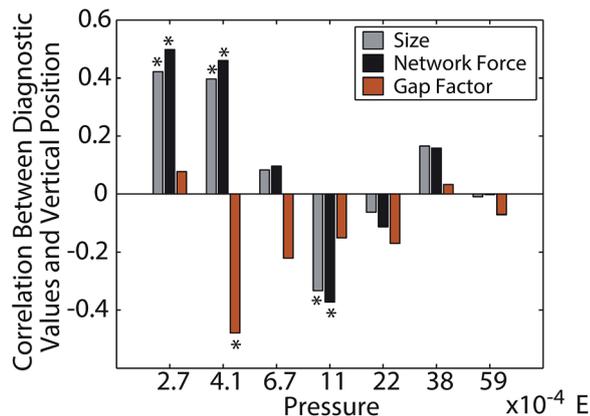


Fig. 13 Spearman correlation between network diagnostics (community size, network force, and gap factor) and the mean vertical position of the community as a function of pressure. Asterisks indicate significant correlations at the level of $p < 0.05$. (The p -value is uncorrected).

indicate that gravity can play a role in the shape of force chains, particularly at low pressures.

Acknowledgements

DSB was supported by the Alfred P. Sloan Foundation, the Army Research Laboratory through contract number W911NF-10-2-0022 from the U. S. Army Research Office, and the National Science Foundation award #BCS-1441502. KED and ETO were supported by NSF-DMR-0644743 and NSF-DMR-1206808, and they thank James Puckett for the development of the photoelastic image-processing code. MLM was supported by the Alfred P. Sloan Foundation, NSF-BMMB-1334611, and NSF-DMR-1352184. MAP was supported by the James S. McDonnell Foundation (research award #220020177), the EPSRC (EP/J001759/1), and the FET-Proactive project PLEXMATH (FP7-ICT-2011-8; grant #317614) funded by the European Commission. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

References

- P. Dantu, in *Proceedings of the fourth International Conference on Soil Mechanics and Foundation Engineering*, London, 1957, pp. 144–148.
- C. H. Liu, S. R. Nagel, D. A. Schecter, S. N. Coppersmith, S. Majumdar, O. Narayan and T. A. Witten, *Science*, 1995, **269**, 513.
- F. Radjai, M. Jean, J. J. Moreau and S. Roux, *Physical Review Letters*, 1996, **77**, 274.
- D. M. Mueth, H. M. Jaeger and S. R. Nagel, *Physical Review E*, 1998, **57**, 3164.
- D. Howell, R. P. Behringer and C. Veje, *Physical Review Letters*, 1999, **82**, 5241.
- T. S. Majumdar and R. P. Behringer, *Nature*, 2005, **435**, 1079.
- M. Saadatfar, A. P. Sheppard, T. J. Senden and A. J. Kabla, *Journal of the Mechanics and Physics of Solids*, 2012, **60**, 55.
- S. Mukhopadhyay and J. Peixinho, *Physical Review E*, 2011, **84**, 2.
- J. Brujić, S. F. Edwards, D. V. Grinev, I. Hopkinson, D. Brujić and H. A. Makse, *Faraday Discussions*, 2003, **123**, 207.
- J. Zhou, S. Long, Q. Wang and A. D. Dinsmore, *Science*, 2006, **312**, 1631.
- K. W. Desmond, P. J. Young, D. Chen and E. R. Weeks, *Soft Matter*, 2013, **9**, 3424.
- M. Barthelemy, *Phys Rep*, 2011, **499**, 1.
- H. A. Makse, N. Gland, D. L. Johnson and L. M. Schwartz, *Physical Review Letters*, 1999, **83**, 5070.
- R. Cruz Hidalgo, C. U. Grosse, F. Kun, H. W. Reinhardt and H. J. Herrmann, *Physical Review Letters*, 2002, **89**, 205501.
- E. Somfai, J.-N. Roux, J. H. Snoeijer, M. van Hecke and W. van Saarloos, *Phys Rev E*, 2005, **72**, 021301.
- E. T. Owens and K. E. Daniels, *Europhysics Letters*, 2011, **94**, 54005.
- J. Peters, M. Muthuswamy, J. Wibowo and A. Tordesillas, *Physical Review E*, 2005, **72**, 041307.
- R. Arévalo, I. Zuriguel and D. Maza, *Physical Review E*, 2010, **81**, 041302.
- L. Kondic, A. Goulet, C. S. O'Hern, M. Kramár, K. Mischaikow and R. P. Behringer, *Europhysics Letters*, 2012, **97**, 54001.
- M. Kramár, A. Goulet, L. Kondic and K. Mischaikow, *Physical Review E*, 2014, **90**, 052203.
- S. Ardanza-Trevijano, I. Zuriguel, R. Arévalo and D. Maza, *Phys. Rev. E*, 2014, **89**, 052212.
- H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction*, American Mathematical Society, 2010.
- F. Radjai, D. E. Wolf, M. Jean and J. J. Moreau, *Physical Review Letters*, 1998, **80**, 61.
- M. E. J. Newman, *Networks: An Introduction*, Oxford University Press, 2010.
- M. Herrera, S. McCarthy, S. Slotterback, E. Cephas, W. Losert and M. Girvan, *Physical Review E*, 2011, **83**, 1.
- A. Tordesillas, D. M. Walker and Q. Lin, *Physical Review E*, 2010, **81**, 011302.
- D. M. Walker and A. Tordesillas, *International Journal Of Solids And Structures*, 2010, **47**, 624.
- S. Slotterback, M. Mailman, K. Ronaszegi, M. van Hecke, M. Girvan and W. Losert, *Phys Rev E*, 2012, **85**, 021309.
- D. S. Bassett, E. T. Owens, K. E. Daniels and M. A. Porter, *Phys Rev E*, 2012, **86**, 041306.
- J. H. Lopez, L. Cao and J. M. Schwarz, *Phys. Rev. E*, 2013, **88**, 062130.
- R. Arévalo, L. A. Pugnaloni, D. Maza and I. Zuriguel, *Philosophical Magazine*, 2013, **93**, 4078.
- D. M. Walker, A. Tordesillas, M. Small, R. P. Behringer and C. K. Tse, *Chaos*, 2014, **24**, 013132.
- R. Navakas, A. Diugys and B. Peters, *Physica A: Statistical Mechanics and its Applications*, 2014, **407**, 312.
- M. A. Porter, J.-P. Onnela and P. J. Mucha, *Not Amer Math Soc*, 2009, **56**, 1082.
- S. Fortunato, *Phys Rep*, 2010, **486**, 75.

- 36 P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro and Z. Nussinov, *The European Physical Journal E*, 2011, **34**, 105.
- 37 F. Radjai and E. Azéma, *European Journal of Environmental and Civil Engineering*, 2009, **13**, 203.
- 38 E. T. Owens and K. E. Daniels, *Soft Matter*, 2013, **9**, 1214.
- 39 J. G. Puckett and K. E. Daniels, *Physical Review Letters*, 2013, **110**, 058001.
- 40 C. O'Hern, S. Langer, A. Liu and S. Nagel, *Phys. Rev. Lett.*, 2002, **88**, 075507.
- 41 C. S. O'Hern, L. E. Silbert, A. J. Liu and S. R. Nagel, *Physical Review E*, 2003, **68**, 011306.
- 42 M. L. Manning and A. J. Liu, *Physical Review Letters*, 2011, **107**, 108302.
- 43 For finite systems, there isn't a bijection between pressure and packing fraction. Because the packing fraction is fixed in these simulations, there is some variance (of roughly 9×10^{-5}) in the pressure between different packings.
- 44 M. E. J. Newman, *Nat Phys*, 2012, **8**, 25.
- 45 M. E. J. Newman, *Physical Review E*, 2006, **74**, 036104.
- 46 D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson and P. J. Mucha, *Chaos*, 2013, **23**, 013142.
- 47 M. E. J. Newman and M. Girvan, *Phys Rev E*, 2004, **69**, 026113.
- 48 M. E. J. Newman, *Proc Natl Acad Sci USA*, 2006, **103**, 8577.
- 49 D. M. Walker and A. Tordesillas, in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. June, 2014, pp. 1275–1278.
- 50 U. Brandes, D. Delling, M. Gaertler, R. Görke, M. Hofer, Z. Nikoloski and D. Wagner, *IEEE Trans on Knowl Data Eng*, 2008, **20**, 172.
- 51 I. S. Jutla, L. G. S. Jeub, and P. J. Mucha, *A generalized Louvain method for community detection implemented in MATLAB*, 2011–2012.
- 52 B. H. Good, Y. A. de Montjoye and A. Clauset, *Phys Rev E*, 2010, **81**, 046106.
- 53 J. Reichardt and S. Bornholdt, *Phys Rev E*, 2006, **74**, 016110.
- 54 J.-P. Onnela, D. J. Fenn, S. Reid, M. A. Porter, P. J. Mucha, M. D. Fricker and N. S. Jones, *Phys Rev E*, 2012, **86**, 036104.
- 55 M. Mailman, C. F. Schreck, C. S. O'Hern and B. Chakraborty, *Physical review letters*, 2009, **102**, 255501.
- 56 Z. Zeravcic, N. Xu, A. J. Liu, S. R. Nagel and W. van Saarloos, *Europhysics Letters*, 2009, **87**, 26001.
- 57 P. A. Cundall and O. D. L. Strack, *Geotechnique*, 1979, **29**, 47.
- 58 S. Papanikolaou, C. S. O'Hern and M. D. Shattuck, *Physical review letters*, 2013, **110**, 198002.
- 59 M. Wyart, L. Silbert, S. Nagel and T. Witten, *Phys Rev E*, 2005, **72**, 051306.
- 60 C. P. Goodrich, W. G. Ellenbroek and A. J. Liu, *Soft Matter*, 2013, **9**, 10993.
- 61 A. J. Liu and S. R. Nagel, *Annual Review of Condensed Matter Physics*, 2010, **1**, 347.
- 62 A. L. Traud, E. D. Kelsic, P. J. Mucha and M. A. Porter, *SIAM Rev*, 2011, **53**, 526.