Granular Matter (2022) 24:24
https://doi.org/10.1007/s10035-021-01184-5

ORIGINAL PAPER

The hydra string method: a novel means to explore potential energy surfaces and its application to granular materials

Christopher Moakler[1] · **Katherine A. Newhall**[2]

Received: 2 February 2021 / Accepted: 26 October 2021 / Published online: 15 December 2021

Abstract
We present a novel means to understand granular materials, the Hydra String Method (HSM). This is an efficient and autonomous way to trawl an arbitrary potential energy surface (or any similarly high dimensional function) that enumerates the saddle points, minima, and minimum energy paths between them. In doing so, it creates a reduced dimensional network representation of this surface. We also present a series of tests to choose optimized parameters for the application of the HSM. We apply this to the potential energy function of a granular system consisting of a configuration of bi-disperse, frictionless, soft spheres. Future work will make use of the found ensemble of transition pathways to statistically predict the dynamics of a system of grains.

Keywords Transition path · Frictionless soft spheres · Stable packings · Computational algorithm · Energy landscape

1 Introduction

The class of granular materials is second only to water in materials used by humans and despite recent advancements to understand their unusual behavior, scientists still lack an adequate way to describe the full range of their behavior [1–4]. Attempts at firmly grasping the dynamics and mechanics of granular materials have been as fruitful as firmly grasping a handful of sand, it quickly runs away from us.

This work presents the foundation of a new approach to predict the dynamic response of purely repulsive-interaction granular systems from a defined energy landscape. We take a soft-sphere quadratic potential between overlapping 2D bidisperse frictionless disks and map out energy-minimizing jammed states as well as the paths between neighboring states that minimize the energy, the minimum energy paths (MEP). Unlike other work that seeks to enumerate all uniquely jammed states [5, 6], we only seek a representative sample of states as full enumeration is computationally infeasible for even small systems of 12 particles. Note that the neighboring states and pathways between them are dependent on the protocol generating the packing (c.f. Ref. [7] that allows the box size to change).

We expect this energy-based approach will lead to advances in understanding both thermal and athermal granular systems. In small colloidal packings, random thermal motion can have a significant effect on rheological properties [8], and even athermal granular systems can be subject to thermal-like fluctuations such as in vibration-fluidized systems [9]. Energy landscape exploration methods are already being implemented for the thermal jamming transition [10, 11] and shearing using an extended potential energy landscape to account for the external strain [12]. Thus we anticipate our efficient energy landscape exploration algorithm will have predictive power by speeding up energy landscape exploration and creating transition networks like those generated from observing rearrangements under strain as in Ref. [13].

The idea of exploring an energy landscape is not itself new. In fact the building blocks of our automated "Hydra" method are the string method [14] and its climbing variant [15] (two of many such methods [16–23]) developed for use in computational chemistry to probe the evolution of a single reactant to a product state with possibly many intermediate

✉ Christopher Moakler
cmoakler@unc.edu

Katherine A. Newhall
knewhall@unc.edu

[1] Department of Physics and Astronomy, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

[2] Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

transition states. These transitions are often diagrammatically sketched as a graph of energy vs. reaction coordinate called reaction coordinate diagrams [24]. Intermediate transition states along this path appear as maximums in energy, but correspond to saddle points in the full high-dimensional phase space of the system. Such a saddle point sits on top of the local lowest energy barrier separating two states and therefore chemical reactions are most likely to proceed over this barrier near a MEP (c.f. [25]). We show a sample MEP in Fig. 1 for a simple 2D potential (a and b) and for a more complicated soft-sphere model (c) discussed in more detail in Sect. 3.

Our automated method reduces the energy landscape to a network representation of transition pathways between neighboring states, through saddle points, like those presented for smaller systems such as six colloidal particles with depletion attraction [26], nematic liquid crystals on 2-D hexagons [27], small atomic clusters interacting with Lennard-Jones potential [28, 29] and 2D lattice polymers [30]. Unlike previous approaches, we find not just the transition state but the full transition path. The method sends out multiple climbing strings to find saddle points that are ensured to be on the edge of the basin of attraction (for gradient-descent dynamics) around a minimum by the monotonicity of the energy along the string. From the unique saddle points found, strings are descended to find neighboring minima. This process is repeated and, like a Hydra, this "string" will grow new exploring heads to continue mapping the energy landscape.

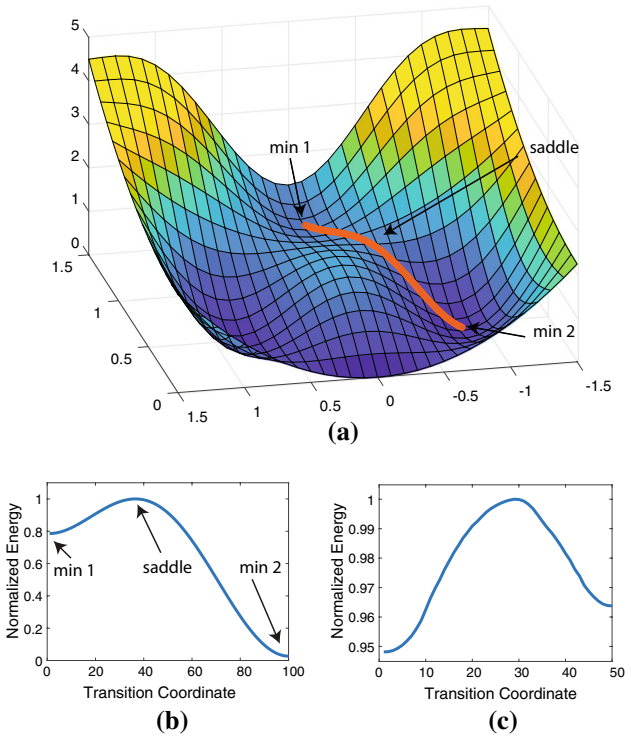

Fig. 1 **a** A simple potential function with a MEP shown in orange and **b** the normalized energy along this transition path or "coordinate". **c** Example of an energy versus transition coordinate for a granular system moving between minima

The main focus of this paper is laying out a sequence of tests that can be applied to any system with an energy landscape to determine parameters that balance accuracy with computational cost to be able to efficiently sample higher-dimensional more complex systems. These tests seek general characteristics of the basins of attractions (of gradient descent dynamics) around each energy-minimizing state in order to efficiently climb to the edge of the basins and thereby sample the neighboring states. The first test dictates how precisely each state must be resolved to determine if two states are unique. The next test characterizes the radial size of a basin and informs how to resolve the smallest features of the landscape and how far away to look for saddle points on the edge of this basin. The last test determines how many climbing string should be sent out to have a reasonable sample of unique saddle points on the edge of the basin of attraction.

The remainder of the paper is organized as follows. In Sect. 2 we briefly review the string and climbing string methods and present the details of our Hydra string method. We explain the three main tests to apply to a system in order to determine parameters for the Hydra string method that will balance accuracy with computational efficiency in Sect. 4. In Sect. 5 we apply these tests to the granular system given in Sect. 3. We conclude the paper with discussions of various parameter choices in Sect. 6 and conclusions in Sect. 7.

2 The hydra string method and algorithm

The Hydra String Method (HSM) is designed to explore the potential energy surface (PES) in a systematic way, creating a network of connected minima and intermediate saddle points. Its base pieces are the existing string [14] and climbing string [15] methods that find an MEP between two minima, or a minima and unknown saddle point, respectively. We briefly review the string and climbing string method before proceeding to describe our algorithm in more detail.

2.1 The string method

The String Method [14] finds the MEP between two minima but is more computationally efficient and exhibits greater stability than the popular Nudged Elastic Band (NEB) method [31]. Both methods evolves copies, or "images", of the system in configuration space by gradient descent. To address the issue of the images falling towards the minima on either end of the path and bunching up there, the NEB method introduced springs between the images to keep them separated along the MEP. These spring forces couple the

images that would otherwise independently undergo gradient descent. The String Method decouples these competing forces on the images and treats them as two different steps, thus enhancing efficiency. The String Method first allows the images to follow a gradient descent for one time step and then interpolates the images to equally space them along the string in arclength.

One of the benefits of decoupling these two processes is that it allows the user to change the interpolation method "on the fly". A linear interpolation method can easily be changed to, for example, a natural spline method or any other desired interpolation method. The decoupled images also do not effect one another before the interpolation, so they can be computed in parallel.

The string method can be implemented with either one or both ends of the string held fixed. When both ends are held fixed the string methods converges to a MEP between those two points while keeping only one end point fixed allows the string method to converge to a MEP between a known minimum and a newly found minimum. This method was designed as a tool for computational chemistry where one may know both end points (known reactant and product states) or with only a single known end point (known reactant or product state).

2.2 The climbing string method

The Climbing String Method [15] is a modification of the String Method that finds saddle points from a single known minimum by allowing the final image along the string to climb up against the gradient with the following force acting on it:

$$\mathbf{F}_{i_{cli}} = -\nabla V(\mathbf{r}_{i_{cli}}) + \nu(\nabla V(\mathbf{r}_{i_{cli}}), \hat{\tau}_{i_{cli}})\hat{\tau}_{i_{cli}}. \quad (1)$$

Here, ∇V is the gradient of the potential energy, $\mathbf{r}_{i_{cli}}$ is the configuration of the system at the climbing image, ν is a tunable parameter (generally set to 2) that controls the climbing speed of the string, and $\hat{\tau}_{i_{cli}} = (\mathbf{r}_{i_{cli}} - \mathbf{r}_{i_{cli}-1})/|\mathbf{r}_{i_{cli}} - \mathbf{r}_{i_{cli}-1}|$ is the approximate unit tangent to the string at the climbing image.

The first image on the string is kept fixed at a minimum and the intermediary images follow a gradient descent at each time step. As in the String Method, the intermediary images are redistributed via an independent interpolation. A computationally cheap option for this redistribution is a simple linear interpolation. But, other methods can easily be substituted.

To ensure that the string remains in the basin of attraction of the starting minimum, if, at any time step, the energy along the string fails to be monotonically increasing, the string is cut at the point of non-monotonicity. The images are then re-interpolated along this shorter string and the process is continued until it converges to a MEP.

While at first glance the climbing string method appears inefficient relative to non-string methods such as the two-image climbing dimer method [18] or improved eigenvector-following methods such as r-ARTn and d-ARTn [23] due to the number of copies of the system that are evolved at each time step, we see a number of advantages. First, both these methods require post-checks to ensure the found saddle point is located on the edge of the original minima's basin of attraction. As just stated above, the climbing string method ensures the saddle point is on the edge of the desired basin, thereby saving computation time. Additionally, these other methods suffer from convergence issues ([20] improves the convergence for the climbing dimer, and the improvement in [23] only increased the rate to 87%), something we rarely see with the climbing string method that converges more than 98% of the time.

Second, we desire the entire MEP, not just the energy barrier between the minimum and saddle point, in order to expand beyond thermally-activated processes for granular systems. These saddle-point-only finding methods would require finding the MEP as part of the post-check. As with the String method, the MEP is sensitive to the initial guess. That is, a strict linear interpolation between the minimum and saddle may converge to a path through many minima, while a monotonic MEP may exist along some other curving and twisting path. Therefore finding the MEP directly while also finding the saddle point appears more reliable.

Last, we are willing to sacrifice accuracy in favor of computational efficiency, explaining in the remainder of this paper how to efficiently choose parameters for the Climbing String method. If greater accuracy is required, each approximate saddle point found could be further refined by using the inexact Newton method, as was done in [15], utilizing the climbing dimer or eigenvector following method, or using a non-linear spacing of the climbing string images to better approximate the tangent to the string at the climbing image.

2.3 The hydra string method

We describe the Hydra string method in more detail, which is illustrated in Fig. 2 and summarized by the pseudo-code in Algorithm 1. Its MatLab implementation can be found on GitHub; see Code Availability at the end of the paper.

An initial minimum is found by choosing a random point and then performing a gradient descent. This minimum is added to a running list of minima found on the PES. From this initial minimum, a number of new strings, $n_{strings}$, with n_{images} images along each string are extended a given distance, $dist_{ext}$, in a random direction and allowed to climb following the climbing string method described above. (A set of initial strings evolving are shown in Fig. 2a). Those strings that converge to a saddle point (see the converged strings in Fig. 2b) have their saddle points added to a running list of saddle points. The images on these

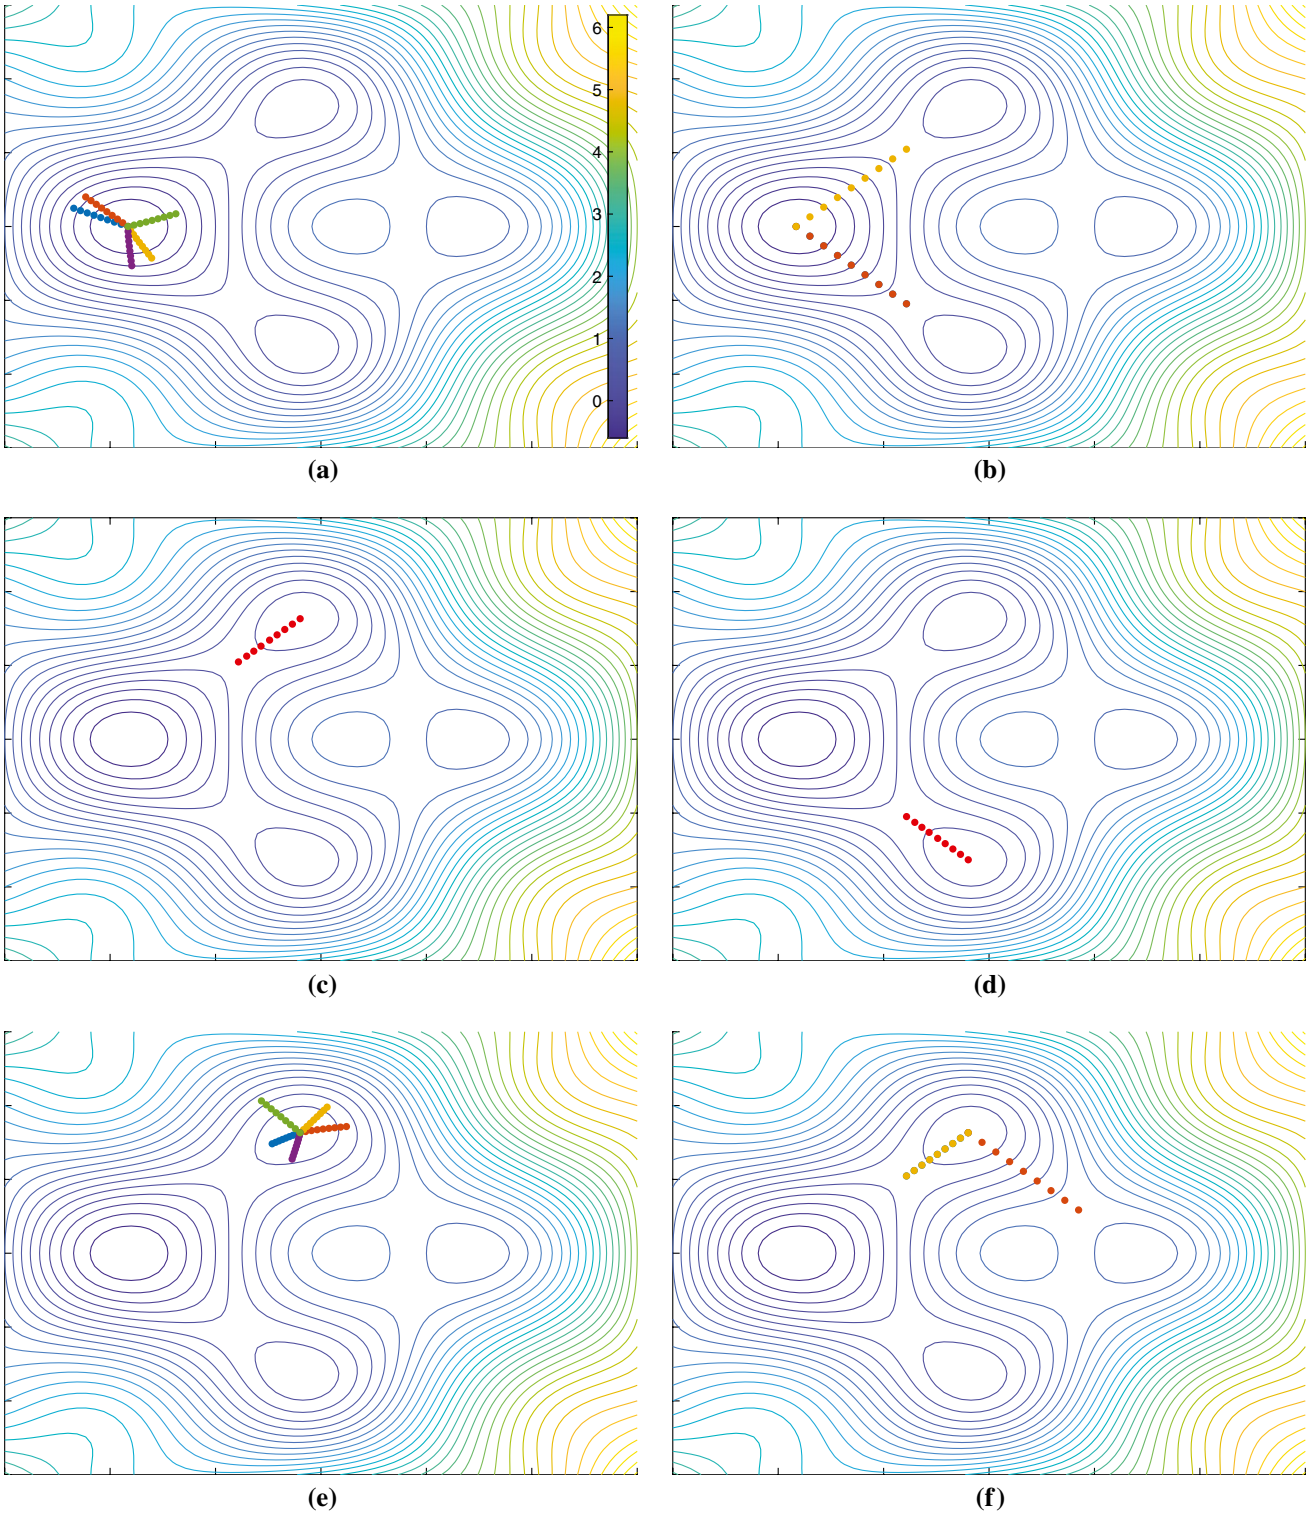

Fig. 2 Depiction of the HSM on the 2D PES, $E = x^2 + y^2 + \sin(\pi x) + \cos(\pi y)$, shown as a contour plot. **a** Evolving climbing strings with one end pointed pinned at the original minimum. **b** Convergence of the strings to minimum energy paths ending at two new saddles. **c** and **d** Descent from these saddles to new minima. **e** Evolving and **f** convergence of climbing strings from one of the newly found minima

saddle points are then perturbed, away from the minimum from which they were found, and allowed to follow a gradient descent until they converge to another minimum (see examples of such a converged path in Fig. 2c and d). These new minima are also added to the running list of minima. An unexplored minimum is chosen from the list and the process begins again (see new evolving strings in Fig. 2e and converged strings in Fig. 2f). This procedure is repeated until either no new minima/saddles are found or a predetermined number of minima are explored. In practice, one only wants to add unique minima(saddles) to the list of minima(saddles). Thus at every step where new minima(saddles) are found, they must be compared to the pre-existing list of minima(saddles) with a tolerance, $tol_{sad/min}$, to determine uniqueness before being added to the list.

Algorithm 1 Hydra String Algorithm

1: x_{random} %randomly chosen position
2: $x_{initial}$ = Descend(x_{random}) %descend to initial minimum in system
3: $x_{min_{new}} = x_{initial}$ %add first minimum to list of unexplored minima
4: $x_{min_{unique}} = x_{initial}$ %add first minimum to list of unique minima

5: **while** $x_{min_{new}}$ is not empty **do**
6: $x_{min} = x_{min_{new}}(1)$ %Pick next minimum to explore
7: delete $x_{min_{new}}(1)$ %Remove minimum to be explored from unexplored list

8: **for** $i = 1{:}n_{strings}$ **do**
9: $x_{min_{extended}}$ = Extend(x_{min},$dist_{ext}$) %Extend a string in a random direction
10: $possible_{unique\_sad}(i)$ = Climb($x_{min_{extended}}$)

11: **if** $possible_{unique\_sad}(i)$ is unique **then**
12: $x_{sad_{temp}}$.append($possible_{unique\_sad}(i)$)
13: **else**
14: continue
15: **end if**

16: **end for**

17: **for** $i = 1{:}Length(x_{sad_{new}})$ **do**
18: $possible_{unique\_min}(i)$ = Descend($x_{sad_{new}}(i)$)

19: **if** $possible_{unique\_min}(i)$ is unique **then**
20: $x_{min_{temp}}$.append($possible_{unique\_min}(i)$)
21: **else**
22: continue
23: **end if**

24: **end for**

25: $x_{min_{new}}$.append($x_{min_{temp}}$) %List of unexplored minima
26: $x_{min_{unique}}$.append($x_{min_{temp}}$) %Full list of unique minima
27: $x_{sad_{unique}}$.append($x_{sad_{temp}}$) %Full list of unique saddles
28: **end while**

2.4 Parallelizability

Two of the major benefits of this algorithm are its high parallelizability and its high degree of flexibility in both execution and application. Each minimum can be explored independently from the others, the strings all act independently from one another, and all of the images along each string act independently from each other during the gradient descent step. Thus, the while loop in Line 5 of Algorithm 1 can be run as a parallel process, the for loops in lines 8 and 17 can also be run as parallel processes, and the string evolution in the `Climb` and `Descend` functions can have each image evolve in parallel. It is difficult to parallelize all of these process at the same time and is disallowed in many programming languages. So a choice of where to parallelize must be made.

Parallelizing over the while loop requires significant inter worker communication as it requires the list of potential new saddles and minima to be transferred at the end of each iteration. However, parallelizing here is beneficial for investigating smaller systems where the minima/saddle lists are small in memory or for a system that is very large or has a complicated potential function such that the time spent finding the minima and saddles is relatively long compared to the time spent transferring data between workers. The two for loops are easier to parallelize and parallelizing here would be a natural choice for systems with simple potential functions where the time spent climbing and descending strings is relatively quick. These options in the implementation of the HSM allows it to be adapted to efficiently explore any system.

3 Example system

As a test of the HSM, we have applied it to a simple example system. This test system we analyze is made up of 24 bidisperse soft spheres in a 2D periodic domain of unit length. The particles have a radii ratio of 1.4 to prevent crystallization with half having a radius $R_L = 0.1336$ and half having a radius $R_S = 0.0954$. This arrangement of particles is shown in Fig. 3.

The periodic boundaries allow us to simulate a larger system with comparatively few particles. However, the periodicity introduces two symmetries to the system which are undesirable for the analysis we wish to perform; critical points become 2-dimensional sheets rather than points. These symmetries correspond to concerted shifts of all particles in the x or y directions. To break this symmetry, we fix one of the particles so it is not allowed to move even when other particles exert a force on it. This removes the symmetries and reduces the dimensionality of the system to 46.

The potential energy of the collection of these soft spheres is simply the sum of the pairwise spring potentials between the N particles, given by:

$$V = \sum_{i=1}^{N} \sum_{j=i}^{N} \boldsymbol{\kappa}_{i,j} \left(1 - \frac{|\mathbf{x}_i - \mathbf{x}_j|}{(R_i + R_j)} \right)^2 J_{ij}. \tag{2}$$

In Eq. 2, N is the number of soft spheres, $\boldsymbol{\kappa}_{i,j}$ is the stiffness tensor (analogous to the spring constant) between particles i and j, $\mathbf{x}_i$ is the position vector of the center of particle i, R_i is the radius of particle i, and J_{ij} is zero if particles i and j do not overlap and one if they do. By setting $\kappa = 1$ we find the energy-minimizing configurations, like the one shown in Fig. 3, appear similar to photoelastic disk systems studied by experimentalists studying shearing particles [32] (a review of the use of photoelastic particles in studying granular materials is available in [33]) and PNIPAM colloidal particles like those in [8] that also interact elastically.

The number of minima of such a system scales exponentially with N [6] and most physical systems will have N much larger than the 24 we have used here. This system is large enough to prohibit an exact enumeration of minima/saddles yet not untractable as a test system for the HSM.

4 Applying the method

The HSM requires several parameters, summarized in Table 1, to be set to efficiently generate an accurate and useful network. To choose these parameters, we need information about the "basic structure" of the PES. By "basic structure" we mostly refer to the structure of the basins of attraction (for gradient-descent dynamics) around the minima in the PES since the Hydra String Method finds minima, and first order saddle points, which lie along the ridges of the basins. Of chief interest is the distance between basins, a characteristic radial size of the basins, and the number of saddles connected to a minimum. Knowing this structure tells us how far away from minima to look for saddles, how accurately we need to resolve the basin, and how many strings we need to extend from a minimum to find the saddles.

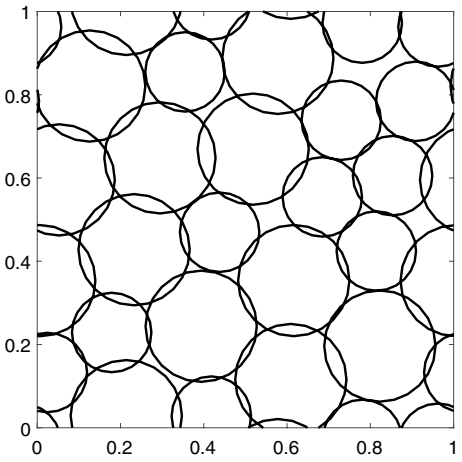

Fig. 3 A sample energy minimizing configuration of the 24 sphere system

Since the structure of the basins of an arbitrary function are generally not known a priori, we propose three numerical experiments as a way to deduce this basic structure. Each numerical experiment we propose addresses one of the characteristics of the basins we are interested in. We will make extensive use of these experiments to determine values for the various parameters used in the Hydra String Method which will be discussed in Sect. 5.

In our experience, determining this basic structure requires an iterative approach. That is, the easiest way to pick values for the parameters of the Hydra String Method is to first run the Hydra String Method. This can be done with very conservative choices for the various parameters which can then be used to refine those initial choices. This process can be repeated several times to more fully understand the structure.

In this section, we will show the results of these numerical experiments when run on the PES that arises from the system described in Sect. 3, a collection of 24 bi-disperse soft spheres that interact with pairwise linear spring forces contained in a unit box with periodic boundaries.

4.1 Determining the distance between basins

When applying the Hydra String Method, the user must designate what the tolerance is to call minima(saddles) distinct. This arises because the climbing strings may approach the same saddle point or the hydra may curl back on itself and find the same minima over again. The strings may approach this minima(saddle) from a slightly different direction and due to the flatness of the PES at these critical points, they may approach slightly different values. One should appeal to the physics of the underlying system to help determine which minima(saddles) should be considered distinct.

To investigate this, we propose the following numerical experiment. Determine a collection of minima or saddle points. These can be found by either randomly sampling your PES and performing a gradient descent to find local minima. Or the Hydra String Method can be implemented with conservative values for the parameters required (small $dist_{ext}$, small $tol_{sad/min}$, large n_{string}, and large n_{images}) as mentioned above. These parameters will be discussed at length in the following section. With either approach, all that is necessary to perform the analysis is a collection of possibly degenerate critical points.

Table 1 Hydra string method parameters

Parameter	Description	Value
ν	Controls how quickly the climbing image climbs towards a saddle, must be greater than 1	2
dt	The time step used in the string evolution ODE solver	10^{-3}
tol_{diff}	The tolerance used to stop the string evolution ODE solver	10^{-8}
$n_{step_{max}}$	The maximum number of time steps allowed in the string evolution ODE solver	10^5
$dist_{ext}$	The distance a string is initially extended away from a minimum or saddle	0.15
n_{images}	The number of images along a given string	10
$n_{strings}$	The number of strings to send out from each minimum to find new saddles	24
$tol_{sad/min}$	The tolerance to call a saddle/minimum unique	10^{-2}

Now calculate the pairwise distance between all of the minima(saddles) and plot them on a histogram (as in Fig. 4). Most of the pairwise distances should be quite large and indicate unarguably distinct critical points. A second population may arise near the stopping tolerance from the climbing string method ODE solver. Other populations may appear between these two pairwise distance groupings. By exploiting the physics that gives rise to the PES the user can set a minimum distance between basins(saddles) to call them unique.

4.2 Determining the radial size of a basin

Knowing the characteristic radial size of a basin allows the user to pick a distance to initially extend the climbing string. Unlike eigenvector following methods, the climbing string will always evolve towards a saddle point for any arbitrarily small initial displacement away from the minimum. However, near the minimum this evolution will be slow, as the gradient is near zero. Therefor efficient implementation requires an initial extension large enough to leave this locally flat region. The characteristic radial size of the basin also gives the user a general sense of the smallest features of interest of the PES. This knowledge can also be used to pick the number of images along each climbing string. These images are what ultimately detect these small scale features of the PES and ensure the MEP is resolved and has not left the basin of the minimum.

To determine this characteristic radial size of the basins, we propose the following numerical experiment. Determine a representative collection of minima of the PES, as in Sect. 4.1. Perturb these minima a small distance from the minimal state in many different directions. Then take these perturbed points, gradient descend them, and determine if these perturbed points return to their original minimal state by calculating the distance between the gradient descended configuration and the original minimum. If this is equal to or smaller than $tol_{sad/min}$, we consider it to have returned. The number of perturbed points that return to the original minimum is then determined. This procedure is repeated for many different perturbation distances. This should also be repeated for many different minima as the basins may not all be identical. These results can be plotted as they are in Fig. 5.

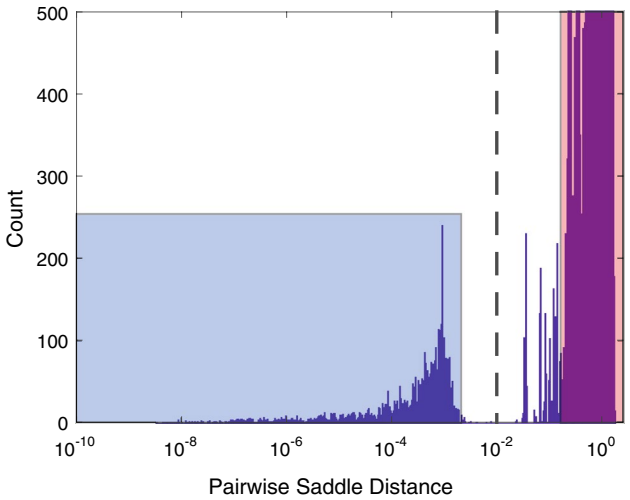

Fig. 4 Histogram showing the pairwise distances between saddle points in a PES that arises from the system described in Sect. 3. The population between 10^{-1} and 10^{0} (red) are the clearly unique saddles and the population smaller than approximately 10^{-3} (blue) is the population of degenerate saddles. The value we have chosen for $tol_{sad/min}$ is indicated by the vertical dashed line (color figure online)

This analysis is similar to that done in [6] where the authors describe the basins as containing a central spherical region within which all states can be relaxed back to the originating minimum. Beyond this spherical region, there are spindles where configurations lying in the spindles relax back to the minimum, but other configurations that do not. The radius of the central spherical region is what we consider the "characteristic size of a basin" because within this region, all perturbations will relax back to the original state. This distance can be seen in Fig. 5 as the point when a single line departs from 1. This analysis can also tell the user the size of the smallest features of interest in the PES as well as reveal interesting features about the structure of the basins.

This may be of interest in and of itself, but this analysis will be necessary to implement the Hydra String Method because we need these results to determine both $dist_{ext}$ and n_{images}, which will be discussed in the next section.

4.3 Determining the number of connected saddles

The Hydra String Method calls for extending some number of climbing strings from each minimum to find new saddles. To choose the number of strings to extend out, it is useful to know how many unique saddles are connected to a given minimum. That is, how many saddle points lie on the ridge of the basin of attraction for any given minimum.

To determine this, we propose the following numerical experiment. Again, find a collection of minima to study. From each minimum extend out one string and note the saddle it converges to, then send a second, a third, etc. Note how many unique saddles are found for the number of extended strings. In practice, this can be accomplished by sending out a large number of strings all at once and determining how many converged to unique saddles after the fact. This should be done for a large number of minima because, as before, each basin need not be the same as the others. Plotting these data as a series of histograms for various numbers of strings extended as shown in, Fig. 6, with the number of unique saddles found on the x-axis and the number of minima with that many unique saddles found on the y-axis will give the user an idea of how many unique saddles the basins in the PES have.

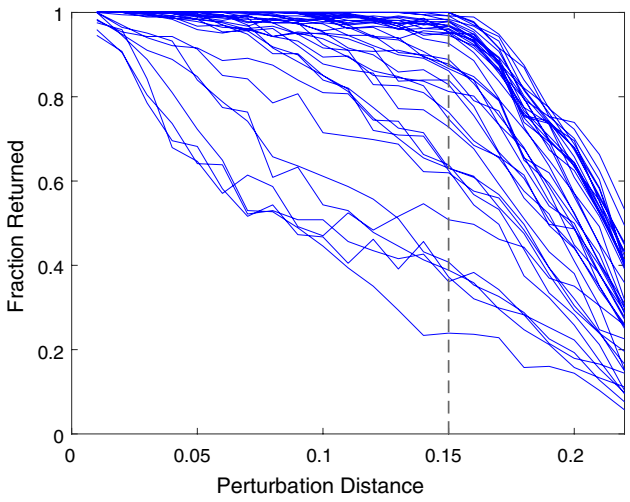

Fig. 5 The results of a radial basin size analysis for the PES that arises from the system described in Sect. 3. Each blue line represents a different minimum that was analyzed for its radial size. This analysis indicates most basins have a radius of approximately 0.15, denoted by the dashed line, because beyond this perturbation distance, a rapidly increasing fraction of perturbed points fail to return to the original minimum

4.4 Computational eficiency

Once optimized, the HSM is quite efficient in its execution. For our system, with the parameters in Table 1, an individual string converges to a saddle after an average of 3.25×10^4 time steps of the climbing string method. Furthermore the strings converge more than 98% of the time. We ran the HSM on a cluster of 24 parallel workers for one week. In that time, it explored $\approx$ 1650 minima finding more than 14,000 pathways. The test to determine the parameters in Table 1, also require a non-trivial amount of computations. Determining $dist_{est}$ and n_{images} take, by far, the most computational effort requiring almost 10^{10} time steps of the string method in our system. These force calculations can be run in parallel and on our system of 24 workers, can be completed in about a day.

The strings will sometimes converge to previously found saddles, this is discussed in Sect. 5.3.2, and when extending 24 strings, as we do in this implementation, the strings find unique saddles about 50% of the time. This can potentially be improved by extending strings in specified directions instead of in random directions, however at this point what these specified directions should be remains unknown.

5 Parameters involved in the hydra string algorithm and how to choose them

Equipped with the numerical experiments from Sect. 4, we will look at each parameter tabulated in Table 1 and describe how to choose a value for it. Specifically, we will discuss how to choose these parameters for the system described in Sect. 3. We find it convenient to group them into three categories of parameters: those used to find a saddle, those used to ensure the saddle is in the original basin, and those used to generate a network of saddles and minima. We discuss each group in turn.

5.1 Parameters used to find a saddle

These are the parameters needed to find an individual saddle, i.e. those needed to implement the climbing string method. They relate to numerically integrating the ordinary differential equation (ODE) describing the gradient descent of the images, and do not require specially designed tests to find these parameters.

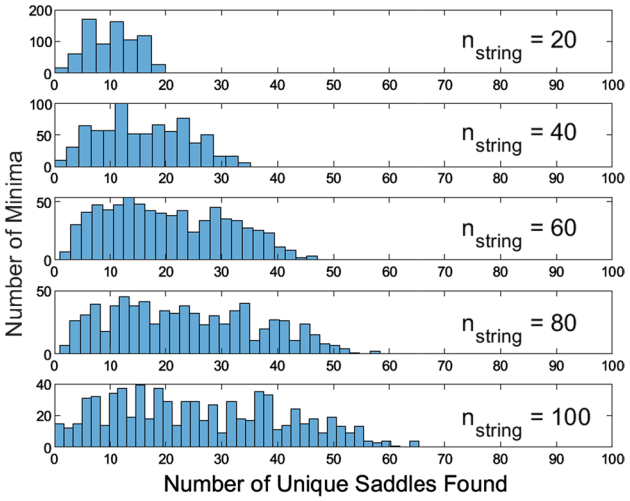

Fig. 6 A series of histograms showing the number of minima from which a search was started and how many unique saddles were found connected to the minimum with a given number of strings extended

5.1.1 *ν*

The value of ν controls how quickly the climbing image climbs and must be greater than 1, but not too large so as to cause the ODE solver to become unstable, preventing the climbing image from converging. Too small a value will cause a slower convergence. We use the typical value of 2 since this reflects the descent direction into the ascent direction, and the values of the other parameters can be tailored around this choice.

5.1.2 *dt*

The size of a time step of the climbing string and gradient descent ODE solver should be chosen by considering the normal accuracy, convergence, and stability issues: not too small to have an excessive time to convergence and not too large to become unstable. The value of dt should also be chosen such that the maximum movement of an image along the string during a gradient descent is small enough to not jump over any important features of the energy landscape. We used $dt = 10^{-3}$.

5.1.3 tol_{diff}

The criterion to stop the evolution of the climbing or descending string is set by the difference in euclidean distance of an image along the string between two consecutive time steps. That is, how much an image moves between two consecutive time steps (including both ODE movement and interpolation of the string images). This parameter should be smaller than the smallest feature of the PES one wishes to resolve but also large enough that the string converges in a reasonable amount of time. The results of the experiment in Sect. 4.2 can help determine the sizes of these smallest features. We found that 10^{-8} was a good balance between these conflicting goals. Note we choose not to terminate the string evolution early if the end is nearing an already-found minima/saddle point, as the calculation of distances is computationally expensive and potentially undesirable depending on how the algorithm is parallelized.

5.1.4 $n_{step_{max}}$

The maximum number of time steps allowed in the string evolution ODE solver needs to be large enough to allow convergence for most strings sent out during the snaking string algorithm. However, too large of a value leads to unnecessary computational load from runaway strings that have gone awry for one reason or another. This parameter is directly affected by dt, ν, and $dist_{ext}$. Thus, it should be chosen after suitable values of those parameters have been chosen. After finding the maximum number of time steps needed for a typical string evolution, via computational trials, a moderate margin should be added to this maximum to get a starting value of $n_{step_{max}}$. The competing goals of minimizing computation time and allowing enough strings to converge to obtain a representative sample of the saddles in a system must be balanced in choosing a value of this parameter. We used 10^5 in our system.

5.2 Parameters to ensure a saddle is in the original basin

These two parameters, $dist_{ext}$ and n_{images}, ensure that the initial climbing string is perturbed far enough away from the originating minimum (or initial string from a saddle) to allow for efficient climbing speed away from the locally flat region surrounding the minimum (or saddle), but also that the string has enough images to resolve the basin along this initial perturbation to be truncated if the string is not monotonic in energy. These two values vary based on a granular system's parameters such as particle radii, number of particles, etc. and will vary for different sorts of PESs. It is at this point that the analysis described in Sect. 4.2 should be performed. Looking at the results of this analysis in Fig. 5, we need to decide what the characteristic radial size across all of these basins is. Each blue line corresponds to a single minimum. The characteristic radial size across all these basins is taken to be 0.15, or the maximum characteristic basin size of all the minimum (i.e. for no minimum do all perturbed points return to the original minimum).

5.2.1 $dist_{ext}$

A good choice is to pick the extension distance as the characteristic radial size across all of the basins. In this case, from Fig. 5 we have that $dist_{ext} = 0.15$. Beyond this perturbation distance, a rapidly increasing fraction of perturbed points fail to return to the original minimum, which would require truncating the initial string for an increasingly large fraction of initial perturbations. One could choose larger distances provided n_{images} is chosen appropriately. If there are enough images, the basin is resolved well enough to locate its edge. If a string is initially extended beyond the original basin, it will be truncated leaving enough images in the original basin so that the string can be re-interpolated and continue climbing.

5.2.2 n_{images}

The number of images, n_{images}, should be chosen large enough to properly approximate the MEPs. However, too many images will lead to an increase in the time spent running the algorithm. For larger systems, this can quickly become computationally prohibitive. We recommend choosing n_{images} such that the string resolves the smallest features of interest in the landscape. Figure 5 shows that there are larger and smaller basins, so sending out a string the distance of the maximum characteristic basin size may miss these smaller basins. To prevent this, one should interpolate many images along the string to capture these smaller features. In this case, we chose $n_{images} = 10$ so that an image falls inside these smaller basins; we potentially ignore some of the very small basins.

In this case, most of the basins have approximately the same structure. That is, they are relatively flat and spherical up until about 0.15. If instead there were many basins with different structures, one may need to have more images to accurately capture the complicated structure that appears on the ridges of these basins. Our basins are quite uniform and so we can choose a small number of images without missing many small features. We may not accurately resolve the MEP but we are confident we find the minima and saddles of interest to our analysis. If one desires a more accurate MEP, they should interpolate many more images to resolve it more finely. Additionally, the number of images can be dynamically changed and clustered around regions of interest. For example, the number of images near the climbing image can be increased to obtain a more accurate tangent approximation.

The Interplay between $dist_{ext}$ and n_{images} Since the energy of the climbing string at each image is kept monotonic by cutting the string if it ever fails to maintain monotonicity, the $dist_{ext}$ can be chosen to be larger than the characteristic radial size of the smallest basin as long as the string resolves the smallest basins of interest since the string will be cut if it fails to remain monotonic in energy. For example, the $dist_{ext}$ can be chosen to be 0.2 instead of 0.15. We can then interpolate more images, say $n_{images} = 15$, and if the string leaves a smaller basin, the string should detect that it has left the basin, because the energy fails to be monotonic, and cut itself. Overestimating the $dist_{ext}$ allows the strings to quickly leave the region with a small gradient and having a large n_{images} prevents the over-extended string from leaving the basin.

5.3 Parameters to make a network of saddles

The Hydra String Method generates a network where saddles and minima are nodes and the minimum energy paths between them are the edges. To generate this network, we need to decide which minima and saddles are unique from one another (as discussed in Sect. 4.1) and we need to find the saddles between these minima.

5.3.1 $tol_{sad/min}$

Often, the Hydra String Algorithm finds minima and saddles that are very near to each other in euclidean space with nearly exactly the same energies. The question arises, are these differences from numerical tolerances or are they in fact distinct points? To answer this we proposed the numerical experiment in Sect. 4.1 to determine a numerical tolerance to call these granular configurations degenerate.

The results of such a numerical simulation are shown in Fig. 4. Due to the number of saddles included, we zoomed in to clearly see the important features. The region with pairwise distances between $10^{-1.5}$ and 10^{0} are the obviously distinct saddles which can be seen by plotting the two configurations for a given pair in this region. The region between $10^{-2.5}$ and 10^{-8} are the obviously degenerate saddles. Again this can be seen by plotting both configurations for any given pair in this region and noting they overlap almost perfectly. From this, we chose $tol_{sad/min}$ to be 10^{-2}. All configurations closer than 10^{-2} in euclidean distance would be considered degenerate.

This histogram is likely to appear wildly differently for different PESs. In the case of a granular material, these pairwise differences will change in magnitude based on the radii of the particles and the number of particles themselves, and thus the dimension of the system. The distance function begins to be less useful at higher dimensions and so different means to differentiate saddles and minima may be necessary for sufficiently large systems. One possible substitute for the case of granular materials is to look at the contact matrices of the constituent particles and defining configurations with different contact matrices to be different. We propose using

the distance between configurations as the deciding factor instead of the contact matrix because recent work with soft spheres found that contact changes do not always correspond to a saddle point in the energy landscape [34].

It is important to note how variable this histogram may appear because this numerical experiment requires the user to apply some knowledge about the PES and the physics of the system that gives rise to that PES to interpret this plot and choose a value for $tol_{sad/min}$. In our application, we are not overly concerned with precisely finding the saddles or minima of the granular system. Rather we prefer to find many saddles and minima quickly to map out the PES and network of connected low energy extrema. Systems described by different physical laws or different applications may require different considerations.

5.3.2 $n_{strings}$

The number of strings to be sent out from each minimum to find new saddles should be large enough to find a statistically representative sample of the saddles connected to a minimum but small enough that not many of the strings converge to the same saddle. The analysis described in Sect. 4.3 should be performed to help choose the value of this parameter. The number of cores on the system should also be considered when choosing a value of $n_{strings}$, due to the parallel nature of the snaking algorithm, a multiple of the number of cores may make the execution more efficient.

Following the method outlined in Sect. 4.3, five histograms are shown in Fig. 6 generated for 20, 40, 60, 80 and 100 strings. The histograms might immediately indicate how many unique saddles exist on the ridge of a basin. However, we have found that as more strings are extended, there are generally more saddles to be found. We instead find it useful to look at the efficiency of these strings which we define as the fraction of strings that find unique saddles. This efficiency is plotted in Fig. 7. Choosing an efficiency of 50% (half the maximum) would dictate $n_s = 20$. Our computing hardware has nodes with 12, 24 or 36 cores, so we chose $n_{strings} = 24$ to take advantage of the parallelizability of the Hydra String Method.

One may also want to apply an analysis that reveals how good of a sample of the surrounding saddle points has been obtained relative to some physically meaningful quantity to determine the number of extended strings. For example, has a good sampling of possible energy barriers between the originating minima and saddle points been obtained? We plot the running average energy barrier of all the unique found saddles as a function of the number of extended strings in Fig. 8 for a few different originating minima. The flatness of the line is a rough estimate of the error of the average; a relatively flat line indicates enough samples have been obtained. From this figure, we can see that the average energy barrier found levels off at different values for different originating minima, but overall, the lines level off around 30 or so extended strings. As before, multiples of 12 are convenient so we might choose $n_{strings} = 36$ in this case.

In this system, the histograms from Fig. 6 each appear largely uniformly distributed. However, we have also analyzed systems with bi-modal or heavily tailed distributions. So, one should look at these underlying histograms before immediately creating the efficiency plot to find a value for $n_{strings}$ as one may need a different measure of the average than a simple mean to effectively create this efficiency plot.

One must also consider what is more important in their analysis of their PES. Is it more important to find every connection between all minima of the PES? Or does one favor exploring more of the PES at the cost of missing a few saddles and minima? In our case, we were not concerned with finding every hard to find saddle but we do want to find the many possibly interconnected saddles between minima. So, we chose an intermediate goal, finding many but not all minima and the connecting saddles.

6 Discussion

Choosing the various parameters in the Hydra String Method allows it to be adapted to different goals such as short range accuracy or long range exploration. Thus one can choose if it is more important to painstakingly map out every minimum and saddle point in your network? Or is it more important to find many minima and saddle points far from the initial point from which the Hydra String Method begins? These conflicting goals appear several times in the Hydra String Method. Does one choose a short extension distance to ensure no

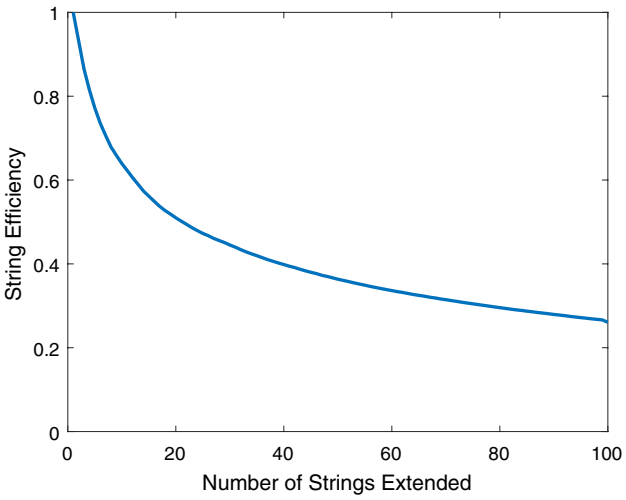

Fig. 7 This plot shows the efficiency of the strings sent to find unique saddles around a basin. The y-axis is what fraction of strings converged to new saddles and the x-axis is how many strings were extended. This is an average over 750 minima

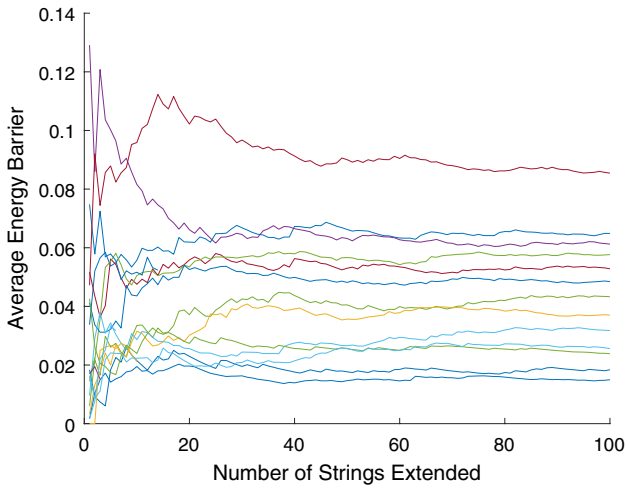

Fig. 8 Each line is for a different originating minima; it shows the average energy barrier between this minimum and the unique found saddles on the edge of the basin of attraction as a function of the number of extended strings. Only approximately 20 lines of the 750 are shown for clarity

nearby saddles are missed? Or does one pick a large value to quickly find further away saddles? Does one extend a large number of strings to find as many connected saddles as possible? Or a few to find the most common saddle points and move further into the PES? And so on.

In our case, we are analyzing the PES of a collection of soft spheres to map out various transition pathways the system can undergo. That is, starting in a stable configuration, what are the nearby stable configurations and what unstable configurations, saddle points, do they pass through as the system moves between these stable configurations? In this application, we favor exploring more of the PES instead of finding every possible minimum or saddle point. Many of those configurations may be unlikely for the granular system to reach and are therefore unimportant to our future analyses.

However, for other systems the Hydra String Method might be applied to, it may be more important to locate every minimum or saddle. For example, when studying a chemical system, one might be interested in determining various by-products or possibly dangerous intermediate products of a chemical reaction. In that case, it might be more important to find every possible minimum or saddle point.

7 Conclusion

In this paper we presented the Hydra String Method, a novel computational method to autonomously and efficiently map the minima and first order saddle points of a PES. In doing so, we presented a systemic approach to tailor the various system specific parameters of the HSM to arbitrary systems and demonstrated this approach on an example soft-sphere granular system. The results of the above numerical experiments may be very different for other systems with different potential functions, such as a pairwise Lennard-Jones potential between particles, commonly used in bubbles [35] and as a model for molecules/atoms in chemistry [36]. This broad applicability is one of the major assets of this method.

In our case, we intend to apply this method to granular systems to map out the various stable configurations and determine the MEPs that connect them. We believe that, these MEPs approximate the transition paths the system undergoes when sheared slowly and with sufficient damping. We hypothesize that the transitions with the lowest energy barriers will be the most likely transitions the system will undergo when slowly sheared with sufficient damping.

Since this method utilizes the String Method, it inherits all of the benefits that the String Method has over other saddle finding schemes. Of note are the modular re-interpolation of images along the string and the improved stability. As previously discussed, the interpolation of the images along the string can be easily changed from a linear interpolation to a cubic spline or any other method. If desired, the images can be adaptively added to further refine highly serpentine paths, removed to save computational costs, or redistributed to be "bunched up" around areas where the landscape may have more intricate structure or "thinned out" in relatively barren regions.

Finally, the advantage this method gains from its highly parallelizable nature and autonomous execution can hardly be overstated. The method efficiently realizes computational speed ups with growing numbers of parallel cores. The time spent climbing and descending strings is much greater than the time spent with overhead memory transference and using a list of unexplored minima to direct the workers mitigates wasted computation time spent searching already explored regions of the PES. The upfront time spent in picking suitable parameters for the method is quickly recouped in the autonomous execution of the search. With some intuitive precautions such as a maximum number of time steps allowed on the climbing/descending functions (to prevent strings from becoming stuck) and a maximum allowed energy along a string (to prevent run away strings) the Hydra will happily branch out and explore any energy surface.

Acknowledgements The authors thank Eric Vanden-Eijnden for discussions on the implementation of the string method as well as Karen Daniels, Ryan Kozlowski and Jack Featherstone for their feedback on an earlier draft of this manuscript.

Funding This work was supported by the National Science Foundation DMS-1816394.

Data availability All processed data is presented in the manuscript figures.

Code availability https://github.com/knewhall/Hydra_String.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

1. Franklin, S.V., Shattuck, M.D.: Handbook of Granular Materials. Taylor & Francis Ltd., Milton Park (2016)
2. Jaeger, H.M., Nagel, S.R., Behringer, R.P.: Granular solids, liquids, and gases. Rev. Mod. Phys. **68**(4), 1259–1273 (1996)
3. Richard, P., Nicodemi, M., Delannay, R., Ribière, P., Bideau, D.: Slow relaxation and compaction of granular systems. Nat. Mater. **4**(2), 121–128 (2005)
4. van Hecke, M.: Jamming of soft particles: geometry, mechanics, scaling and isostaticity. J. Phys. Condens. Matter **22**(3), 033101 (2009)
5. Gao, G.-J., Blawzdziewicz, J., O'Hern, C.S.: Geometrical families of mechanically stable granular packings. Phys. Rev. E **80**(6), 061303 (2009)
6. Ashwin, S.S., Blawzdziewicz, J., O'Hern, C.S., Shattuck, M.D.: Calculations of the structure of basin volumes for mechanically stable packings. Phys. Rev. E **85**(6), 061307 (2012)
7. Tuckman, P.J., VanderWerf, K., Yuan, Y., Zhang, S., Zhang, J., Shattuck, M.D., O'Hern, C.S.: Contact network changes in ordered and disordered disk packings. Soft Matter **16**(41), 9443–9455 (2020)
8. Basu, A., Ye, X., Tim Still, P.E., Arratia, Z.Z., Nordstrom, K.N., Rieser, J.M., Gollub, J.P., Durian, D.J., Yodh, A.G.: Rheology of soft colloids across the onset of rigidity: scaling behavior, thermal, and non-thermal responses. Soft Matter **10**(17), 3027 (2014)
9. D'Anna, G., Mayor, P., Barrat, A., Loreto, V., Nori, F.: Observing Brownian motion in vibration-fluidized granular matter. Nature **424**(6951), 909–912 (2003)
10. Maiti, M., Schmiedeberg, M.: Ergodicity breaking transition in a glassy soft sphere system at small but non-zero temperatures. Sci. Rep. **8**(1), (2018)
11. Maiti, M., Schmiedeberg, M.: The thermal jamming transition of soft harmonic disks in two dimensions. Eur. Phys. J. E **42**(3), (2019)
12. Markus B.-B., and Andreas H.: Shearing small glass-forming systems: a potential energy landscape perspective. Phys. Rev. E **98**(3), (2018)
13. Mungan, M., Sastry, S., Dahmen, K., Regev, I.: Networks and hierarchies: How amorphous materials learn to remember. Phys. Rev. Lett. **123**(17), 178002 (2019)
14. Weinan, E., Ren, W., Vanden-Eijnden, E.: String method for the study of rare events. Phys. Rev. B **66**(5), 052301 (2002)
15. Ren, W., Vanden-Eijnden, E.: A climbing string method for saddle point search. J. Chem. Phys. **138**(13), 134105 (2013)
16. Henkelman, G., Jónsson, H.: Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. J. Chem. Phys. **113**(22), 9978–9985 (2000)
17. Henkelman, G., Uberuaga, B.P., Jónsson, H.: A climbing image nudged elastic band method for finding saddle points and minimum energy paths. J. Chem. Phys. **113**(22), 9901–9904 (2000)
18. Henkelman, G., Jónsson, H.: A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. J. Chem. Phys. **111**(15), 7010–7022 (1999)
19. Olsen, R.A., Kroes, G.J., Henkelman, G., Arnaldsson, A., Jónsson, H.: Comparison of methods for finding saddle points without knowledge of the final states. J. Chem. Phys. **121**(20), 9776–9792 (2004)
20. Heyden, A., Bell, A.T., Keil, F.J.: Efficient methods for finding transition states in chemical reactions: comparison of improved dimer method and partitioned rational function optimization method. J. Chem. Phys. **123**(22), 224101 (2005)
21. Barkema, G.T., Mousseau, N.: Event-based relaxation of continuous disordered systems. Phys. Rev. Lett. **77**(21), 4358–4361 (1996)
22. Malek, R., Mousseau, N.: Dynamics of Lennard–Jones clusters: a characterization of the activation–relaxation technique. Phys. Rev. E **62**(6), 7723–7728 (2000)
23. Jay, A., Huet, C., Salles, N., Gunde, M., Martin-Samos, L., Richard, N., Landa, G., Goiffon, V., De Gironcoli, S., Hémeryck, A., Mousseau, N.: Finding reaction pathways and transition states: r-ARTn and d-ARTn as an efficient and versatile alternative to string approaches. J. Chem. Theory Comput. **16**(10), 6726–6734 (2020)
24. Müller, K.: Reaction paths on multidimensional energy hypersurfaces. Angewandte Chemie International Edition in English **19**(1), 1–13 (1980)
25. Vanden-Eijnden, E.: Transition Path Theory, pp. 453–493. Springer, Berlin, Heidelberg (2006)
26. Perry, R.W., Holmes-Cerfon, M.C., Brenner, M.P., Manoharan, V.N.: Two-dimensional clusters of colloidal spheres: ground states, excited states, and structural rearrangements. Phys. Rev. Lett. **114**(22), 228301 (2015)
27. Yucen, H., Jianyuan, Y., Pingwen, Z., Apala, M., Lei, Z.: Solution landscape of a reduced landau-de gennes model on a hexagon. arXiv (2020)
28. Doye, J.P.K., Massen, C.P.: Characterizing the network topology of the energy landscapes of atomic clusters. J. Chem. Phys. **122**(8), 084105 (2005)
29. Doye, J.P.K.: Network topology of a potential energy landscape: a static scale-free network. Phys. Rev. Lett. **88**(23), 238701 (2002)
30. Schön, J.C.: Energy landscape of two-dimensional lattice polymers. J. Phys. Chem. A **106**(45), 10886–10892 (2002)
31. Mills, G., Jónsson, H., Schenter, G.K.: Reversible work transition state theory: application to dissociative adsorption of hydrogen. Surf. Sci. **324**(2–3), 305–337 (1995)
32. Majmudar, T.S., Behringer, R.P.: Contact force measurements and stress-induced anisotropy in granular materials. Nature **435**(7045), 1079–1082 (2005)
33. Zadeh, A.A., Barés, J., Brzinski, T.A., Daniels, K.E., Dijksman, J., Docquier, N., Everitt, H.O., Kollmer, J.E., Lantsoght, O., Wang, D., Workamp, M., Yiqiu Z., Hu, Z.: Enlightening force chains: a review of photoelasticimetry in granular matter. Granular Matter **21**(4) (2019)
34. Morse, P., Wijtmans, S., van Deen, M., van Hecke, M., Manning, M.L.: Differences in plasticity between hard and soft spheres. Phys. Rev. Res. **2**(2), (2020)
35. Zhang, K., Kuo, C.-C., See, N., O'Hern, C., Dennin, M.: Stable small bubble clusters in two-dimensional foams. Soft Matter **13**(24), 4370–4380 (2017)
36. Chill, S.T., Stevenson, J., Ruehle, V., Shang, C., Xiao, P., Farrell, J.D., Wales, D.J., Henkelman, G.: Benchmarks for characterization of minima, transition states, and pathways in atomic, molecular, and condensed matter systems. J. Chem. Theory Comput. **10**(12), 5476–5482 (2014)