



Cite this: *Soft Matter*, 2025, 21, 4839

Sampling the large-dimensional energy landscape of a 2D granular system with the hydra string method

Katherine A. Newhall 

In this work, I improve upon the existing hydra string method [C. Moakler and K. A. Newhall, *Granular Matter*, 2021, **24**, 24] to systematically sample the energy landscape of a low friction 2D granular system. This method climbs in random directions out of a minimum energy state, finding unique saddle transition points and the neighboring minimum energy states only to repeat the process from the newly found minima. The data is saved as a network with nodes representing the energy-minimizing states and edges representing transition pathways that are parallel to the gradient of the energy at each point along the path. I show how the hydra string method is able to produce a better sample of transition pathways between stable states compared to just randomly sampling the system. The method is also modified to take into account energy minima that are not points caused by non-mechanically stable individual particles and skip past entire configurations that are not mechanically stable. The samples reveal that the energy of the states correlates with the size of the energy barriers between them. Neighboring state energies are also correlated, with correlations decreasing with distance as measured by path length on the network.

Received 11th November 2024,
Accepted 14th March 2025

DOI: 10.1039/d4sm01337a

rsc.li/soft-matter-journal

1 Introduction

Many intriguing dynamical properties of systems, such as metastability or resistance to applied forces, emerge from an underlying energy landscape.¹ For systems that dissipate energy, like gradient or Langevin models, stable configurations are local energy minima. Applied forces including thermal fluctuations can push the system from one energy-minimizing state to another. Traditional applications of energy landscapes include finding transition paths between two known stable configurations, such as in chemical reactions,² or finding a global energy-minimizing state, such as in protein folding.³ Much like complex molecules, granular matter consists of pair-wise interactions between grains, yet energy-based methods have not been tried, most notably due to the influence of static friction. Recent work suggests a bridge between frictional jammed packings and frictionless hard-sphere glassy systems, with friction stabilizing higher energy configurations.⁴

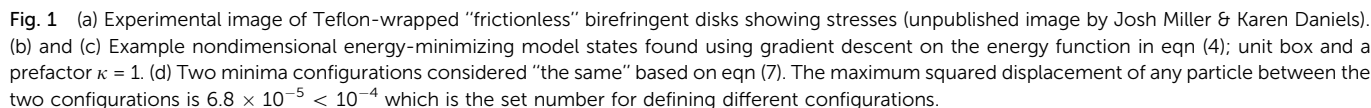
Glassy systems are a low friction environment in which energy landscapes have been explored. Their fractal nature may explain soft material behavior, such as power-law rheology and non-diffusive bubble motion in the Ostwald ripening of bubbles.⁵ (Although their fractal nature is under question with different sampling methods giving different results.⁶)

Landscapes may also prove useful in other relatively slow dynamics such as quasistatically sheared glassy systems⁷ or the slow loading and unloading of frictionless soft particles.⁸ Unlike the protein energy landscape that has a funnel-like structure and one or a few energy-minimizing states of interest, glassy systems are more uniform with deep canyons and roughness corresponding to large numbers of plausible configurations within these basins.^{9–11} Combined with the fact that enumeration of all states is possible for only small systems,^{12,13} methods to sample and explore the energy landscape are desirable over methods to find global minimizers or completely map the entire landscape. For this reason, together with Chris Moakler, I developed the hydra string method to sample a localized region of an energy landscape, creating a network of neighboring energy-minimizing states.¹⁴

The hydra string method is based on the string-method¹⁵ and its saddle-point-finding climbing-string variant.¹⁶ These efficient methods find transition pathways for gradient systems that are parallel to the gradient of the energy at each point along the path. The hydra string method sends out climbing strings, finding a set of unique saddle points on the edge of the basin of attraction of the starting minima. These Morse-index one saddle points have only one unstable direction and thus separate two energy-minimizing configurations. The network of transition paths between these configurations is recorded, and the process repeats from the newly found minima.

Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, USA. E-mail: knewhall@unc.edu





After detailing the model for the soft-sphere particles in Section 2, the original hydra string method, together with the sampling methods explored in this paper and the modifications to address the issues above are presented in Section 3. Rattler particles are ignored when determining the unique set of configurations visited, and unjammed configurations are removed from the list of future minima to explore. I compare sampling methods and parameters in Section 4. Forming a local network of neighboring minima has the advantage of finding saddle points on the edge of the basin of attraction from both paths that exit the basin as well as enter the basin. I argue this forms a more representative sample than just exiting the basin alone. I also compare different computational parameters to argue that a ‘good enough’ sample has been obtained. I also discuss features

$$\vec{f}_{ij} = -\nabla E_{ij} = -\kappa \delta_{ij}^{5/4} \frac{\vec{x}_i - \vec{x}_j}{|\vec{x}_i - \vec{x}_j|} \quad (2)$$




Fig. 2 For a set of minima found by randomly choosing initial positions for each sphere and then gradient descending to tolerance, a histogram of energies for those determined to be 'jammed' vs. 'floppy'. Jammed configurations are determined by first removing 'rattlers' (particles with 2 or fewer contacts) and then requiring the eigenvalues of the Hessian matrix to all be non-zero, within a tolerance of 1×10^{-12} . Changing the zero tolerance to 1×10^{-5} made no difference in the histogram separation of the states. Note the histogram bars appear as slightly different colors when they overlap.

For disks that overlap one of the four walls of the box, I retain the linear force, treating each wall as a cylinder with infinite radius. The top and bottom walls produce forces in the y -direction and the side walls produce force in the x -direction. The resulting forces on particle i are

$$\begin{aligned} f_i^1 &= \kappa(R_i - x_i) & \text{if } x_i < R_i \\ f_i^2 &= -\kappa(R_i - (1 - y_i)) & \text{if } 1 - y_i < R_i \\ f_i^3 &= -\kappa(R_i - (1 - x_i)) & \text{if } 1 - x_i < R_i \\ f_i^4 &= \kappa(R_i - y_i) & \text{if } y_i < R_i \end{aligned} \quad (3)$$

for the left, top, right, and bottom walls, respectively. Together, the total energy of the system is

$$E = \sum_{i=1}^N \left[\sum_{j=1}^N C_{ij} E_{ij} - \sum_{k=1}^4 \frac{(f_i^k)^2}{2\kappa} \right] \quad (4)$$

where the contact matrix $C_{ij} = 1$ if particles i and j overlap, and zero otherwise, and the net force acting on particle i is given by

$$\begin{aligned} F_i^x &= \sum_{j \neq i} C_{ij} f_{ij}^x + f_i^1 + f_i^3 \\ F_i^y &= \sum_{j \neq i} C_{ij} f_{ij}^y + f_i^2 + f_i^4 \end{aligned} \quad (5)$$

for $i, j = 1 \dots N$.

Experimentally observed static configurations of frictionless particles correspond to energy-minimizing states of the model system. In order to create an ensemble of such configurations, the positions of the particles are first randomly chosen inside the unit box, and then updated by gradient descent dynamics,

numerically integrating the equations

$$\frac{d}{dt} \vec{x}_i = -\nabla E = (F_i^x, F_i^y) \quad (6)$$

with time step $\Delta t = 0.1$, where the energy E is given in eqn (4) and net forces on particle i , F_i^x and F_i^y , by eqn (5). The numerical integration was stopped when the absolute values of F_i^x and F_i^y were less than 10^{-7} for all $i = 1 \dots N$ particles. Fig. 1(b) and (c) show two energy-minimizing configurations of the model system with substantially different energies. Due to this large range in possible energies, throughout the remainder of the paper, the log base 10 of the energies will always be presented.

The substantially different energies in fact come from a bimodal distribution of energies across an ensemble of minima shown in Fig. 2 explained by the existence of both mechanically stable jammed states and effectively zero energy unjammed or floppy states. Numerically, mechanically stable states are identified by first recursively removing rattler particles with fewer than three contacts before finding the eigenvalues of the Hessian matrix for the remaining particles. If these eigenvalues are larger than 1×10^{-12} , (no zero eigenvalues) the configuration is labeled jammed. If no particles remain after recursively removing rattlers, the configuration is labeled floppy.

While the two configurations shown in Fig. 1(b) and (c) are clearly distinct, the two overlaid configurations in Fig. 1(d) do not show substantial differences in their contacts or overall arrangement. In general a tolerance must be set to consider numerically-found states to be distinct. Here, I take configurations to be distinct if the Euclidean distance between each numbered particle in the two configurations are all less than 10^{-2} . Specifically, configuration (\vec{x}^A, \vec{y}^A) and (\vec{x}^B, \vec{y}^B) are unique if

$$(\vec{x}_i^A - \vec{x}_i^B)^2 + (\vec{y}_i^A - \vec{y}_i^B)^2 > 10^{-4} \quad \text{for all } i = 1 \dots N. \quad (7)$$

Note that swapping two particles of the same size does result in a different configuration with this definition. The tolerance of 10^{-2} was chosen because it is about 10% of a particle radius.

3 Hydra string method

I review the string method¹⁵ and its climbing variant¹⁶ for finding the minimum energy path (MEP) between two minimizers of a gradient system. These are the fundamental pieces of the hydra string method¹⁴ for exploring complex energy landscapes. This algorithm is also summarized, and the modifications are discussed to account for rattler particles and floppy states.

To find the MEP, the string method uses steepest descent dynamics to evolve a curve $\phi(\alpha, t)$ parameterized by $\alpha \in [0, 1]$, the normalized arc length. For the model granular system of Section 2, $\phi(\alpha, t)$ is a $2N$ -dimensional vector of the particles' positions. The two ends of the string, $\phi(0, t)$ and $\phi(1, t)$, converge to two energy-minimizing states or can be fixed as these points from the start. As a numerical method, finding the MEP amounts to discretizing the string into a sequence of N_s images of the system, $\{\phi_1^k, \dots, \phi_{N_s}^k\}$ where $\phi_j^k \approx \phi((j-1)\Delta\alpha, k\Delta t)$.



Here, $\Delta\alpha = 1/(N_s - 1)$ and Δt is a chosen time step size. For each time step, the images are updated independently according to a gradient descent given by

$$\phi_j^{k+1} = \phi_j^k - \nabla V(\phi_j^k) \Delta t \text{ for } j = 1 \dots N_s. \quad (8)$$

For the model granular system, $-\nabla V$ is the $2N$ vector of forces F_i^x and F_i^y given by eqn (5) for $i = 1 \dots N$. Then, the images are redistributed along the string using interpolation to keep them equally spaced in arc length. Note the output path is dependent on the initial condition of the string in the case that there is more than one MEP in the energy landscape.

The string method is augmented to be a saddle-point search method by allowing one end of the string to climb uphill.¹⁶ Specifically, the uphill dynamics of the image $\phi(1, t)$ are given by

$$\partial_t \phi(1, t) = -\nabla V(\phi) + \nu (\nabla V(\phi), \hat{\tau}) \hat{\tau} \quad (9)$$

where (\cdot, \cdot) denotes the inner product, $\hat{\tau} = \partial_\alpha \phi / |\partial_\alpha \phi|$ is the unit tangent vector to the string, and the “climbing speed” ν is a parameter chosen to be greater than 1. Thus, this last image climbs uphill on the energy surface in the direction tangent to the string, but continues to perform steepest descent in the directions perpendicular to the string. It therefore converges to a Morse-index-one saddle point, *i.e.* one unstable direction.

In implementing this method, for each time step the last image $\phi_{N_s}^k$ is updated by

$$\phi_{N_s}^{k+1} = \phi_{N_s}^k - \nabla V(\phi_{N_s}^k) + \nu (\nabla V(\phi_{N_s}^k), \hat{\tau}^k) \hat{\tau}^k \quad (10)$$

where the unit tangent vector $\hat{\tau}$ is approximated by

$\hat{\tau}^k = (\phi_{N_s}^k - \phi_{N_s-1}^k) / |\phi_{N_s}^k - \phi_{N_s-1}^k|$ and the remaining images for $j = 1 \dots N_s - 1$ are updated according to eqn (8). The images are redistributed after every iteration to keep them equally spaced in arc length. To ensure that the found saddle point is on the edge of the basin of attraction of the starting minimum, the energy along the string is also computed at each time-step, and the string is truncated if the energy is found to be non-monotonic. Note that while faster-converging methods than gradient descent exist for finding critical points, the string along the minimum energy path is found by gradient descent so not much computational speed up is expected by switching to a faster-converging method once the pathway has converged.

Two methods for sampling the energy landscape are explored. Both rely on exploring the basin of attraction around a single minimum with climbing strings searching for saddle points. This process is described next. The two methods differ by whether the exploring strings are sent out from distant randomly-found energy minimizing states, or using the hydra string method¹⁴ to obtain a more accurate and complete sample by locally sampling the network of pathways connecting nearby energy-minimizing states.

The method for finding saddle points on the edge of the basin of attraction of a single minimum is as follows. From an energy-minimizing state, multiple climbing strings with $N_s = 30$ images are initialized by linearly interpolating between the energy minimizing state, and the state in which a particle is chosen at random and perturbed an amount 0.04 in a random




Fig. 3 Example transition paths found with the climbing string starting from the same energy-minimizing state (solid black circles) followed by continuation past the saddle point (dashed blue circles) to the energy-minimizing state on the other side of the barrier (dotted green circles). Underneath is the energy along the string indexed by image number, which are linearly spaced in arc length along the string. Notice the scale of the y-axis: (a) low energy barrier with minimal particle movement (b) medium energy barrier (c) high energy barrier with significant rearrangement.



direction. This method of perturbing a single particle is chosen as it aligns with the idea that saddle points in a granular system are typically local rearrangements of particles. Eqn (8) and (10) are iterated a maximum of 5×10^5 times with a step size of $\Delta t = 0.1$. Iterations are stopped when the maximum displacement of the particles in the both the x and y directions are less than 1×10^{-7} . Since each climbing string is independent, this process is done in parallel, collecting only the climbing strings that converged within the maximum number of iterations.

Collecting all the climbing strings that converged to a saddle point, only a unique list of saddle points is kept by ensuring the Euclidean distance between each numbered particle in the two configurations are all less than 10^{-2} , as in eqn (7). The full pathway to the energy minimizing state on the other side of the saddle point is found by extending the climbing string past the saddle point in the direction tangent to the string, $\hat{\tau}$, and then following the dynamics of eqn (8) for each image along the string. A maximum of 1×10^5 iterations with step size $\Delta t = 0.1$ are taken, stopping when the maximum displacement of the particles in both the x and y directions are less than 1×10^{-7} . To avoid losing resolution, the number of images along the string is doubled for this part. As each saddle must be compared to a running list of unique saddles, and the gradient descent dynamics are found to converge faster than the climbing string, this part is done in serial.

Fig. 3 depicts three different transitions found using the above-described method from the same minimum (black solid line) to the saddle point (blue dashed line) to the new energy-minimizing state (green dotted line). The evolution of the energy along each path is shown underneath. Note the change in scale of the y -axis; each of these transition paths starts from the same minimum, therefore at the same (non-zero) energy as computed with eqn (4).

There are multiple ways to quantify the distance between states in this system. The most natural mathematical distance would be Euclidean distance in N^2 space, but this lacks physical information about how far individual particles move. One could look at the maximum distance any particle is displaced, or the summed distances all the particles are displaced. However, the underlying assumption of an energy-landscaped-based description of the system is that the energy barrier, not the physical distance, between states impedes travel. Therefore in this work, I find it natural to discuss nearby states based on the number of barriers that must be crossed, which is path length in the network representation.

The first method for exploring the energy landscape consists of repeating the above-described procedure of finding transition paths for multiple randomly found energy-minimizing states that are found completely independently as described at the end of Section 2. The second method for exploring the energy landscape, as we will see obtains a more accurate and complete sample by locally sampling the network of pathways connecting nearby energy-minimizing states. The hydra string method¹⁴ starts as described above, finding unique saddle points around an energy minimum. It then repeats the process not from independent minima, but from the newly-found




Fig. 4 The hydra string script is run to explore 150 minima total (taking roughly 4.3 days on 32 cores). The resulting graph shows distinct minima as nodes with connecting edges representing energy barriers found. The starting-node for the algorithm is the middle node, with found connected minima in the next ring, and so on. Connections to minima not explored were removed for better visual clarity. The minima are colored by the log base 10 value of their energy indicated by the colorbar.

minima, keeping a running list of the unique minima and unique saddle points found, as well as the network of MEPs connecting these points. Each MEP network edge connecting two minima contains exactly one saddle on it since only Morse-index one saddle points are found. Thus a growing network of transition paths is found spreading out from the starting minimum.

The resulting network from the hydra string method is depicted in Fig. 4 when 150 minima are explored. Each node is an explored minimum and each edge is a transition path through a unique saddle; multiple edges between nodes indicate multiple unique saddles along different transition paths. Starting from the center node, the next ring of nodes represents the uniquely found minima under the first iteration of the algorithm. The third ring of nodes are new minima found when each of the second-ring nodes are explored. Note that a few self loops are present, and may be a result of how “the same” is determined in the algorithm.

The hydra string method while developed for a model granular system, is applicable to any gradient system. Here, I also consider a modification to the hydra string method specific to granular systems. One feature of granular systems is the existence of “rattler” particles that have fewer than 3 contacts and thus have lateral translational degrees of freedom leading to degenerate minimum states. If the rattler particle(s) are in different locations, the hydra string method would consider these different states, while really they should be considered the same. Additionally, if all particles have fewer than 3 contacts, then the entire system is in a zero-energy state and is considered “unjammed”. In both cases, the hydra string method would spend a large time exploring and enumerating essentially the same state over and over again. For example, the large number of low-energy minima connected to the low-energy (blue) minimum in the second ring in Fig. 4 could be



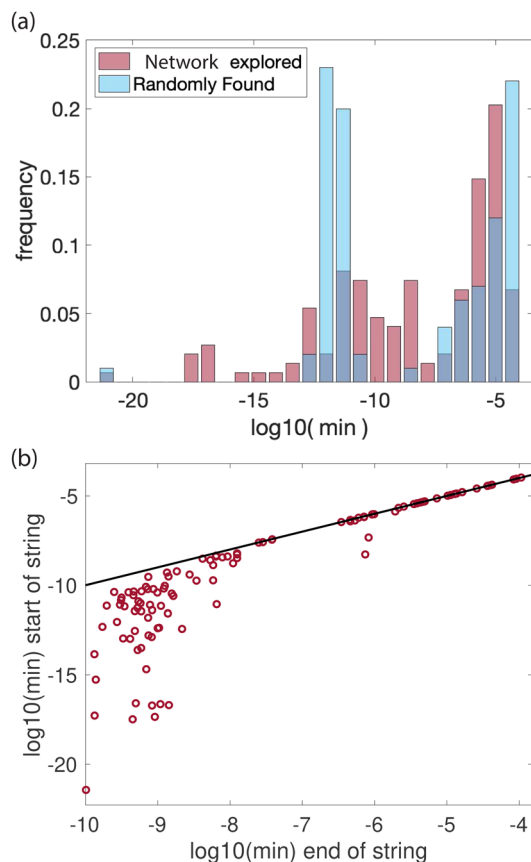


Fig. 5 (a) Histogram of the \log_{10} of the energy of the minima found either by randomly selecting initial points for the particles and then gradient descending until tolerance is reached or using the hydra string method from one randomly selected minimum. (b) Comparison of the energy of the same minimum when it appears at the start of a string vs. the end of a string, showing energies below about 10^{-8} are not accurately resolved.

an example of the latter case. To account for both these cases, the hydra string method is modified to (1) only compare non-rattler particles to determine if states are the same and (2) only add a new state to the list of states to be explored if it is jammed (no zero eigenvalues of the Hessian matrix after recursively removing rattler particles). The detailed algorithm is:

- In parallel, climb strings to first order saddle points truncating the string if it becomes non-monotonic (leaves the basin of attraction of the minimum). Check for convergence and one negative eigenvalue over the entire system. Rattler particles are included because the climbing string does not ignore rattlers; rattlers may transition back into the jammed backbones of the system.

- In serial, process saddle points.

- Recursively remove rattlers (fewer than 3 contacts) and compare remaining particle locations to list of found saddles. If configuration is unique, proceed. If not, continue to next saddle.

- Extend the end of the string a short way past the saddle, in the direction tangent to the end of the string. Evolve the string with gradient descent until convergence.

- Check that the string contains only one transition. If not, truncate and evolve again to convergence.
- Recursively remove rattlers and compare the end point (minimum reached) to the list of found minima.
- If the configuration is unique, add minimum to list of found minimum and if the configuration is jammed (no zero eigenvalues of the Hessian matrix of non-rattler particles) add to the list of minima to be explored.
- Add the edge to the list of edges and proceed to the next saddle.
- Once all saddle points have been processed, proceed to exploring the next minimum on the to-be-explored list.

Code is available on GitHub.²³

4 Results

With any sampling technique, one would want to know if a reasonably representative sample has been obtained. I first present a comparison between randomly sampling distinct minima and the hydra string network sampling technique to argue that the hydra string method does a better job. Then, I will discuss specific features of the granular system discovered using the hydra string sampling method.

4.1 Random sampling vs. hydra string sampling

Fig. 5(a) shows a histogram of the energies of the minimum states found using both random samples and with the hydra string method. While not identical, both show the bimodal distribution of energies expected. The hydra string method allows for an additional check on the accuracy of the numerical method as the same energy-minimizing state can be re-found. In Fig. 5(b) I compare the found minima energies on the end of the string with the starting minima energy at the start of a different string. These points deviate significantly from the line $y = x$ below about 10^{-8} , indicating these energies are not properly resolved. This is consistent with the positions of the particles not being resolved below 10^{-7} as the error in the energy is on the order of $\delta_{ij}^{5/4}\epsilon$ where ϵ is the error in the positional difference, $x_i - x_j$, and δ_{ij} is typically small. The fact that the starting end minimum is lower in energy is consistent with the fact that it has undergone considerably more iterations of gradient descent than the ending end minimum. Returning to Fig. 5(a), the discrepancy between the histograms below about 10^{-8} is likely due to these numerical resolution issues rather than the sampling technique. It appears both methods adequately sample possible minima in the system.

I proceed to investigate the sampled saddle points around one minimum. For any high-dimensional energy landscape, computational cost will likely prohibit finding all saddles and transition paths. Rather I investigate if a reasonable sample of available transition paths has been found. Fig. 6(a) shows the number of unique saddles found surrounding different minima with their energy plotted on the x -axis for different parameter values. As expected, sending out 256 climbing strings from each of the 25 independent minima (red triangles)



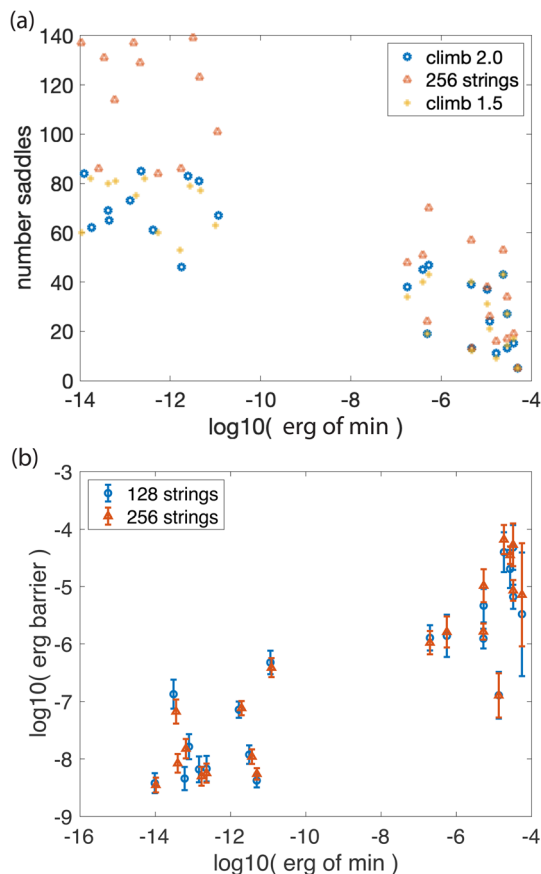


Fig. 6 For the same 25 minima, climbing strings are sent out to compare different parameters. (a) Scatter plot of the number of unique saddles found vs. energy of the minimum with climbing parameters 2.0 and 1.5 (128 strings) and increasing the number of strings sent out to 256 (climbing parameter 2.0). (b) Average energy barrier vs. energy of the minimum when both 128 and 256 strings are sent out (climbing parameter 2.0). Error is Monte Carlo error, std/\sqrt{n} .

finds more unique saddle points than sending out 128 climbing strings (blue circles), especially for lower-energy minima. For comparison, 128 strings with a slower climbing speed $\nu = 1.5$ are also included (yellow plus signs) with limited effect on the number of saddles found. I argue that a reasonable sample means the energy barrier heights to saddle points have been well sampled, rather than converging in the number of saddle found. Fig. 6(b) shows that although more unique saddles are found, the average energy barrier seen is not significantly changed. Therefore, 128 climbing strings are sufficient to sample available transition paths from a minimum.

Another way to assess if a reasonable sample of the energy landscape has been obtained is to ensure there is no bias to the energy barriers found. That is, looking at the statistics of the outgoing energy barriers from a given minimum should be the same as the statistics of the incoming barriers if the collection of pathways from multiple independent minima are to represent typical pathways across the entire landscape. Fig. 7(a) immediately shows bias in the climbing string method, as it tends to find pathways leading to higher energy states; the outgoing energy barrier is higher than the incoming energy barrier. This




Fig. 7 (a) and (b) For the same 25 minima, climbing strings are sent out with three sets of parameters changing climbing parameter ν and number of strings. (a) For each transition, the outgoing and incoming energy barrier is compared. (b) Histogram of the distance from the line $y = x$ is plotted, where negative indicates below the line. The solid line at 0 is shown for reference to the dashed line at the mean distance from the line $y = x$. (c) and (d) Same, but for all hydra string found energy barriers. (e) and (f) Same, but only considering energy barriers from the hydra string method that start or end at the 150 explored minima.

suggests that perhaps the string is climbing too fast, preferring higher-energy barriers. However, a reduction in the climbing speed from $\nu = 2$ (blue circles: fully inverting the gradient in the climbing direction) to $\nu = 1.5$ (yellow stars: half inverting the gradient in the climbing direction, as 1 is no climbing at all) does not noticeably change the distribution of pathways found. This is further confirmed in Fig. 7(b) that depicts the distances of each point in Fig. 7(a) from the line $y = x$, showing all three parameter sets produce the same mean off-set indicated by the black dotted lines that are indistinguishable. This is in sharp contrast to the sample produced by the hydra string method (Fig. 7c–f) that show significantly less bias. The difference is that the hydra string sampling allows for finding pathways back to minima that were already explored with out-going climbing strings. As the hydra string method only samples a small local region of the energy-landscape, many minima remain that have only been found as end-points on pathways. Removing these periphery minima, and only considering pathways that either start or end at the 150 explored minima returns a non-biased sample of pathways (Fig. 7e and f).

Overall, the hydra string method is creating a reasonable picture of the minima and available pathways in the network by






Fig. 8 Hydra string method, comparing non-rattler particles to determine "same" and only exploring mechanically jammed states. Triangles indicate floppy states. The colorbar indicates the log of the energy of the state.

performing a detailed sampling of one region, as opposed to independently sampling pathways around one minimum but picking minima from different areas on the energy landscape.

4.2 Granular energy landscape

Having established a method for sampling minima and transition pathways, I now show some features of the simple granular system's energy landscape. First, I switch to using the aspects of the hydra string method developed specifically for granular systems. Notably removing rattler particles before determining if states are the same and further testing if the system is jammed. Due to the degeneracy of minima for floppy unjammed states, only minima that are jammed are further explored with the hydra string method. The resulting network of explored minima is shown in Fig. 8, with floppy states shown as triangular nodes and jammed states as circular nodes.

One feature of the energy landscape revealed in Fig. 7(c) and (e) is that neighboring minima must have similar magnitudes of energies, as the incoming and outgoing energy barriers are clustered around the line $y = x$. The minima on both sides of a pathway must have comparable energies, as their difference




Fig. 9 The correlation of the log 10 of the energy of nodes in the network as a function of their shortest path-length separation.

corresponds to the difference between the forward and backward energy barriers (provided the barriers are not too much larger than the energies, which is validated next). To quantify this similarity, I take the correlation of the minima energies separated by the same path length in the network. The path length is the smallest number of edges that must be traversed to transit between two nodes. This correlation is plotted in Fig. 9 and shows the correlation remains until about path length 4. Note that the larger path lengths 8 and 9 appear less frequently in the sampled network and are subject to sampling error. This suggested the energy landscape might have some structure, like 'islands' of jammed states at higher energies surrounded by a 'sea' of floppy states at numerically zero energy.

Not only are the energies of neighboring minima correlated, but the energy barrier heights are correlated to the energy of the minimum. In Fig. 10 each histogram of energy barrier heights is formed from outgoing transition paths from minima with energy within the bin whose center is given by the text in the panel. There is a decrease and also spreading of energy barrier heights as the energy of the minima decreases. In fact, the mode of the histogram appears to be about equal to the minimum energy bin. Since this is a log 10 scale, this is similarity in order of magnitude. Perhaps one expects that a more tightly packed system (higher minimum energy) would require changes on that same order of magnitude to make small rearrangements of the particles to new states.

5 Conclusions

I have demonstrated that sampling the space of possible transitions for a high-dimensional energy landscape is better accomplished with a detailed sampling of a local area, forming a network of minima connected by minimum energy paths with the hydra string method. Bias towards higher-energy barriers potentially introduced by the climbing string method is removed with this method allowing for the discovery of new transition paths that arrive at previously-explored minima. This also provides a network representation of the energy landscape, with the energy-minimizing states as nodes and the minimum energy transition pathways between states as edges.

I demonstrated how this method reveals new features of an experimentally-based 19-particle granular system. Energy barriers between neighboring states appear correlated to the energies of the states themselves, leading to correlations between state energies that decay as a function of path length on the network. Unjammed states have zero energy, and surround the 'rugged mesas' of higher energy jammed states.

The computational costs will increase as the number of particles increases, as will the number of unique saddles surrounding a given minimum. The latter is likely to add little additional computational cost, as the method seeks to sample the distribution that these saddles form rather than exhaustively enumerate them. However, this does rely on the assumption that extremely rare saddles or pathways are not the driving mechanism behind the physical behavior of a granular material.



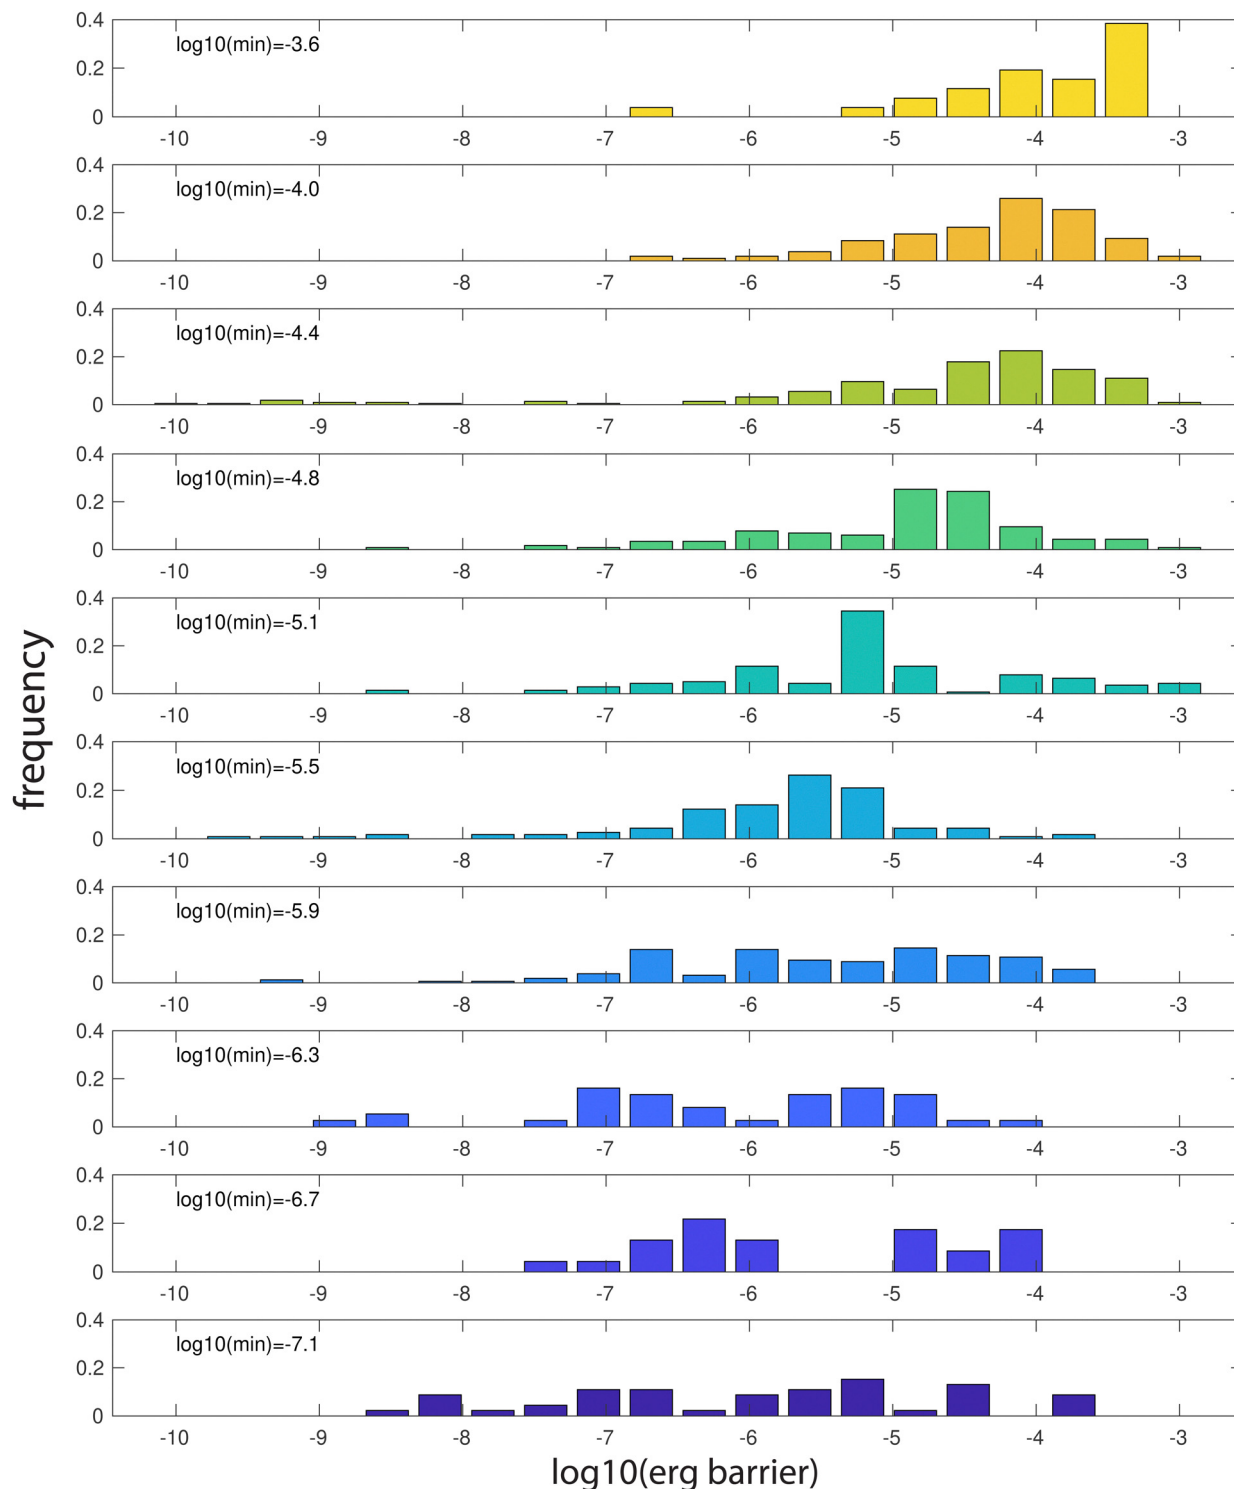


Fig. 10 For the energy barriers found using the hydra string method resulting in the network depicted in Fig. 8, histograms of the energy barrier sizes are created after grouping the starting states' energy into bins with centers indicated by the text in each plot. The y-axis indicates the frequency of energy barriers given that the starting minima fell into the given plot's range.

Even with the 19-particle system studied here, there are many features of the energy landscape that can still be explored. For example, the transitions shown in Fig. 3 seem to indicate that higher energy barriers between states involve

more particles moving. One could build on this by looking for pathways between states that minimize either particle movement or summed energy barrier crossings, rather than minimizing the number of barriers crossed (path length). This could



allow for an exploration of how ‘close’ any given state is to a failure event, which in the small model system considered here is a transition to an unjammed state. As the disk sizes increase, one might expect a given state to become farther from an unjammed state until no unjammed states are possible. Thus studying the energy landscape could be a way to examine how far a system is from the jamming point, together with some notion of distance be it Euclidean distance, transition path length, or something else.

Data availability

The code is publicly available on GitHub, as indicated in the manuscript.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

I thank Karen Daniels and Silke Henkes for useful discussions and Josh Miller for providing the experimental picture. This work is supported by NSF grant number DMS-2307297.

Notes and references

- 1 D. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses*, Cambridge University Press, Cambridge, 2004.
- 2 H. Eyring, *J. Chem. Phys.*, 1935, **3**, 107–115.
- 3 O. M. Becker and M. Karplus, *J. Chem. Phys.*, 1997, **106**, 1495–1517.
- 4 Y. Xing, Y. Yuan, H. Yuan, S. Zhang, Z. Zeng, X. Zheng, C. Xia and Y. Wang, *Nat. Phys.*, 2024, **20**, 646–652.
- 5 H. J. Hwang, R. A. Riggleman and J. C. Crocker, *Nat. Mater.*, 2016, **15**, 1031–1036.
- 6 P. Suryadevara, M. Casiulis and S. Martiniani, Mirages in the Energy Landscape of Soft Sphere Packings, *arXiv*, 2024, preprint, arXiv:2409.12113, DOI: [10.48550/arXiv.2409.12113](https://doi.org/10.48550/arXiv.2409.12113).
- 7 M. L. Manning and A. J. Liu, *Phys. Rev. Lett.*, 2011, **107**, 108302.
- 8 S. Luding, K. Taghizadeh, C. Cheng and L. Kondic, *Soft Matter*, 2022, **18**, 1868–1884.
- 9 K. Fritsch, J. Friedrich, F. Parak and J. Skinner, *Proc. Natl. Acad. Sci. U. S. A.*, 1996, **93**, 15141–15145.
- 10 P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani and F. Zamponi, *Nat. Commun.*, 2014, **5**, 3725.
- 11 A. Thirumalaiswamy, R. A. Riggleman and J. C. Crocker, *Proc. Natl. Acad. Sci. U. S. A.*, 2022, **119**, e2210535119.
- 12 N. Xu, J. Blawdziewicz and C. S. O'Hern, *Phys. Rev. E:Stat., Nonlinear, Soft Matter Phys.*, 2005, **71**, 061306.
- 13 G.-J. Gao, J. Blawdziewicz and C. S. O'Hern, *Philos. Mag.*, 2007, **87**, 425–431.
- 14 C. Moakler and K. A. Newhall, *Granular Matter*, 2021, **24**, 24.
- 15 W. E, W. Ren and E. Vanden-Eijnden, *J. Chem. Phys.*, 2007, **126**, 164103.
- 16 W. Ren and E. Vanden-Eijnden, *J. Chem. Phys.*, 2013, **138**, 134105.
- 17 P. Cao, M. P. Short and S. Yip, *Proc. Natl. Acad. Sci. U. S. A.*, 2019, **116**, 18790–18797.
- 18 P. Cao, M. Li, R. J. Heugle, H. S. Park and X. Lin, *Phys. Rev. E:Stat., Nonlinear, Soft Matter Phys.*, 2012, **86**, 016710.
- 19 A. Kushima, X. Lin, J. Li, J. Eapen, J. C. Mauro, X. Qian, P. Diep and S. Yip, *J. Chem. Phys.*, 2009, **130**, 224504.
- 20 E. T. Owens and K. E. Daniels, *EPL*, 2011, **94**, 54005.
- 21 J. E. Kollmer and K. E. Daniels, *Soft Matter*, 2019, **15**, 1793–1798.
- 22 K. L. Johnson, *Contact Mechanics*, Cambridge University Press, 1985.
- 23 K. Newhall, github.com/knewhall/Hydra_String, https://github.com/knewhall/Hydra_String.

