Candidate number: 501601
Word count: 6257

TP19: Models of granular networks in two and three dimensions

Supervisor: Dr Mason Porter

I develop several new models that attempt to reproduce the structure of two-dimensional granular materials, and compare them with experiment using diagnostics from network science. I discuss one model that treats granular networks as a perturbation of a close-packed hexagonal lattice, and two models that are modifications of the random geometric graph (RGG). One of these modified RGGs simulates the forces that compressed particles exert on each other and displaces them accordingly. This model is found to be a good model of granular materials, in that its network properties appear to be very close to those of real granular networks; however, its effectiveness does not seem to depend upon the form of the force law implemented. I then briefly examine the three-dimensional analogues of two of the models.

1 Introduction

1.1 Granular materials

A granular material consists of a large number of macroscopic particles (for example, grains of sand) [1]. Understanding flows of granular materials is important in predicting geophysical hazards [2], as is understanding the nature of the *jamming transition* that occurs when a granular material becomes rigid under increasing pressure [3]. Further important yet poorly understood phenomena are the propagation of sound and the distribution of forces in granular packings. It has been proposed that the appearance of *force chains* may be a barrier to our understanding in this area [4]. The phrase 'force chain' refers to the observation that stresses in granular materials are transmitted through the system along chains of particles [5]. These distinctive chains only include a fraction of all the particles, and they snake through the material in an inhomogeneous and seemingly unpredictable manner (see Fig. 1). In [6], the structure of force chains was approached from a network-science perspective. It was found that the mean force on a particle is significantly correlated with its *intracommunity strength z-score*, which is a measure of the particle's connectivity within its local network community (see Sections 1.2 and 1.5 for descriptions of network connectivity and communities). The results in [6] indicate that a network-science perspective can provide powerful tools with which to probe the structure of granular materials. Motivated by this, I use diagnostics from network science as a means to assess the effectiveness of different models of granular materials.

This report is organised as follows. In the remainder of this section, I briefly describe some relevant concepts from network science and describe the origin of the experimental data that I will be comparing the models to. In Section 2, I describe each model in detail and explain its motivation. I also discuss each model's performance when compared with experiment using network diagnostics. In Section 3, I analyse the force-modified RGG model in more detail, comparing the distribution of forces and compressions that it predicts with results from experiment, and in Section 4 I investigate how the effectiveness

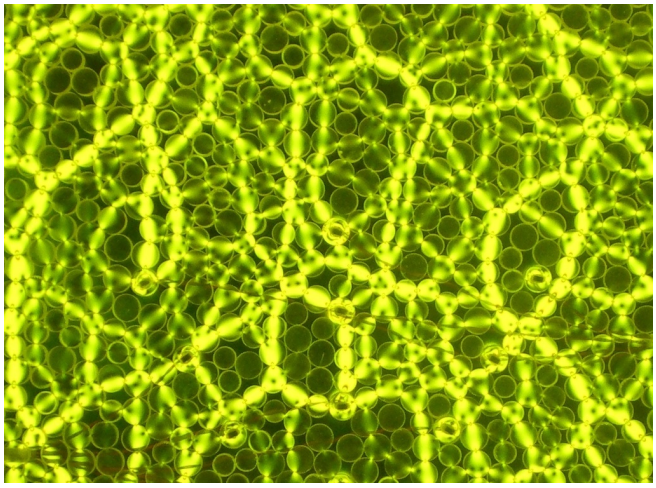

Figure 1: Force chains observed in a granular medium consisting of photoelastic disks, which emit light when compressed [7]. See Section 1.4. Image credit to Eli Owens and Karen Daniels (used with permission).

of this model is dependent on the precise form of the force law it implements. In Section 5, I extend two of the models to three dimensions and compare the results to the two-dimensional case. Finally, in Section 6, I discuss my findings and conclusions.

1.2 Spatial networks

A network is a set of *nodes* and an associated set of *edges* that define pairwise connections between the nodes. For a network with N nodes, we label the nodes $i = 1, 2, 3, \ldots, N$ and label edges using the unordered pair of integers (i, j), which then represents an undirected connection between nodes i and j[1].

A spatial network is a network whose nodes are embedded in some space possessing a metric [8]. A granular material can be used to define a network embedded in a 2D Euclidean plane, with particles represented by nodes that are connected by an edge if and only if the particles are in contact. This representation is a particularly natural one, so it seems reasonable to expect a network-science perspective to provide enlightening ways of analysing the structure of granular materials.

It is important to clarify what is meant when two particles are said to be "in contact". If two particles have positions $\mathbf{x}_i$ and $\mathbf{x}_j$ and respective radii r_i and r_j when uncompressed, we define them to be in contact if and only if $|\mathbf{x}_i - \mathbf{x}_j| \leq r_i + r_j$.

1.3 Some definitions

A network with N nodes can be described by its *adjacency matrix* [9], an $N \times N$ matrix whose elements take the values

$$A_{ij} = \begin{cases} 1, & \text{if } i \text{ and } j \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

A *weighted network* has a weight w_{ij} associated with the edge (i, j) that connects nodes i and j. These weights might come from the physical distance between two connected nodes, or can be the force that two particles exert on one another. For a weighted network, the weighted adjacency matrix is:

$$W_{ij} = \begin{cases} w_{ij}, & \text{if } i \text{ and } j \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

The adjacency matrix encodes all the information about the network and allows direct calculation of many useful quantities. For example, the *degree* k_i of node i is the number of nodes to which it is connected. In terms of the adjacency matrix,

$$k_i = \sum_j A_{ij}. \tag{3}$$

A *path* from node i to node j is a sequence of edges that connect a sequence of nodes starting at node i and terminating at node j. The length of the path is the number of edges in the path (or for weighted networks the sum of the weights of those edges). A *geodesic* path is a path between two nodes with the shortest possible length. Geodesic paths between two nodes are in general not unique.

A *planar* network is a network whose nodes can be placed in the plane in such a way that the edges do not intersect each other (except at the nodes).

1.4 Experimental data

Owens and Daniels [7] performed experiments on a set of bidisperse 2D granular packings. A *bidisperse* medium contains particles of two different sizes, while a *monodisperse* medium contains particles of identical size. A collection of disks were cut from Vishay PSM-4 photoelastic material, with thickness 6.35 mm and radii $r_1 = 4.5$ mm and $r_2 = 5.5$ mm. Approximately 1000 particles were packed into a container with an open top, so that they were confined only by gravity. The average *packing fraction* of these arrangements was 0.84 ± 0.01, where the packing fraction is defined as the fraction of the total area occupied by the particles. In total 17 different arrangements were studied, each of which was obtained by manually rearranging the disks. The photoelasticity of the particles means that the contact forces can be estimated by comparing photographs of the system to calibration images for known forces [6]. Data was thereby obtained for both the contact networks and the weighted force networks.

1.5 Network diagnostics

In network science, many diagnostics have been introduced as means of quantifying the properties of networks [9]. The beauty of applying network science to granular materials is that even though these diagnostics were not introduced with granular networks in mind, they can give insightful information about the structure of the network and hence the structure of the material.

[1] We can also define a *directed* network, in which edges are labelled using an *ordered* pair of integers (i, j), which represents a directed connection from node i to node j.

Network diagnostics can be defined for both weighted and unweighted networks. In my analyses, I compare only diagnostics of the unweighted contact networks. Below I list the diagnostics that I used. Diagnostics (with the exception of communicability and mean shortest distance) were calculated using code from the Brain Connectivity Toolbox [10]. See [6, 9] for precise definitions.

- *Node (edge) betweenness centrality.* Node (edge) betweenness centrality of a given node (edge) measures the number of geodesic paths that pass through that node (edge). This diagnostic gives an indication of a node's (or edge's) importance within a network — for example, a station in a rail network has a high node betweenness centrality if a large number of rail routes pass through it.

- *Clustering coefficient and transitivity.* These two diagnostics are closely related, and give information about local clustering within a network. A node with a high clustering coefficient has a high number of connections between its neighbours. Transitivity is the proportion of connected triplets of nodes that also form triangles.

- *Assortativity.* This quantifies the extent to which the degrees of connected nodes are correlated.

- *Global and local efficiency.* Global efficiency quantifies how well a signal transmits through a network, whereas the local efficiency of a node quantifies how well a signal transmits within a local subgraph that includes that node.

- *Mean shortest distance.* This is simply the mean length of the geodesic paths in a network.

- *Maximised modularity.* Modularity is related to the problem of community detection. The aim of community detection is to partition a network into non-overlapping communities such that the edge density within communities is large whereas connections between communities are sparse. Community detection algorithms maximise a quantity known as modularity — the maximised value then gives a measure of how well the network can be partitioned into communities[2].

- *Subgraph centrality.* This diagnostic quantifies the extent to which a given node participates in the subgraphs of a network.

- *Communicability.* The communicability of a pair of nodes attempts to describe the ease of communication between the nodes by counting the number of paths between them while down-weighting the contributions from longer paths.

In addition to these diagnostics, I also developed two new measures in the course of my investigations:

- $\mathrm{tr}(e^A - I)$. This quantity is closely related to communicability except that rather than counting paths between different nodes, it counts loops that start and end at the same node.

- *Path-weighted betweenness.* This diagnostic is identical to node betweenness centrality except that it down-weights the contribution from each path by a factor $1/l!$, where l is the length of the path.

For diagnostics that give a value for each individual node (or edge), I calculate the mean over all the nodes (or edges) to provide a single number to characterise the whole network.

2 Models

2D granular networks are highly constrained by some basic physical principles:

- They are embedded in 2D, so they must be planar;

- The compression between two particles in contact with one another is generally small compared to the radii of the particles, so there is a limit to how close any two nodes can be;

- There is a geometrical restriction on the maximum degree of each node: for monodisperse disks — or bidisperse disks whose radii ($r_2 > r_1$) are in the ratio $\frac{r_2}{r_1} < 1.3$ — no particle can be in contact with any more than six other particles[3].

[2] I use the Louvain algorithm [11] to optimise the modularity.

[3] Let r_1 and r_2 be such that n circles with radius r_1 can be packed around a circle with radius r_2. Then the centres of two adjacent smaller circles subtend an angle $\theta = 2\arcsin\frac{r_1}{r_1+r_2}$ at the centre of the larger circle (by simple trigonometry). Hence $\frac{r_2}{r_1} = \mathrm{cosec}\frac{\pi}{n} - 1$. In this case $n = 7$, so $\frac{r_2}{r_1} \approx 1.3$.

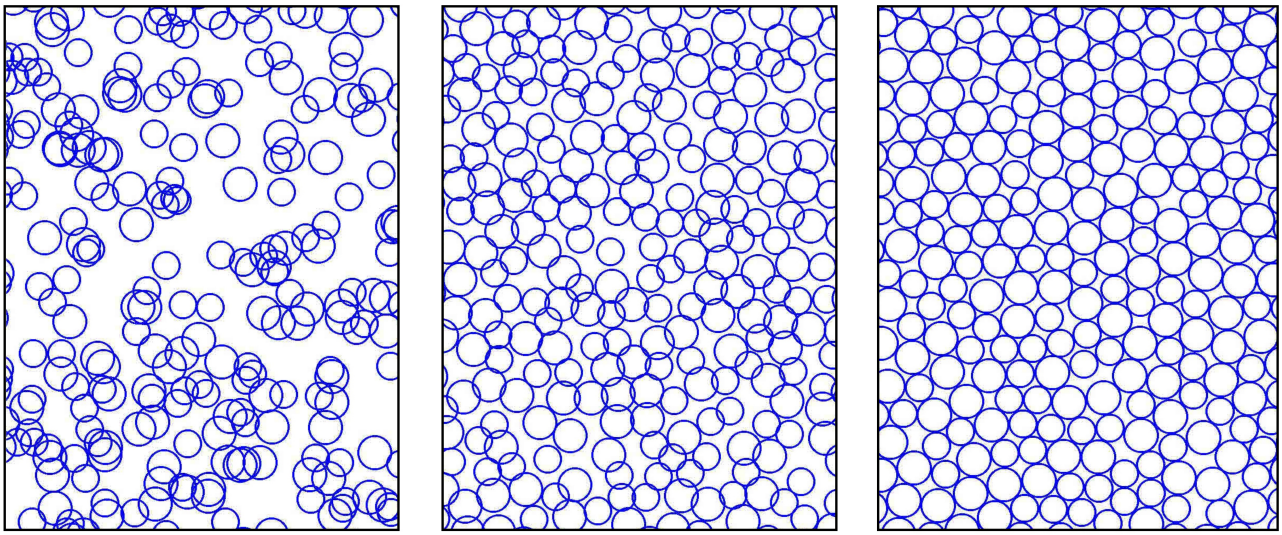

Figure 2: From left to right, sections of packings generated by the bidisperse RGG, the $p(\rho)$-modified RGG, and the force-modified RGG.

I develop different models to focus on each of these constraints. In the following subsections, I present a description of each model and a brief examination of its performance. I give the network diagnostics for each model along with experimental data in Table 1 in Appendix B.

2.1 Random geometric graph

A very crude model is the random geometric graph (RGG) in 2D. Here we place N particles uniformly at random within a rectangular subset of $\mathbb{R}^2$ (whose dimensions approximately match the dimensions of the experimental setup) and connect them by an edge if they are separated by a distance of less than $2R$. We choose R so that the *edge density* (the ratio of the number of edges to the number of nodes) of the resulting network matches the experimental data. Note that these RGGs will in general not be planar networks. Therefore, the only constraint captured by the RGGs is the dimensionality of the space in which the network is physically embedded, because we use the 2D Euclidean metric to define distance between nodes. This model has already been studied in the context of granular networks in [6], where the RGG is used as a null-model network with which to compare network diagnostics computed from the experimental data.

We can also crudely model a bidisperse system by assigning each particle a radius r_i that can take one of two values. We then connect two particles by an edge if the distance between them is $d_{ij} < R(r_i + r_j)$, where again R is a parameter that we can change to match the edge density with experiment.

The RGG is a poor model of granular networks. As has already been noted in [6], the network diagnostics are significantly different for an ensemble of RGGs versus the experimental networks (see Table 1 in Appendix B). We find that RGGs are locally more connected, with higher values of clustering coefficient and local efficiency than we find in the experimental data. This makes sense, since we have not imposed any constraint on the maximum degree of a node, and the non-planarity of the RGGs allows connections between particles that are physically impossible in a real granular network. Interestingly, the bidisperse version of the RGG fared no better in matching the experimental data than the monodisperse version.

However, the RGG results are useful because they give a baseline against which to compare the results from the other models. Comparing any diagnostic against its value for the RGGs allows us to get a sense of what order of magnitude to expect, and what constitutes a close agreement (or otherwise) with experiment. We will see that all of the other models perform significantly better than the RGGs. This is to be expected, because each of the models imposes constraints that are motivated by real effects in granular materials.

2.2 $p(\rho)$-modified RGG

There are many ways to modify the RGG model to more accurately reflect the structure of granular materials. One of the most evidently unphysical aspects of the simple RGG model is that two particles can

be arbitrarily close to one another. Particles exert a repulsive force when compressed, so there is a limit to how close we expect to find particles to each other, for a given packing fraction.

Define the *proximity* ρ, given an arrangement of pre-placed particles, as the distance from a given point to the nearest particle. Proximity is thus a function of position in the plane. We can then modify the RGG model by placing particles one by one — a point is chosen uniformly at random and a particle is placed there with a probability $p(\rho)$ — a function of the point's proximity. If $p \to 0$ for small proximities, then the resulting arrangement has a constraint on how close particles can be positioned to one another.

A simple choice for the function $p(\rho)$ is

$$p(\rho) = \begin{cases} 0, & \text{for } \rho < 2\alpha r_i, \\ 1, & \text{for } \rho \geq 2\alpha r_i, \end{cases} \tag{4}$$

where r_i is the radius of the particle being placed. The value of α puts a limit on the number of particles of a given radius that can be placed in a given area. I placed bidisperse particles into a rectangular region, matching the dimensions and the sizes of the particles to experiment. Given this setup, the highest value of α that allowed the appropriate number of particles to be placed was $\alpha \approx 0.8$. It is reasonable to use the highest value of α (< 1) as possible, because, as I have already noted, compressions between particles are generally small compared to their radii and so we should enforce as strict a constraint as possible on how close particles can be to one another.

Particles are then connected using a similar procedure to that used for the simple RGG model. Figure 2 shows particle arrangements generated from the simple RGG, the modified RGG, and the force-modified RGG (see Section 2.4). As might be expected simply from a visual comparison of these models, the $p(\rho)$-modified RGG performs much better than the original RGG in almost every diagnostic (see Table 1 in Appendix B).

2.3 Modified lattice

A very important property of 2D packings of disks is their tendency to crystallise into lattice-like structures [12]. Under high enough pressures, a monodisperse packing can crystallise into a hexagonal lattice, which is the optimal packing arrangement for circles in 2D. Crystallised packings do not exhibit the force chain structure that we observe in irregular packings. This is the main reason that Owens and Daniels used bidisperse particles — particles with different radii are much less likely to crystallise into regular structures.

Nevertheless, there is a definite tendency for bidisperse packings to locally approximate hexagonal packings. From visual inspection, we see that 2D packings include some regions that are reminiscent of hexagonal lattices (See Fig. 3). This motivates a modified lattice model, which treats granular networks as a perturbation of a hexagonal lattice. Starting with the contact network for a rectangular section (dimensions to match experiment) of a perfect close-packed lattice, I remove edges uniformly at random until the edge density matches experiment.

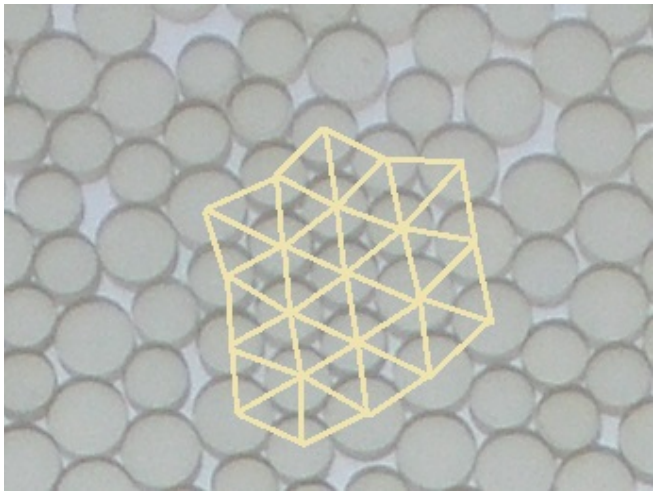

Figure 3: On a small scale, bidisperse particles can still pack into a regular hexagonal structure. Original image from Owens and Daniels [7] (used with permission).

Note that all the other models give a set of particle positions as their output. This set contains enough information to construct the contact network (once R is fixed). The modified lattice model is different in that it directly outputs the contact network. This network does not contain enough information to reconstruct a set of particle positions. This means that, while the other models can be visualised by plotting the positions of the particles (as in Fig. 2) the modified lattice model is not amenable to a visual representation. It is a model *only* of the contact network, and makes no attempt to model the precise physical arrangement of particles.

Despite this, the network diagnostics of this model match very well to those of the experimental data (see Table 1). The match is particularly good for diagnostics which characterise global, system scale properties of the network — specifically, node and edge betweenness centrality, mean shortest distance,

and global efficiency.

2.4 Force-modified RGG

An obvious constraint on granular materials is that they are formed of particles that obey physical laws. In particular, the particles are acted on by gravity (unless the plane is horizontal) and by forces from the other particles and the walls of the container. In Hertzian contact theory, the force between two compressed particles takes the form

$$f \propto \delta^{\beta}, \tag{5}$$

[7] where δ is the total compression and β is an exponent which depends on the geometry of the particles. For the particles in the experimental data I used, the exponent β was found to be approximately $5/4$ [7].

I modify the crude RGG model by calculating the inter-particle forces using force law (5) for a random arrangement of particles. I then allow each particle to move a displacement $\mathbf{d}_i = \varepsilon \mathbf{f}_i$, where $\mathbf{f}_i$ is the total force on particle i, and ε is a resolution parameter. I then recalculate the forces, and repeat the process.

There are subtleties involved in choosing an appropriate value for the parameter ε. Ideally ε would be taken to be arbitrarily small[4] so that the system could move smoothly and incrementally into its equilibrium arrangement. However, reducing ε dramatically increases the computation time.

A larger value of ε is useful at the beginning of the simulation, because the particles start out clustered together and need to move significant distances in order to fill the spaces in between. However, this misses the fine detail needed to home in on the equilibrium arrangement, and the particles end up hopping back and forth around their equilibrium position without ever getting there.

To address this issue, I monitor the potential energy of the system. Since the model is supposed to allow the particles to relax under the forces they exert on one another, the potential energy should decrease with every step of the simulation. The code was modified so that if after any step the potential energy becomes higher than it was after the previous step, the resolution parameter is tuned down by a constant factor. This allows me to start the simulations with a relatively high value of ε, which is then modified as the simulation progresses so that the system is always sensitive at the appropriate distance scales. The potential energy will eventually tend to an asymptotic value, and there will come a point at which small adjustments in particle positions are unimportant. One method I use to determine the moment at which this occurs is to calculate the adjacency matrix at every step. If the adjacency matrix remains the same for a significant number of iterations, I judge it appropriate to end the simulation. See Fig. 2 for a particle arrangement generated using this model. For more details on the implementation of the model, see Appendix A.

This model is the most sophisticated one that I investigated, and is the closest to real-world granular networks because it includes more physical constraints than the other models. Therefore, we expect it to perform particularly well under our network diagnostics. This is indeed what we find, especially compared with the simple RGG and $p(\rho)$-modified RGG models. The diagnostics for which the force-modified RGG performs best appear to be ones that are especially sensitive to local structure — clustering coefficient, local efficiency, modularity, subgraph centrality, transitivity, communicability, $\mathrm{tr}(e^A - I)$, and path-weighted betweenness (Table 1). The modified lattice model still seems to better capture the global properties of the networks.

For the diagnostics that are more sensitive to global properties, the force-modified RGG is actually not much better than the $p(\rho)$-modified RGG model. It is interesting that the improvement from the cruder

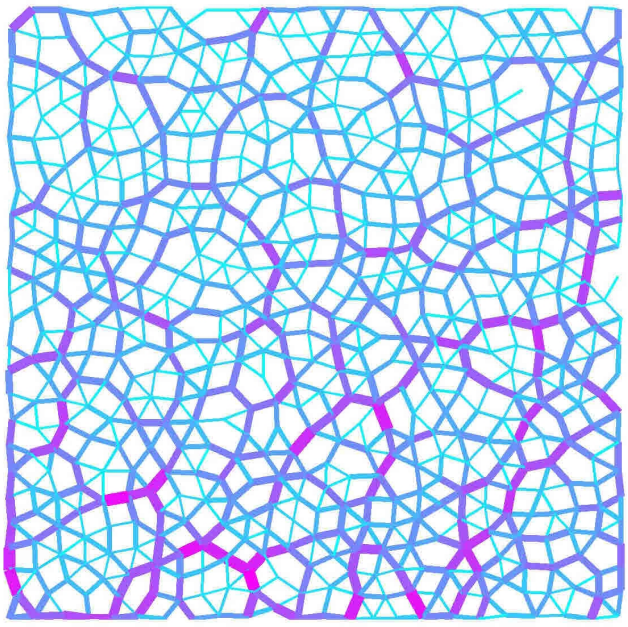

Figure 4: Force chains in the force-modified RGG model. Thickness of the lines indicates the strength of the inter-particle forces.

[4]Here, 'small' means that the mean distance moved by each particle is small compared to the dimensions of the container.

model to the more sophisticated one does not show up significantly in the global network diagnostics. The conclusion is striking: modelling the forces between particles causes a significant improvement in the model's local structure, but seems to have little effect on the global structure.

3 Physical analysis of the force-modified RGG

Thus far, I have focussed on using network diagnostics to assess the effectiveness of the models. Since the force-modified RGG model incorporates a force law, it makes sense to analyse this model from a physics perspective as well.

Figure 4 is a plot of the force chains for a particle arrangement generated by this model. I calculate the forces between particles using force law (5), with $\beta = 5/4$. Encouragingly, the force-modified RGG model reproduces force chains that are reminiscent of experiment (see Fig. 1).

To investigate these matters more quantitatively, I also examine the force distributions. In Fig. 5, I plot the force distribution as measured in experiment, compared to the force distribution for force-modified RGG packings generated at the same packing fraction. Note the presence of more than one peak in the experimental distribution. The force distribution generated by the force-modified RGG has only one distinct peak, so some aspect of the physics has clearly been overlooked by this model. To understand what the model is missing, we return to the force law (5). Including the constant of proportionality, we write it as

$$f_{ij} = \alpha \delta_{ij}^{\beta}, \tag{6}$$

where we recall that δ_{ij} is the total compression between particles i and j. In Hertzian contact theory, the constant α is generally a function of the particles' geometry. In the force-modified RGG model, the particles are displaced in increments proportional to the resolution parameter ε, so the constant α is unnecessary if we assume it takes the same value for all particles. I therefore did not incorporate it into the force-modified RGG model. However, we would actually expect α to be a function of the radii of the two particles. Therefore,

$$f_{ij} = \alpha(r_i, r_j) \delta_{ij}^{\beta}. \tag{7}$$

To generate a physically more realistic force distribution, we need to determine the three values: $\alpha(r_1, r_1)$, $\alpha(r_1, r_2)$, and $\alpha(r_2, r_2)$, where r_1 and r_2 are the two different radii of the bidisperse particles. We then obtain a force distribution that is a sum of three (single-peak) distributions, with one peak for each distinct value of α. This may be the mechanism behind the more complex force distribution that we observe in experiment.

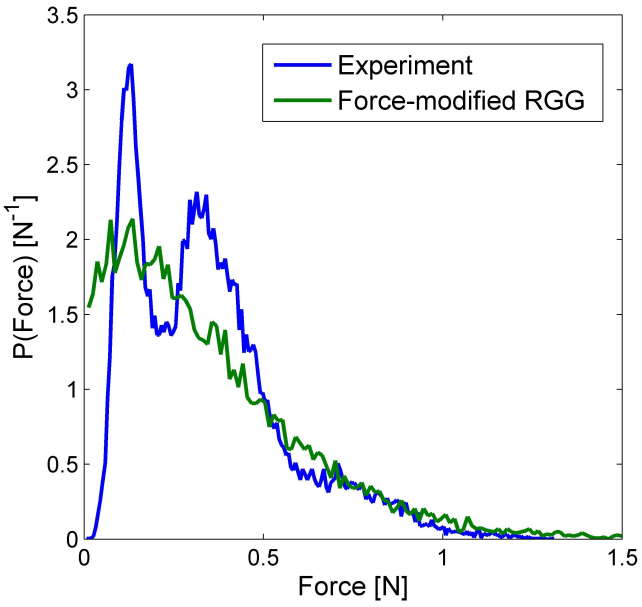

Figure 5: Normalised force distributions for experiment (averaged over 17 different packings) and the force-modified RGG (averaged over 20 realisations). The distribution from the model is scaled so that the mean force matches experiment.

In an attempt to improve the force-modified RGG model, one can try to determine the values of α by fitting to experiment. Indeed, including as much of the known physics as possible seems an obvious way of improving the model. However, as I shall discuss in Section 4, the precise force law used in the generative stage seems to have little effect on the model's effectiveness, particularly its effectiveness at modelling network properties.

Another relevant distribution is the distribution of compressions (i.e. the values of δ_{ij} in the above equations). The δ-distribution is simpler in that it does not depend on the different values of α, so we can use it instead of the force distribution as a means of comparing the model to experiment.

In Fig. 6, we compare the δ-distributions from experiment and from the force-modified RGG. I used sets of packings which were generated at the same packing fraction as the experiments. Interestingly, the peak of the experimental distribution is at a significantly higher compression. For non-identical

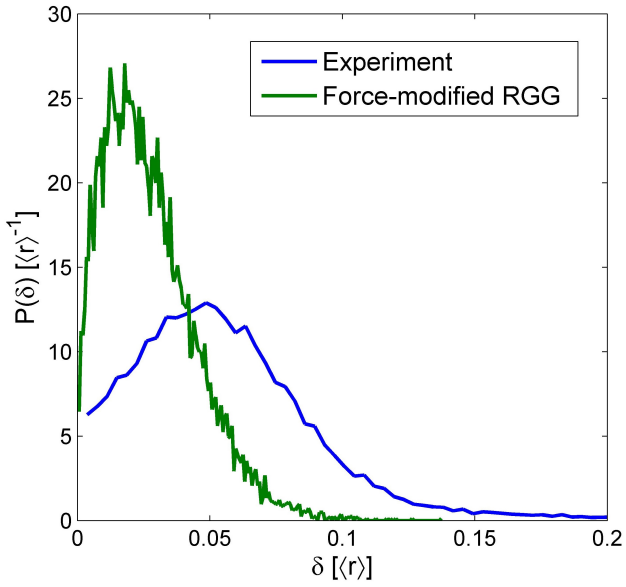

Figure 6: Normalised δ-distributions for the force-modified RGG (averaged over 20 realisations) compared to experiment (averaged over 17 different packings), at the same packing fraction. I give values of δ in units of $\langle r \rangle$, the mean particle radius.

packing fractions, we might expect to find a higher peak compression in one than the other. However, since both are at the same packing fraction, this is an undesirable result. A lower peak compression suggests that the force-modified RGG model generates packings with too low a potential energy.

To investigate further, I calculate the mean potential energy of the packings generated by the force-modified RGG model and compare to experiment. The potential energy in both cases can be calculated from the particles' positions. In this calculation I assume all values of α to be the same and equal to 1, so the values are not strictly accurate and are useful only for comparison[5]. I obtain the mean value 8.88×10^{-6} (with a standard deviation of 0.99×10^{-6}) for the experimental packings and 1.74×10^{-6} (standard deviation 0.47×10^{-6}) for the force-modified RGG packings.

4 Sensitivity to the force law

We have found that including a model of the forces between particles significantly improves the RGG's effectiveness at describing granular network structure (especially locally). It is important to ask whether the model's effectiveness relies on using the correct force law. As we saw in Section 3, the force law could have been made more realistic by choosing correct values for the coefficients $\alpha(r_i, r_j)$, instead of assuming that they were all equal. Would the model perform better with such an improvement?

To investigate how sensitive the network properties of the model are to the force law, I generated particle arrangements for different values of the parameter β. As I have mentioned, for the particles used in the experiments, $\beta \approx 5/4$. In Table 2 in Appendix B, I show the results for the various network diagnostics for $\beta = 5/4$ and $\beta = 0$. The latter is an extreme case, in which the force between a pair of particles is independent of their compression and particles all exert the same force on one another if they are in contact.

Interestingly, the results for the two cases are nearly identical. All differences are well within one standard deviation. In the force-modified RGG we might as well have used a force law that is simpler both analytically and computationally, as it would still generate networks that are good models of real granular networks (in so far as the chosen diagnostics can reveal).

This suggests that the important constraint on granular networks is not the exact form of the force law between particles but rather the fact that they exert *some* force on each other. The network properties that I studied do not seem to be sensitive to the forces between particles — however, it *is* important that some repulsive force is modelled.

5 Extension to 3D

The simplicity of the models that I have developed allows a straightforward extension to three-dimensional granular materials. Without experimental data in 3D, the best we can do is to compare the results from the models. In Table 3 in Appendix B, I present the network diagnostics for the force-modified RGG versus the modified lattice model. I also give results from a 3D version of the simple RGG for comparison. I use the same number of particles in each model and match the edge density as before. For the force-modified RGG and simple RGG, I use a container whose sides are all of equal length; for the modified lattice, I start with an approximately cubic section of a 3D hexagonal close-packed lattice.

We no longer observe the particularly good agree-

[5]The explicit expression I use to calculate the potential between particles i and j is $V_{ij} = \delta_{ij}^{\beta+1}$. This is strictly only *proportional* to the potential, but since I use it only for comparison, this technicality is unimportant.

ment between the force-modified RGG and modified lattice that we saw in 2D. This is especially true of the more local diagnostics, such as clustering coefficient and local efficiency, for which the disparity between the two models has significantly increased.

The force-modified RGG imposes the most realistic physical constraints — this suggests that the results from this model will be the most reliable and in best agreement with real 3D granular materials. Experimental data is necessary to further investigate this possibility.

6 Discussion and conclusions

I have investigated three novel models of granular networks. The main analysis I performed on these models was a comparison of network diagnostics for the unweighted contact networks. The most sophisticated of these models, the force-modified RGG, provides a model not only of the contact network but also of the weighted force network. I compared this model's prediction for the distributions of forces and compressions to experiment. Then I extended the force-modified RGG and the modified lattice models to three dimensions and repeated the analysis of their network diagnostics, comparing their behaviour in 3D and 2D.

The network analysis of the RGG-based models was interesting because it illustrated how including more physical constraints in a model improved its performance in comparison to experiment. The simple RGG model included no constraints other than the fact that it embedded the network in a two-dimensional space. The $p(\rho)$-modified RGG model attempted to crudely model the constraint that particle compressions are small compared to their radii by preventing particles from being closer than 0.8 times the sum of their radii. Finally, the force-modified RGG model attempted to simulate the actual forces between particles. The performance of the models improved with each constraint — the simple RGG performed worst, followed by the $p(\rho)$-modified RGG, while the force-modified RGG was the closest match to experiment (see Table 1).

An enlightening distinction that a network analysis allows us to make is the distinction between local and global properties of the system. Each network diagnostic is sensitive at a characteristic size scale. For example, betweenness centrality is related to the number of geodesic paths that go through a given node/edge, so its value depends on the structure of the rest of the network. In this sense its value for each node/edge is sensitive to global properties of the network. The main reason I introduced *path-weighted betweenness* was to create a related measure that was more sensitive to local properties. As described in Section 1.5, the contribution of each path to this diagnostic is down-weighted by a factor $1/l!$, where l is the length of the path. Consequently, this diagnostic is much less sensitive to large-scale structure, and the leading contribution comes from short paths in the vicinity of each node. A related diagnostic is $\mathrm{tr}(e^A - I)$, which counts the number of loops that start and end at a given node, and in a similar way down-weights the contributions from longer loops.

As was discussed in Section 2, the improvement from the $p(\rho)$-modified RGG to the force-modified RGG mainly shows up in the diagnostics which are sensitive to local structure. As far as the global diagnostics are concerned, the force-modified RGG is only a slight improvement on the $p(\rho)$-modified RGG. The natural conclusion to draw is that modelling in detail the forces between particles significantly improves the local structure of a model, while accurate global properties can be modelled by much simpler constraints, like the one the $p(\rho)$-modified RGG is based upon. Since it is always desirable to have a model that gives accurate predictions with minimal physical constraints, it is valuable to be able to identify which constraints different properties of granular networks depend on.

An especially interesting result was found during the investigations described in Section 4. Although modelling the forces between particles improves the local structure of the RGG model, none of the diagnostics I used appear to be affected by the precise form of the force law. Whether or not there are any diagnostics of the unweighted network that are sensitive to the choice of force law is a question that further investigations might shed light on.

The physical analysis of the force-modified RGG in Section 3 brought to light one issue with this model. It generates packings which appear have a lower potential energy than the packings studied in experiment. One method of overcoming this problem might be to stop the simulation when the potential energy reaches a certain value. This would introduce an extra parameter to the model, but might generate packings which more closely resemble real granular materials. Repeating both the network and physical analyses of Sections 2 and 3 could support this suggestion.

To conclude, diagnostics from network science have

made it possible to perform a much more detailed analysis of granular materials — and models thereof — than a purely physical approach would have allowed. I have been able to examine the effectiveness of different models at different size scales, using a set of diagnostics that are sensitive to a wide range of system properties. The diagnostics I have used are, however, far from a complete set, and more work could be done examining how granular networks and their models perform under different network diagnostics.

As discussed in Section 2, the modified lattice model has network diagnostics that match very well to experiment (especially on global properties). Whether or not there are network diagnostics for which the modified lattice shows up as a poor model is a question that might be answered by further investigation. The modified lattice does not attempt to model the physics of granular networks to the extent that the force-modified RGG does, and perhaps this is showing up in the fact that the force-modified RGG outperforms the modified lattice on local diagnostics (in particular communicability and $\mathrm{tr}(e^A - I)$). There may however be other diagnostics for which this deficiency shows up more significantly.

Network diagnostics appear to be useful and highly sensitive tools for understanding the structure of granular materials. Further work might uncover more appropriate network diagnostics for application to this area, as well as lead to a better understanding of their dependence on the physical properties of granular materials.

Acknowledgements

I would like to thank Karen Daniels for providing experimental data and insightful comments, and Lisa Manning for providing MATLAB code and data from simulations. I would especially like to thank my supervisor Mason Porter for many useful comments and invaluable discussions.

References

[1] K. Hutter and K. R. Rajagopal. *On flows of granular materials*. Continuum Mech. Thermodyn. 6 819 (1994).

[2] P. Jop, Y. Forterre, and O. Pouliquen. *A constitutive law for dense granular flows*. Nature Lett. 441 727 (2006).

[3] M. van Hecke. *Jamming of soft particles: geometry, mechanics, scaling and isostaticity*. J. Phys.: Condens. Matter 22 033101 (2010).

[4] H. A. Makse, N. Gland, D. L. Johnson, and L. M. Schwartz. *Why Effective Medium Theory Fails in Granular Materials*. Phys. Rev. Lett. 83 5070 (1999).

[5] D. M. Mueth, H. M. Jaeger, and S. R. Nagel. *Force Distribution in a Granular Medium*. Phys. Rev. E 57 3164 (1998).

[6] D. S. Bassett, E. T. Owens, K. E. Daniels, and M. A. Porter. *Influence of network topology on sound propagation in granular materials*. Phys. Rev. E 86 041306 (2012).

[7] E. T. Owens and K. E. Daniels. *Sound propagation and force chains in granular materials*. Europhys. Lett. 94 54005 (2011).

[8] M. Barthélemy. *Spatial Networks*. Physics Reports 499 1 (2011).

[9] M. E. J. Newman. Networks: An Introduction. (Oxford University Press, Oxford, 2010).

[10] M. Rubinov and O. Sporns. *Complex network measures of brain connectivity: Uses and interpretations*. NeuroImage 52, 1059 (2009).

[11] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre. *Fast unfolding of communities in large networks*. J. Stat. Mech. (2008) P10008

[12] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. *Jamming in Hard Sphere and Disk Packings*. J. Appl. Phys. 95 (3) 989 (2004).

Appendices

A Force-modified RGG algorithm

The equation I used for the force acting on particle i is

$$\begin{aligned}\mathbf{f}_i = \sum_{j\neq i} \Big[\big(\frac{1}{2}(r_i + r_j) - \frac{1}{2}|\mathbf{x}_i - \mathbf{x}_j|\big)^\beta \frac{\mathbf{x}_i - \mathbf{x}_j}{|\mathbf{x}_i - \mathbf{x}_j|} \Big] \\ +(r_i - x_i)^\beta \hat{\mathbf{x}} - (r_i + x_i - L_x)^\beta \hat{\mathbf{x}} \\ +(r_i - y_i)^\beta \hat{\mathbf{y}} - (r_i + y_i - L_y)^\beta \hat{\mathbf{y}}\end{aligned} \tag{8}$$

where it is understood that the terms raised to the power β are only to be evaluated if they are positive.

In this equation, r_i is the radius of the i^{th} particle and x_i and y_i are its coordinates, being the components of the vector $\mathbf{x}_i$. Two walls of the container lie along the x and y axes, while the other two lie along the lines $x = L_x$ and $y = L_y$. The unit vectors along the x and y axes are $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, respectively. The terms in the sum give the forces from other particles, while the final four terms are the forces from each of the walls.

For the realisations analysed in this report, I used $L_x = 0.29$ and $L_y = 0.38$. As described in Section 2.4, the particles — whose positions are initially distributed uniformly at random — are moved a distance $\mathbf{d}_i = \varepsilon \mathbf{f}_i$ in every step of the simulation. A different starting value of ε is appropriate for different values of β. I found that, for $\beta = 1.25$ and $\beta = 0$, good starting values are $\varepsilon = 8$ and $\varepsilon = 0.001$ respectively. Higher values often caused the potential energy of the system to quickly diverge, while lower values made the simulation unnecessarily slow. At every step of the simulation the potential energy of the system was calculated. The formula used for this was

$$\begin{aligned} V = \sum_{j \neq i} \Big[\big(\frac{1}{2}(r_i + r_j) - \frac{1}{2}|\mathbf{x}_i - \mathbf{x}_j|\big)^{\beta+1} \Big] \\ +(r_i - x_i)^{\beta+1} - (r_i + x_i - L_x)^{\beta+1} \\ +(r_i - y_i)^{\beta+1} - (r_i + y_i - L_y)^{\beta+1}. \end{aligned} \tag{9}$$

Let V_i be the potential of the system after the i^{th} iteration. If $V_i > V_{i-1}$, then we replace $\varepsilon \to \varepsilon' = 0.9\varepsilon$. This ensures that the rearrangement of the particles becomes more and more precise as the system gets closer to an equilibrium arrangement.

B Network diagnostics

Table 1: Network diagnostics for the experimental networks and models. I take means and standard deviations over 17 experimental networks and 20 realisations of each of the models.

Diagnostic	Experimental networks		Force-modified RGG		Modified lattice	
	Mean	SD	Mean	SD	Mean	SD
Edge density	2.101	0.0466	2.1012	0.0201	2.0972	0.0227
Node betweenness	20220	200	21275	254	20429	88
Edge betweenness	19.0082	0.1335	19.6592	0.2310	18.9846	0.1128
Clustering coefficient	0.2582	0.0108	0.2315	0.0071	0.2963	0.0084
Transitivity	0.2688	0.0078	0.2566	0.0058	0.2932	0.0062
Assortativity	0.1384	0.0218	0.2490	0.0209	0.0632	0.0270
Global efficiency	0.0763	0.0008	0.0754	0.0007	0.0769	0.0004
Local efficiency	0.3242	0.0170	0.2909	0.0111	0.3755	0.0108
Mean shortest distance	19.1953	0.1335	19.7059	4.7942	19.0068	0.0622
Maximised modularity	0.8473	0.0027	0.8449	0.0021	0.8543	0.0029
Subgraph centrality	6.8551	0.6044	7.4001	0.1720	7.8492	0.2474
Communicability	0.0690	0.0055	0.0695	0.0026	0.0727	0.0033
$\mathrm{tr}(e^A - I)$	7163.4	460.9	7296.1	196.1	7760.6	168.3
Path-weighted betweenness	28.7262	0.6722	28.9275	0.2751	28.9261	0.3346
Diagnostic	$p(\rho)$-modified RGG		Monodisperse RGG		Bidisperse RGG	
	Mean	SD	Mean	SD	Mean	SD
Edge density	2.1015	0.0181	2.1086	0.0398	2.1053	0.0550
Node betweenness	21495	208	4290.7	4005.6	6096.4	6746.8
Edge betweenness	19.8526	0.1838	3.9432	3.6016	5.5745	6.0640
Clustering coefficient	0.3265	0.0065	0.5503	0.0088	0.5443	0.0108
Transitivity	0.3326	0.0039	0.5915	0.0079	0.5887	0.0119
Assortativity	0.2212	0.0207	0.5907	0.0400	0.5735	0.0537
Global efficiency	0.0742	0.0006	0.0190	0.0032	0.0209	0.0050
Local efficiency	0.4284	0.0099	0.6332	0.0088	0.6261	0.0136
Mean shortest distance	19.8804	0.1706	19.4228	5.8312	20.0990	7.5789
Maximised modularity	0.8684	0.0037	0.9546	0.0021	0.9541	0.0036
Subgraph centrality	8.1630	0.1826	35.5901	22.0596	40.7119	34.3601
Communicability	0.0738	0.0031	0.4103	0.3469	0.5033	0.5832
$\mathrm{tr}(e^A - I)$	8165.8	208.2	39433	25148	45272	39171
Path-weighted betweenness	27.4245	0.3510	19.1524	0.5257	19.3028	0.6974

Table 2: Comparison of network diagnostics for the force-modified RGG model with different values of β (averaged over 20 realisations in each case).

Diagnostic	$\beta = 5/4$		$\beta = 0$	
	Mean	SD	Mean	SD
Edge density	2.1012	0.0201	2.1010	0.0223
Node betweenness	21275	254	21020	259
Edge betweenness	19.6592	0.2310	19.4323	0.2319
Clustering coefficient	0.2315	0.0071	0.2315	0.0076
Transitivity	0.2566	0.0058	0.2586	0.0074
Assortativity	0.2490	0.0209	0.2684	0.0574
Global efficiency	0.0754	0.0007	0.0757	0.0007
Local efficiency	0.2909	0.0111	0.2943	0.0115
Mean shortest distance	19.7059	4.7942	19.5386	0.1733
Maximised modularity	0.8449	0.0021	0.8440	0.0031
Subgraph centrality	7.4001	0.1720	7.5207	0.2227
Communicability	0.0695	0.0026	0.0720	0.0039
$\mathrm{tr}(e^A - I)$	7296.1	196.1	7433.6	253.9
Path-weighted betweenness	28.9275	0.2751	29.0069	0.3404

Table 3: Network diagnostics for the 3D versions of the models (averaged over 20 realisations in each case).

Diagnostic	Modified lattice		Force-modified RGG		RGG	
	Mean	SD	Mean	SD	Mean	SD
Edge density	3.7432	0.0313	3.7385	0.0493	3.7310	0.0905
Node betweenness	5496.1	13.6	6253.4	27.4	7298.9	233.1
Edge betweenness	6.4980	0.0134	7.2524	0.0274	8.2875	0.2354
Clustering coefficient	0.1706	0.0033	0.2900	0.0056	0.5077	0.0063
Transitivity	0.1524	0.0030	0.2809	0.0044	0.4854	0.0062
Assortativity	0.1743	0.0188	0.3001	0.0242	0.4319	0.0447
Global efficiency	0.1879	0.0005	0.1685	0.0006	0.1462	0.0032
Local efficiency	0.2600	0.0086	0.4838	0.0126	0.6643	0.0051
Mean shortest distance	6.5046	0.0100	7.2524	0.0274	8.3740	0.2247
Maximised modularity	0.6653	0.0048	0.7061	0.0035	0.8392	0.0081
Subgraph centrality	55.5779	1.8141	60.7041	4.4417	1017.20	1584.2
Communicability	5.0862	0.4120	3.5704	0.3483	40.3959	63.7762
$\mathrm{tr}(e^A - I)$	54843	2324	59704	4442	1016100	1584000
Path-weighted betweenness	105.4764	0.6966	80.5972	0.8974	59.3594	2.0521