

From canyons to valleys: Numerically continuing sticky-hard-sphere clusters to the landscapes of smoother potentials

Anthony Trubiano and Miranda Holmes-Cerfon 

Courant Institute of Mathematical Sciences, New York University, New York, New York 10012, USA



(Received 26 August 2019; accepted 25 March 2020; published 27 April 2020)

We study the energy landscapes of particles with short-range attractive interactions as the range of the interactions increases. Starting with the set of local minima for $6 \leq N \leq 12$ hard spheres that are “sticky,” i.e., they interact only when their surfaces are exactly in contact, we use numerical continuation to evolve the local minima (clusters) as the range of the potential increases, using both the Lennard-Jones and Morse families of interaction potentials. As the range increases, clusters merge, until at long ranges only one or two clusters are left. We compare clusters obtained by continuation with different potentials and find that for short and medium ranges, up to about 30% of particle diameter, the continued clusters are nearly identical, both within and across families of potentials. For longer ranges, the clusters vary significantly, with more variation between families of potentials than within a family. We analyze the mechanisms behind the merge events and find that most rearrangements occur when a pair of nonbonded particles comes within the range of the potential. An exception occurs for nonharmonic clusters, i.e., those that have a zero eigenvalue in their Hessian, which undergo a more global rearrangement.

DOI: [10.1103/PhysRevE.101.042608](https://doi.org/10.1103/PhysRevE.101.042608)

I. INTRODUCTION

Metastable states of a system of interacting particles determine much of the system’s behavior, yet they can be difficult to study because they are sensitive to the interaction potential between the particles [1–3]. For mesoscale particles such as colloids, the interaction potential is not always well known because it depends on a combination of factors that occurs on a much smaller scale than the particles, such as electrostatic interactions, van der Waals interactions, the presence of impurities in solution, complex surface interactions created by tethered polymers, and other physical effects. Experimentally, the interaction potential is hard to measure because the particles typically interact over a distance much smaller than their diameters [4]. To model attractions between such particles, one typically chooses an interaction potential from a canonical family such as Morse, Lennard-Jones, or square-well potentials, and chooses parameters to fit aspects of the experimental data. Yet, even for these families of potentials, it is not known how sensitive the metastable states are to the choice of potential or parameters, nor how the metastable states for different choices are related to each other.

Conveniently, it has been shown that when particles have short- or even medium-range attractive interactions, aspects of their phase behavior are insensitive to the exact shape of the interaction potential, but rather depend on a single parameter characterizing the potential, i.e., the second virial coefficient [5]. The same is true of the set of metastable states, provided the range is short enough [6–8]. This observation has motivated the study of the energy landscape in the *sticky limit* when the range of the potential goes to zero and the depth goes to infinity, in such a way that the partition function approaches a δ function at the point of contact [9,10]. In

this limit, the metastable states of a system of N identical spherical particles are the set of sphere packings that have the most pairs of spheres in contact. Finding these sticky-hard-sphere (SHS) clusters is a problem in geometry that has been addressed using several techniques, both analytical and numerical [11–14], and the resulting data have given insight into a variety of physical properties of mesoscale particles [6,7,15,16]. However, real experimental colloidal systems do not always lie close enough to the sticky limit for it to be quantitatively accurate, and discrepancies from the predictions of the sticky limit have been observed even for systems as small as $N = 8$ particles [6].

We seek to understand how sensitive the metastable states are to the choice of potential when a system is near, but not exactly at, the sticky limit. Starting with the sticky-sphere landscape, which is thought to be the most rugged and to contain the most local minima, we apply numerical continuation to follow local minima as we slowly increase the range within a family of potentials for systems of $6 \leq N \leq 12$ spheres. This procedure finds most of the local minima for smooth potentials and, in particular, all the known deep local minima. We compare clusters that come from the same SHS cluster using different potentials, and find that clusters are similar for short ranges, up to about 30% of particle diameter, but vary significantly for longer ranges. We keep track of bifurcation events, where local minima split, merge, or disappear, and show that most bifurcations involving rearrangements occur when a pair of nonbonded particles comes within the range of the potential. An exception is bifurcations involving *nonharmonic* SHS clusters (those whose Hessian has a zero eigenvalue which does not extend into a finite floppy mode), which undergo a more global rearrangement whose location cannot be predicted from the starting SHS cluster.

Our study builds on others that have examined how energy landscapes vary as the range of the pair potential is varied. Wales [17] argued that catastrophe theory gives a quantitative relationship between local minima and the nearest saddle points when close to a bifurcation, and empirically showed this relationship holds reasonably well even away from the bifurcation. Trombach *et al.* [18] performed a local optimization in a Lennard-Jones(m, n) potential with varying m, n (varying range) at fixed energy, using SHS clusters as an initial condition for the optimization, and found most of the local minima on the Lennard-Jones landscapes; they showed the ones not found were from a small set of initial “seeds.” Trombach *et al.* [19] followed a similar approach to study the “kissing problem,” which asks how to arrange 12 spheres on the surface of a central sphere, in a family of Lennard-Jones potentials. The latter two approaches are the closest to ours; however, these studies performed a one-step optimization for each value of range, and hence could only compare the number of clusters found. In contrast, we vary the range parameter slowly, using the previously found cluster as the next initial condition, so we can additionally find and study bifurcations.

II. METHODS

We begin with a set of SHS clusters that is thought to be nearly complete, and likely missing only a small number of high-energy, nonharmonic clusters [14]. This set was produced using an algorithm to enumerate clusters that starts with a single rigid cluster, breaks a contact, and follows the one-dimensional transition path until a new contact is formed, producing a new rigid cluster. Iterating over all bonds and then all rigid clusters in the evolving list gives the dataset [14]. We consider how each of these clusters evolves as we slowly smooth out the pair potential into either a Morse or Lennard-Jones potential, given, respectively, by

$$U_M(r) = E(e^{-2\rho(r-d)} - 2e^{-\rho(r-d)}), \quad (1)$$

$$U_{LJ}(r) = \frac{E}{m} \left[m \left(\frac{d}{r} \right)^{2m} - 2m \left(\frac{d}{r} \right)^m \right]. \quad (2)$$

Here, r is the interparticle distance, $E > 0$ is the well depth, ρ and m are parameters governing the inverse range of the potential, and d is the equilibrium bond distance; we choose units so that $d = 1$.

We perform continuation on the set of SHS clusters as follows. We set the initial range parameters to be $\rho = m = 50$, and choose a corresponding energy parameter E . At each step of the continuation, we decrease the range parameter by 0.01, which slowly increases the range, and we update E in a way that we describe momentarily. We then minimize the potential energy (either U_M or U_{LJ}) under the new parameter values using the conjugate gradient algorithm, with the clusters obtained at the previous step as an initial condition (see Appendix A for details). These steps are repeated until the range parameter becomes 1.

Figure 1 illustrates the continuation procedure using a Morse potential for $N = 6$. There are two SHS clusters, which we will call the polytetrahedron and the octahedron. The




FIG. 1. Merging tree for $N = 6$ particles. The top row of the tree contains all SHS clusters (left = polytetrahedron; right = octahedron), and the tree follows unique clusters through the continuation process. The ordered pair (ρ, m) beneath the final cluster gives the range values for which the clusters merged for the Morse and Lennard-Jones potentials, respectively. The polytetrahedron suddenly rearranges and becomes identical to the octahedron. The octahedron remains the same throughout the continuation procedure. The topology of the trees is the same for each choice of κ and for each potential. Green (light) arrows denote a smooth transition between clusters; black arrows indicate a cluster rearranged nonsmoothly. Bars indicate an interparticle distance less than or equal to 1.

octahedron remains the same throughout the continuation process. The polytetrahedron remains nearly the same for range parameters $\approx \rho > 5$, with interparticle distances varying by less than about 3%. As ρ decreases below 5, the central bond in the polytetrahedron (between the blue and cyan particles) stretches slightly, and then the outer bonds (pink-yellow and red-green) stretch slightly too, by about 10%, as the red and pink particles come slightly closer together. At $\rho = 4.09$, the cluster suddenly rearranges, in one optimization step, when the red and pink particles come together to form the octahedron. For $\rho < 4.09$, there is only one Morse cluster, which is identical to the original octahedron.

During the optimization step, it is possible to reach a saddle point rather than a local minimum. This possibility is checked by computing the eigenvalues of a Hessian matrix. If a negative eigenvalue is found, a reoptimization procedure is performed in which the critical point is displaced in both directions along the corresponding eigenvector to obtain new starting points for the conjugate gradient algorithm. The algorithm could then produce two distinct local minima and we keep track of any such splitting events.

After constructing these lists of clusters, we compare each cluster pairwise to determine whether they are unique up to translations, rotations, and permutations (see Appendix B for details.) If two clusters are not unique, we say their “parent” clusters from the previous step have “merged.” For each family of potentials and each choice of energy parameters, we construct a bifurcation diagram showing how clusters merge and split as a function of the range parameter. It is this bifurcation diagram that we study in the remainder of the text.

During the continuation, we must decide how to vary the energy E as a function of the range parameters ρ, m . We do not believe this choice is terribly important to the results. We decided to try to keep the equilibrium distributions for systems along a single continuation path roughly comparable, so we chose E so that the partition function for a bond between a pair of particles on a line remains approximately constant. That is, define the one-dimensional pair partition function to be $\kappa = \int_0^{r_c} e^{-\beta U_{(r)}(r)} dr$, where r_c is a cutoff beyond which $U_{(r)} \approx 0$, and $\beta = (k_B T)^{-1}$ is the inverse of Boltzmann's constant times the temperature. The partition function κ is proportional to the time that a pair of one-dimensional particles in isolation spend within a distance of r_c of each other, when the system is in thermal equilibrium. The parameter κ is known as the *sticky parameter* because it measures how sticky the particles are—larger κ means the particles spend more of their time nearly in contact with each other [10]. It can be shown (see Appendix C) that the sticky parameter is a linear function of the second virial coefficient B_2 , which characterizes the thermodynamic properties of a system through the Law of Corresponding States [5].

The sticky parameter can be approximated using Laplace asymptotics for a potential with a deep and narrow attractive well as $\kappa = \sqrt{2\pi} e^{-\beta U_{(r)}(d)} / \sqrt{\beta U''_{(r)}(d)}$ (see Appendix C for details). Evaluating this expression for the Morse and Lennard-Jones potentials and nondimensionalizing the energy using units of β (so we may set $\beta = 1$ in the above formulas) gives

$$\kappa_M(\rho, E) = \sqrt{\frac{\pi}{E\rho^2}} e^E, \quad \kappa_{LJ}(m, E) = \sqrt{\frac{\pi}{Em^2}} e^E. \quad (3)$$

Although these expressions are meaningful only for very short-range potentials, we use them to determine the relationship between the range parameters ρ, m and the energy parameter E at all ranges. Notice that both of the parameters ρ, m measure the inverse range and they appear in the formulas above in the same way, so we will use these parameters interchangeably hereafter.

We perform continuation for each of U_M, U_{LJ} , and for each of three different values of the sticky parameter, $\kappa_{\text{LOW}} = 23.4$, $\kappa_{\text{MED}} = 49.5$, and $\kappa_{\text{HIGH}} = 100.4$. At each step of the continuation, we solve (3) for E using Newton's method. This gently increases E as ρ or m decreases (the range increases).

III. RESULTS

A. Completeness of the set of continued clusters

First we ask whether this continuation procedure produces all the local minima for a given landscape. We compare the set of Morse clusters obtained by continuation for $\rho = 30$ and $\rho = 6$ to the local minima found by a basin-hopping technique in [20]. The number of unique local minima in each set is given in Table I. Our method finds all local minima in the basin-hopping dataset for $N \leq 8$, but for larger N it misses a few. Upon inspection, the unmatched clusters are mostly high-energy clusters: each unmatched cluster has energy greater than $-(3N - 6)E$ and usually close to $-(3N - 7)E$, whereas a typical matched cluster has energy between $-(3N - 5)E$ and $-(3N - 6)E$. Figure 2 compares the energy distributions

TABLE I. Number of unique SHS clusters with N particles, C_N , as well as the number of Morse clusters found through the continuation procedure, $|\text{SHS} \rightarrow M_\rho|$, and the total number of Morse clusters, $|M_\rho|$, for range parameters $\rho = 6, 30$. The difference between the continued and complete sets, Δ , is also reported. The continued clusters were generated using sticky parameter κ_{MED} . The * indicates that a heuristic algorithm was used to determine the uniqueness of the clusters, described in the Appendix.

N	C_N	$ \text{SHS} \rightarrow M_{30} $	$ M_{30} $	Δ_{30}	$ \text{SHS} \rightarrow M_6 $	$ M_6 $	Δ_6
6	2	2	2	0	2	2	0
7	5	4	4	0	4	4	0
8	13	10	10	0	8	8	0
9	52	30	31	1	17	19	2
10	263	151	170	19	57	61	4
11	1659	866	1127	259	161	170	9
12*	11980	5684	8059	2375	489	506	17

of the clusters that we find at $\rho = 30$ and the basin-hopping data. A smaller fraction of clusters is missing at longer range: at $\rho = 30$, the method missed 11% and 23% for $N = 10, 11$, respectively, whereas for $\rho = 6$, the method missed 6.5% and 5.3%, respectively. The continuation procedure did not find any structures that were not present in the basin-hopping dataset.

Our results are comparable to those of Trombach *et al.* [18], who computed Lennard-Jones clusters with $m = 6$, $E = 1$ ($\kappa = 0.8$) by performing a one-step optimization from a SHS cluster. For $N = 10, 11$, their method failed to find 2/64 (3.1%) and 5/170 (2.9%) for $N = 10, 11$, respectively, i.e., slightly smaller numbers than ours. They also found that most missing clusters were high energy.

If missing clusters are high energy, this suggests that as the range increases, local minima are created on the flat, higher-energy parts of the sticky-sphere landscape, from




FIG. 2. Histogram of the cluster energies E_{cl} scaled by the well depth, E , for $N = 10$ and $N = 11$. We compare the energies of the clusters found at $\rho = 30$ for the Morse potential with $\kappa = \kappa_{\text{MED}}$ using continuation, against the set of all known clusters. Most of the clusters that are missed using continuation are high energy.




FIG. 3. (a) Potential energy of all clusters as a function of $\ln(\rho)$ for $N = (6, 7, 8)$ from top to bottom, respectively, and (b) the average symmetry number of all $N \leq 9$ clusters as a function of $\ln(\rho)$, as well as a scatter plot of the change in symmetry number for each individual cluster as a function of $\ln(\rho)$. Negative values for the scatter plot mean a cluster's symmetry decreased when it merged. The clusters used for these calculations were generated using the Morse potential and $\kappa = \kappa_{\text{MED}}$.

configurations corresponding to floppy clusters with one or more internal degrees of freedom. Such creation of local minima cannot be detected by our procedure. Because we obtain better agreement at longer ranges, we hypothesize that these high-energy local minima disappear at larger ranges. Typically, one is interested in low-energy minima, so we feel confident using our dataset going forward to understand bifurcations in the low-energy parts of the landscape.

As a brief application of our nearly complete data, we show how statistical and geometric properties of the clusters evolve as a function of the range. Figure 3 shows how the Morse potential energy and symmetry number (order of the point group) vary with range. The average potential energy decreases as the range of the potential increases, presumably because particles can interact attractively with neighbors that are farther away. The average symmetry number increases as the range increases: clusters for longer-ranged potentials are more symmetric, on average. Interestingly, the symmetry number does not always increase monotonically following a particular cluster; occasionally a cluster acquires lower symmetry as the range increases.

B. Visualizing bifurcations in the energy landscape

Next we examine bifurcations in the energy landscape and compare bifurcation diagrams for different potentials and parameters. There are two kinds of bifurcations that we can detect: merging events, when two or more local minima become the same cluster (like when the polytetrahedron and octahedron merge in Fig. 1), and splitting events, when one local minimum splits into two or more.

We find many merging events as the range parameter decreases. Interestingly, we find no nonisomorphic splitting events. Splitting events are possible when a cluster hits a saddle point in the optimization. This only happened when we tracked nonharmonic clusters, the smallest of which occurs at $N = 9$. These clusters hit a saddle point initially, and then continued to hit saddle points every so often until $\approx \rho < 30$. However, every time we hit a saddle point and searched both directions of the negative eigenvector, we always found two local minima that were the same up to a rigid rotation or




FIG. 4. Merging tree for $N = 7$, with Morse clusters plotted at nodes of the tree. The top row of the tree contains all SHS clusters and the tree follows unique Morse structures through the continuation process. The topology of the trees is the same for each choice of κ and for each potential. The ordered pairs (ρ, m) beneath each cluster give the range value for which the merge happened for Morse and Lennard-Jones potentials, respectively. Green (light) arrows denote a smooth transition between clusters; black arrows indicate a cluster rearranged nonsmoothly. Bars indicate an interparticle distance less than or equal to 1.

a permutation of the particle labels, and hence, which are identified as the same cluster. This result was unexpected—our original hypothesis was that nonharmonic clusters would lead to nontrivial splitting events—and we do not have an explanation for why we see none.

Because we find only merging events, our data can be represented as a graph with a tree structure. The top row contains all SHS clusters, and clusters which merge are connected at a branch in the tree, with a node representing the cluster they merge into. When two or more clusters merge into one, we call the clusters before the merge the “parents” and the cluster just after the merge the “child.”

The simplest merging tree is the one for $N = 6$ (Fig. 1.) The topology of the merging tree is the same for all potentials and all parameter values that we considered. The single merge event occurs at slightly shorter range for the Morse potential ($\rho = 4.09$) than for the Lennard-Jones potential ($\rho = 4.07$).

For $N = 7$, the merging trees continue to have the same topology for all potentials and parameters, while the range at which the merges occur depends on the potential (Fig. 4.) Lennard-Jones clusters usually merge at smaller values of the range parameter (longer range) than Morse clusters, although for the single merge at large range parameter ($\rho = 38.13, m = 38.23$), the Lennard-Jones cluster merged first. Similar




FIG. 5. Merging trees for the Morse potential with $N = 9$ particles. The top row of the trees contains all SHS clusters and the trees follow unique Morse structures through the continuation process. The black nodes and edges are the same for each κ value. The green nodes and edges are specific to the $\kappa_{\text{MED}} = 49.5$ tree. The blue nodes and edges are specific to the $\kappa_{\text{LOW}} = 23.4$ and $\kappa_{\text{HIGH}} = 100.4$ trees, which are (interestingly) the same. The nonharmonic cluster and its path are shown in red. Note that only the topology of the trees is being shown, i.e., the vertical positions of the nodes are not to scale.

observations hold for $N = 8$ (see Appendix E, Fig. 15): the trees are all topologically the same, but the ranges where merges occur differ slightly between the two potentials.

How do the clusters of $N = 7$ merge? In each merge, there is always one parent cluster that does not change and one that rearranges significantly. We speculate on the mechanism of the more dramatic rearrangement by inspecting the clusters. In the first merge in Fig. 4, the red and yellow particles of cluster 1 (counting clusters from the left) begin to interact with each other and pull toward the center of mass, pushing the central particles apart to form the more symmetric child cluster. This child cluster has a ring of outer particles that are slightly farther apart than they were in the smoothly varying parent SHS cluster, i.e., cluster 2, presumably because spreading apart the ring allows the central blue and brown particles to come closer together. The second and third merges are like the $N = 6$ merge: clusters 3 and 5 each contain a polytetrahedron, which suddenly rearranges into an octahedron, a substructure of cluster 4 that they each merge with. The merges occur at slightly different ranges; one reason could be that clusters 4 and 5 in the second merge both have six symmetries, so they are more similar to begin with than clusters 3 and 4 in the third merge, which have two and six symmetries, respectively. In the final merge, the square base of the rightmost cluster absorbs the red particle to become a pentagon. Overall, only one SHS cluster, i.e., cluster 2, evolves smoothly throughout the whole continuation process; this cluster happens to be the one with the most symmetries initially.

For $N = 9$, the topology of the merging trees varies with both parameters and potential. The Morse trees are identical for short range, but not for long range (Fig. 5). The first difference occurs at $\rho = 3$, where a single cluster merges with two distinct clusters depending on the value of κ . The SHS cluster as well as these two continuation possibilities are shown in Fig. 6. The SHS parent cluster is built from cluster 1 of $N = 7$ by adding two extra particles, red and pink, along nonadjacent edges of the pentagon. The leftmost child cluster looks similar and can be formed by pulling the red




FIG. 6. The SHS cluster 10 for $N = 9$ (top), as well as the two possible results of the continuation procedure at $\rho = 3$ for the Morse potential. Choosing $\kappa_{\text{MED}} = 49.5$ results in the green (right) path, while choosing $\kappa_{\text{LOW}} = 23.4$ and $\kappa_{\text{HIGH}} = 100.4$ results in the blue (left) path, consistent with the coloring in Fig. 5.

and pink particles toward the center of mass, close enough to bond. The rightmost child cluster is quite different, containing an octahedron fragment that the others do not. In addition to the red and pink particles being pulled closer together in this cluster, the dark-purple particle seems to get pushed away from the center of the cluster. This child cluster is the only remaining cluster when the range becomes 1, and has lower energy than the leftmost child cluster. We are not sure why the sticky parameter affects the result in this way or if there is physical intuition behind it. One possibility is that slight perturbations in the potential energy caused the optimization algorithm to find a deeper minimum.

For $N = 10, 11$, merging trees continue to depend on the sticky parameter and potential, although the upper portions of the trees are still independent of parameters. For both potentials, the trees are exactly the same for $\rho > 32$ and $\rho > 40$, respectively, though the differences for $\rho > 30$, corresponding to a range of about 8% of particle diameter, in both cases are minimal: a difference of between 1 to 10 nodes.

For these larger values of N , it is hard to compare the topology of the trees quantitatively, so instead we compare the geometry of the individual clusters. Suppose we wish to compare two sets of clusters at a given range, where each set is obtained by performing continuation with different well-depth parameters E or potential. There is a one-to-one mapping between the sets of clusters because each cluster comes from a unique SHS cluster. We use this mapping to compare clusters: for each pair that comes from the same SHS cluster, we compute the root-mean-square deviation (RMSD; see Appendix B for details.) We then sum the RMSD over all clusters to obtain a metric comparing the sets, which we call the total RMSD. We compute the total RMSD as a function of the range parameter ρ .

Figure 7 shows the total RMSD as a function of ρ for $8 \leq N \leq 11$, for clusters from different potentials or different κ (different well depths E , since the ranges at which they are




FIG. 7. Total RMSD between clusters plotted as a function of $\ln(\rho)$ for (a)–(d) $N = 8, 9, 10, 11$, respectively. Blue (dashed) curves compare low and medium values of κ for the Morse potential, red (dotted) curves compare low and medium values of κ for the Lennard-Jones potential, and yellow (solid) curves compare Morse and Lennard Jones with medium κ value. To interpret the magnitude of these curves, note that a typical RMSD between a pair of distinct clusters is about 1–2.

compared is the same). Remarkably, the total RMSD is nearly zero for all comparisons until relatively long ranges, roughly $\rho < 6$, or about 35% of particle diameter. This means that not only are the topologies of the merging trees nearly the same up to longer ranges, but the geometries of the clusters themselves are also nearly the same—in other words, for a fixed range, it does not matter whether you use a Morse or a Lennard-Jones potential, or what you choose for the well depth (within the limits we considered); the metastable states are nearly identical.

For longer ranges, the total RMSD increases the most for clusters from different potentials: the geometry of the metastable states is sensitive to the choice of potential. The total RMSD increases only a little bit for clusters from the same family of potential but with different well depths; here the clusters have a much more similar geometry. The differences could be caused by slight differences in geometry or by merging at slightly different ranges; we cannot tell the difference using this metric. Table II shows a more extensive comparison than is contained in the figure and supports the observation that there is more variety between families of potentials than within a single family.

Finally, we return to the observation that for merges at small N , exactly one parent transitioned smoothly, and we ask whether this holds true for larger N as well. We consider every merge event for $N \leq 10$ and compute the RMSD between each cluster just before and just after it merged. For each group of merging clusters, we find exactly one cluster with a small RMSD, and all the others have much larger RMSDs

TABLE II. Total RMSD summed over all integer values of ρ , for different comparisons. Left: Morse potentials compared at low (L), medium (M), and high (H) values of κ . Middle: Lennard-Jones potentials, also compared at different values of κ . Right: Morse and Lennard-Jones clusters compared at the same values of κ .

N	Morse			Lennard-Jones			Comparison		
	L-M	M-H	L-H	L-M	M-H	L-H	L	M	H
6	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0.7	0.7	0.7
8	0	0	0	0	0	0	4.1	4.1	4.1
9	0.6	0.6	0.1	0.6	0.1	0.6	12	11	12
10	14	13	21	21	22	18	125	128	132
11	67	65	63	74	78	76	702	713	711

(Fig. 8). Therefore, up to $N = 10$, there is a unique smoothly varying parent cluster for each merge event. This observation also suggests that the merges occur as fold bifurcations, in which a local maximum and local minimum annihilate, leaving no extrema. The annihilated local minimum then jumps abruptly in configuration space upon optimization past the bifurcation. The merges do not appear to occur as pitchfork bifurcations, in which two local minima separated by a local maximum smoothly coalesce into a single local minimum; such a bifurcation would give two smoothly varying parents.

C. Predicting merge events

Next we ask whether there is a geometric criterion that governs when clusters merge. Intuitively, we expect that a cluster rearranges when the range of the pair potential becomes comparable to the distance between a pair of nonbonded particles.

To make this hypothesis quantitative, we define a function $r_i(\rho)$ to be the minimum interparticle distance that is greater




FIG. 8. Scatter plot of RMSD values between a cluster before and after a merge at a given value of ρ , for all merge events for $N \leq 10$, with a log-log scale. For each group of merging clusters, the smallest RMSD is unfilled and the rest are filled in with blue. Each merge group has only one cluster with a small RMSD, with most being nearly 0, implying that all merge events occur by fold bifurcations.




FIG. 9. (a) $r_{13}(\rho)$, $r_9(\rho)$, the minimum interparticle distance greater than 1 for clusters 13 and 9 for $N = 8$. Note that $r_9(\rho)$ varies smoothly, but $r_{13}(\rho)$ jumps at $\rho = 38$, when the clusters merge and become identical to cluster 9. The corresponding SHS clusters (13 and 9) are plotted in (b) and (c), respectively, in a way that minimizes the root-mean-square difference between them. The main difference between the clusters is the distance between particles 1 and 2 and particles 3 and 4. In cluster 9, $r_9(\rho)$ measures the distance between particles 1 and 2, and in cluster 13, $r_{13}(\rho)$ measures the distance between particles 3 and 4. These clusters merge when particles 3 and 4 get close enough to bond.

than 1, for cluster i at range parameter ρ . Our hypothesis is that merges depend on $r_i(\rho)$ in some way.

To investigate how the function $r_i(\rho)$ behaves, we plot it as a function of ρ in Fig. 9 for two Morse clusters at $N = 8$, i.e., clusters 9 and 13 in [14]. The SHS clusters have $r_9(\infty) = 1.0515$, $r_{13}(\infty) = 1.0887$. As ρ decreases, $r_9(\rho)$ decreases smoothly past $\rho = 38$, whereas $r_{13}(\rho)$ jumps at $\rho = 38$, from $r_i(38+) = 1.07$ to $r_i(38-) \approx 1.03$, when the cluster rearranges and becomes identical to cluster 9.

This behavior is the same for all merges that occur at $\rho = 38$ for $7 \leq N \leq 10$. In each merge, the cluster i that rearranges has $r_i(\infty) = 1.0887$ and the cluster j that varies smoothly has $r_j(\infty) = 1.0515$. Just before the merge, the values are also the same across all merges: $r_i(38+) \approx 1.07$, $r_j(38+) \approx 1.03$. Interestingly, the cluster that transitions smoothly is always the one that has the largest minimum eigenvalue in the Hessian at $\rho = 50$.

One might guess from this data that $r_i(\infty)$, the minimum gap in the SHS cluster, directly determines the value ρ_c at which a cluster i first rearranges. Unfortunately, the story is not so simple: among clusters with $7 \leq N \leq 10$, there are 101 clusters with $r_i(\infty) = 1.0887$, with values of ρ_c ranging from 31.60 to 38.76. To show this nonuniqueness, a scatter plot of $r_i(\infty)$ vs ρ_c is shown in Fig. 14.

What is true is that $r_i(\rho_c+)$, the minimum distance for non-bonded particles just before the rearrangement, determines ρ_c . Figure 10 shows a scatter plot of ρ_c versus $r_i(\rho_c+)$ for all clusters i that rearrange, for both potentials. The data are very well fit by the curve $r_i(\rho_c+) = 2.1\rho_c^{-1.05} + 1.03$, which we obtained using nonlinear least squares to fit the data with a function of the form $ax^b + c$. Since the width of the attractive




FIG. 10. The minimum interparticle distance greater than 1 right before a cluster rearranges, $r_i(\rho_c+)$, vs the value ρ_c at which it rearranges for the first time, for both the Morse and Lennard-Jones potentials for $6 \leq N \leq 10$. The scatter plot only includes merge events for clusters that rearrange discontinuously, and such that the parent SHS clusters are harmonic.

well of the potential scales with ρ as $c\rho^{-1} + 1$, where c is a constant, this fit is strong support for the hypothesis that a cluster's first rearrangement occurs when the closest non-contacting pair comes within the range of the pair potential. The distance of this noncontacting pair can change during the continuation, which is why the distance in the SHS cluster does not determine the range at which the cluster rearranges (it does seem to determine it approximately, since the distance does not usually change too much during the continuation. See Appendix D). What is useful about our formula is that it gives a specific number with which to measure the width of the potential—it tells us that when the gap between particles is closer than $2.1/\rho$, their interaction starts to matter.

This formula does not hold for the smoothly transitioning parents, which comprise about 1/3 of the clusters and initially merge without rearranging. If we applied the formula to these clusters anyway, approximating $\rho_i(\rho_c+) \approx \rho_i(\infty)$, it would falsely predict a large ρ_c (short range for rearrangement.) Some of these smoothly transitioning parents do rearrange in later merges. We tried to find a relationship between $r_i(\rho)$ for these clusters and the value of ρ at which they first rearrange discontinuously, but we could not find any relationship.

Another exception to this behavior occurs for the nonharmonic clusters, which rearrange at much *shorter* ranges than predicted by the formula above. Every a nonharmonic SHS cluster for $9 \leq N \leq 11$ has $r_i(\infty) \approx 1.4142$, so the formula above would predict that they merge via rearrangement at $\rho_c \approx 5$; using actual distances when they rearrange, which are closer to $r_i(\rho_c) \approx 1.3$, gives $\rho_c \approx 7$. However, most nonharmonic clusters undergo a large rearrangement at $\rho_c \approx 15-20$, well before the minimum gap is within the range of the potential.

This suggests that nonharmonic clusters rearrange by a more global mechanism. To explore this mechanism, recall that nonharmonic clusters reach saddle points during the minimization for larger values of ρ and a reoptimization procedure




FIG. 11. Plot of the minimum eigenvalue of the Hessian of the Morse potential as a function of ρ for the $N = 9$ nonharmonic cluster. Plots of the SHS cluster taken from [8] are shown as well: the left plot shows particles with unit diameter, and the right plot has arrows on the sphere centers showing the zero mode of the Hessian.

is performed. The result of the reoptimization is a cluster that is structurally very similar to the starting nonharmonic cluster, with a nonzero but very small minimum eigenvalue. The cluster stays close to this configuration until $\rho \approx 15$ – 20 , when it rearranges and merges with harmonic clusters. As an example, Fig. 11 plots the minimum eigenvalue in the Hessian of the energy for the $N = 9$ nonharmonic cluster. As ρ decreases, the minimum eigenvalue slowly increases from 0 until a jump occurs near $\rho = 17$, at which point the cluster merges with a harmonic cluster. Similar behavior is exhibited for all $9 \leq N \leq 11$ nonharmonic clusters except for 4 (of 35) nonharmonic clusters at $N = 11$: they are nearly constant, with a small minimum eigenvalue, until they rapidly rearrange at $\rho \approx 15$ – 20 . The exceptions at $N = 11$ had a minimum eigenvalue that moved away from zero before the cluster merged.

We examine the nonharmonic cluster 6 for $N = 10$ in detail. This cluster, as well as most others, has a planar or near planar set of six particles that attach to each other in a ring and to a seventh central particle. This cluster stays nearly the same until $\rho \approx 30$, when the particles on the outer ring begin to




FIG. 12. Evolution of cluster 6 for $N = 10$, a nonharmonic cluster, during the continuation process. (a) Starting SHS cluster. Note the near planar set of particles surrounding particle 10. (b) Morse cluster at $\rho = 35$. The bonds between the outer ring of particles have been pushed apart. (c) Morse cluster at $\rho = 17$. The planar particles have been pulled down below particle 10 now. At this point, the cluster has merged with another cluster that was not initially nonharmonic. Gray bars represent interparticle distances less than or equal to 1 and blue (dark) bars represent interparticle distances less than $2.1\rho^{-1.05} + 1.03$.




FIG. 13. SHS clusters (a) 34 and (b) 58, for $N = 10$. When a bar is added to cluster 34 between particles 5 and 8, i.e., the two closest nonbonded particles, it has the same adjacency matrix as cluster 58. However, while these clusters both merge at $\rho = 38$, they do not merge together.

separate. This outer ring then begins to be pulled downward until $\rho \approx 17$, where the cluster rearranges and merges with another cluster. Various stages of this process are shown in Fig. 12. This general rearrangement mechanism occurred for most of the nonharmonic clusters.

Despite being able to often predict *when* a cluster will merge based on its geometry, we have not found a way to predict *which clusters* will merge together. One idea was to compute an adjacency matrix for each cluster representing the pairs of particles whose distances are comparable to the minimum distance between nonbonded particles, and then to compare adjacency matrices. This does not work unfortunately; a counterexample is shown in Fig. 13.

IV. CONCLUSION

We used numerical continuation to study the evolution of sticky-hard-sphere clusters as the range of interaction increased, for Morse and Lennard-Jones potentials. This procedure finds most local minima of the smoother energy landscapes; the relatively few unmatched clusters were higher-energy clusters. This suggests that a similar technique could be used to find deep minima of larger systems, since SHS clusters with a maximal number of contacts can sometimes be

found theoretically [21,22], but exploring short-range energy landscapes numerically is a challenge because the potential develops very high gradients.

As the range of interaction increased, distinct clusters merged together, so the total number of unique clusters decreased. Interestingly, we found no nontrivial splitting events, where one cluster split into two unrelated by rotations or permutations. We analyzed the mechanism behind individual merge events in details for small clusters. Many larger clusters contain these smaller clusters as subunits, so these simple merge events provide insight into how larger clusters merge.

We found that all merges involved one cluster that varied smoothly, hence whose structure did not change, while the other clusters rearranged significantly during the merge. This suggests that all merges in our data were fold bifurcations. In addition, we found that the clusters that rearranged did so when the range of the interaction potential became comparable to the minimum distance between particles that were not yet bonded. We found a specific formula to measure the relevant range: when the pairwise distance became about $1 + 2.1/\rho$, these clusters rearranged. This formula may be useful for simulations or experiments that wish to use a threshold distance for saying particles are bonded.

An exception to this formula occurred for the nonharmonic clusters, which rearranged by a more global mechanism.

We compared sets of clusters obtained by continuation for different potentials and different parameters, and found the corresponding clusters to be nearly identical for short and medium ranges, i.e., roughly 30% of particle diameter, but varied quite a lot for longer ranges. There was more variation when we changed the family of potentials than when we changed parameters within the same family of potentials.

Our observations show that for short-range interactions up to about 8% of particle diameter, the exact choice of pair potential and parameters has a negligible effect on the number of accessible ground states and the structure of local minima on the energy landscape; however, these states do differ from SHS clusters beyond a range of about 5% (for the values of N considered.) For longer range potentials, greater than about 30% of particle diameter, the particular choice of interaction potential can affect the structure of these states. An intriguing extension would be to study the dynamics as the range increases and to find a continuation procedure to study transition rates and paths between local minima or other dynamical quantities. Are they also similarly insensitive to the choice of potential?

ACKNOWLEDGMENTS

The authors would like to thank John Morgan for sharing data, Maria Cameron for providing code to determine point group order, and Dennis Shasha for helpful discussions. This work was supported by the US Department of Energy under Award No. DE-SC0012296, and the Research Training Group in Modeling and Simulation funded by the National Science Foundation via Grant No. RTG/DMS - 1646339. M.H.-C. acknowledges support from the Alfred P. Sloan Foundation.

APPENDIX A: OPTIMIZATION ALGORITHM

The conjugate gradient algorithm [23] was used to minimize the potential energy. For N particles, the potential energy is a function of $3N - 6$ position variables. There are six degrees of freedom corresponding to rigid body translation and rotations, which are removed by constraining particle 1 to the origin, particle 2 to the x axis, and particle 3 to the x - y plane. The conjugate gradient algorithm terminates when the change in the objective function is less than 10^{-13} and it is verified that the point is a local minimum via the eigenvalues of the Hessian. The condition number of the minimization problem scales with E , so the convergence rate scales like $1 - \frac{1}{\sqrt{E}}$ for $E \gg 1$. However, for large enough E , $E \sim \ln \kappa$, so the convergence rate does not change much with sticky parameter. This is consistent with what we see in practice.

The conjugate gradient method is unstable from some starting points and can blow up or take very large steps. This depends on the value of the sticky parameter, as a larger sticky parameter corresponds to a deeper well depth, which results in steeper gradients. A check for instability is performed after every optimization and if either possibility occurs, we reset and try one of a variety of methods. The methods are, in order of application, gradient descent followed by conjugate gradients, conjugate gradients with resets every $3N - 6$ iterations, swapping two random particle labels among the last $N - 3$ particles and then applying conjugate gradients, or perturbing the starting point by a random vector of norm 10^{-12} and applying conjugate gradients. If all of these methods fail, the starting point is logged as the minimum. This usually results in a point with potential gradient norm $10^{-7}E$, instead of the usual tolerance of $10^{-13}E$. The fraction of optimizations that results in such an error is (5.6, 3.5, 1.5, 0.74, 0.6)% for $6 \leq N \leq 10$.




FIG. 14. Scatter plot of the minimum distance between non-bonded particles in a SHS cluster, $r_i(\infty)$, vs the range value where the cluster merges, ρ_c , for all harmonic clusters with $6 \leq N \leq 10$. Blue points show clusters which merge smoothly; red points show clusters which merge by rearranging. The blue curve is fit by $1.03\rho_c^{-0.43} + 0.83$, and the red curve is fit by $1.45\rho_c^{-0.63} + 0.94$. The yellow curve shows the fit for $r_i(\rho_c+)$, which sits between the red and blue curves.




FIG. 15. The merging tree for a system with $N = 8$ particles. The top row of the tree contains all SHS clusters and the tree follows unique Morse clusters through the continuation process. The trees are the same for each choice of κ . The values ρ and m in each node give the value of the range parameter for which the merge occurred.

When a saddle point is reached (minimum eigenvalue becomes negative), a reoptimization procedure is applied to reach a local minimum. This involves displacing along the eigenvector corresponding to the negative eigenvalue and reapplying conjugate gradients until a minimum is found. In some cases, a saddle point with more than one negative eigenvalue is reached; usually only 2, but occasionally more. We found that the choice of eigenvector to displace along did not affect the result of reoptimization, so the eigenvector corresponding to the eigenvalue of largest magnitude is chosen for consistency.

A C++ implementation of this continuation procedure using DLIB [24] and additional code used (merge detection, RMSD, etc.) is available on GitHub [25].

APPENDIX B: TESTING UNIQUENESS

To determine when two clusters are the same, we begin by checking that they have the same set of interparticle distances. If so, the Kabsch algorithm [26] is applied to compute an

optimal rotation of one cluster onto the other, and the root-mean-square deviation is computed as

$$\text{RMSD} = \left(\frac{1}{N} \sum_{i=1}^N D_i^2 \right)^{0.5},$$

where D_i is the Euclidean distance between particle i in cluster one and cluster two. If this is less than a tolerance of 10^{-6} , we consider the clusters the same.

If not, we check possible permutations. To do so efficiently, we begin by computing the moment of inertia tensor, diagonalizing it, and using it to rotate each set of particles to the principal axes. Then we need to realign (permute) the particles so both clusters have the same configuration. We do this by solving a minimum cost assignment problem, where the cost matrix is the Euclidean distance between particle i in cluster one and particle j in cluster two. We then reapply the first part of the algorithm.

This still does not account for certain reflectional and rotational symmetries, so we repeat this procedure for all




FIG. 16. Cluster images of the major merge events for the $N = 8$ system, for values of ρ near (a) $\rho = 38$, (b) $\rho = 6$, (c) $\rho = 4$, and (d) $\rho = 2$. The parent clusters are SHS clusters. Green (light) arrows indicate smooth transitions.

swaps and sign combinations of (x, y, z) coordinates, i.e., 48 possibilities. Taking the minimum value among all these possibilities gives the RMSD.

In the case of 12 particles, this procedure becomes too slow to compare all pairs of clusters. In this case, we adopt a heuristic approach where we simply compare a sorted list of interparticle distances between clusters. This is a necessary condition for uniqueness but may not be sufficient. Up to $N = 11$, the heuristic approach gives the same results as comparing RMSD values, so we are hopeful this extends to $N = 12$.

APPENDIX C: STICKY PARAMETER

Here we discuss the origin of the sticky parameter κ and its relation to the second virial coefficient B_2 . In our numerical experiments, we seek to vary the range ρ of potentials, $U(r)$, in our continuation procedure. This leaves a free parameter in the potential function E , the well depth. Keeping E fixed as the range varies changes properties of the potential function at every stage. One such property is how “strong” a bond between two nearby particles is, which is given by the partition function

$$Z = \int_0^{r_c} e^{-\beta U(r)} dr, \quad (C1)$$

where r_c is a cutoff beyond which the pair potential is essentially constant. Thus a natural way to choose E after changing the range is to pick E such that $Z = \kappa$ remains constant. To approximate this integral, we can use the method of Laplace

asymptotics. The potential function has a minimum at $r = d$, such that $U'(d) = 0$, and $U''(d) > 0$. We expand the potential function in a second-order Taylor series to get $U(r) \approx U(d) + \frac{1}{2}U''(d)(r-d)^2$ near the minimum. The integrand then becomes a Gaussian, and since it decays very fast away from the minimum, we may extend the limits of integration to $\pm\infty$. Evaluating the Gaussian integral gives us the expression given for κ in Sec. I,

$$\kappa \approx e^{-\beta U(d)} \int_{-\infty}^{\infty} e^{-\frac{1}{2}\beta U''(d)(r-d)^2} dr = \sqrt{\frac{2\pi}{\beta U''(d)}} e^{-\beta U(d)}. \quad (C2)$$

We can show that our sticky parameter is related to the more commonly used second virial coefficient B_2 . The second virial coefficient can be expressed as [27]

$$B_2(\beta) = -\frac{1}{2} \int_0^{\infty} (e^{-\beta U(r)} - 1) 4\pi r^2 dr, \quad (C3)$$

and appears as the second coefficient in a power-series correction to the ideal gas law for a given interaction potential. Again, we can approximate B_2 by using Laplace asymptotics in the same way. We first note that for $r \gg d$, $U(r) \approx 0$. Therefore, the term $(e^{-\beta U(r)} - 1)$ goes to 0 exponentially fast, contributing nothing to the integral. Near the minimum, $(e^{-\beta U(r)} - 1) \approx e^{-\beta U(r)}$, and we can proceed as before. We make a second-order approximation of the potential and evaluate the second moment of the Gaussian integral. The final

result is

$$\begin{aligned} B_2(\beta) &\approx -\frac{1}{2}e^{-\beta U(d)} \int_{-\infty}^{\infty} 4\pi r^2 e^{-\frac{1}{2}\beta U''(d)(r-d)^2} dr \\ &= -\left[\frac{2\pi}{\beta U''(d)}\right]^{3/2} e^{-\beta U(d)} \\ &= -\frac{2\pi}{\beta U''(d)} \kappa, \end{aligned}$$

showing B_2 is a approximately linear function of the sticky parameter.

APPENDIX D: PREDICTING MERGE EVENTS FROM SHS GEOMETRY

We established a relationship between the value of the range parameter in which a cluster rearranges, ρ_c , and the minimum distance between nonbonded particles just before this rearrangement, $r_i(\rho_c+)$. Ideally, a one-to-one map between ρ_c and $r_i(\infty)$, coming from the sticky-hard-sphere geometry, would allow us to predict when any structure would merge. Figure 14 shows a scatter plot of $r_i(\infty)$ vs ρ_c , with smooth transitioning clusters in blue and rearranging clusters in red. The curves are fit using nonlinear least squares with the form $ax^b + c$.

Unfortunately, we find that the map is not one to one. There are many points for which $r_i(\infty) = 1.0887$ and $r_i(\infty) = \sqrt{2}$, but the corresponding ρ_c values are spread over a large range. Despite this, $r_i(\infty)$ does provide us with an approximation of when clusters will merge. This minimum distance can evolve differently as a function of cluster geometry, but the variations are typically small. Hence, if $r_i(\infty) = 1.0887$ or $r_i(\infty) = \sqrt{2}$, you can safely state that the cluster will merge *around* $\rho = 38$ and $\rho = 6$, respectively.

APPENDIX E: $N = 8$ MERGING TREES

Figure 15 shows the merging tree for $N = 8$. As for $N = 7$, the tree is insensitive to the value of the sticky parameter. Changing the form of the potential affects the range at which clusters merge. Figure 16 separates the tree into the four major merge events, with SHS clusters at the top and green arrows indicating a smooth transition. In Fig. 16(a), we see that the $N = 7$, $\rho = 38$ mechanism is responsible for this merge. In Fig. 16(b), the rearranging clusters have an octahedron as a substructure and the smooth cluster is new, i.e., it does not contain any $N = 6$ or $N = 7$ clusters as substructures. In Fig. 16(c), all the SHS clusters have a polytetrahedron as a substructure and they merge with the resulting cluster from Fig. 16(b). Finally, Fig. 16(d) contains one structure with an octahedron substructure and a new cluster, which rearrange to form the final cluster.

-
- [1] D. W. Brenner, The art and science of an analytic potential, *Phys. Status Solidi B* **217**, 23 (2000).
 - [2] D. Padmavathi, Potential energy curves and material properties, *Mater. Sci. Appl.* **2**, 97 (2011).
 - [3] D. Wales, *Energy Landscapes, Applications to Clusters, Biomolecules and Glasses* (Cambridge University Press, Cambridge, 2004).
 - [4] V. N. Manoharan, Colloidal matter: Packing, geometry, and entropy, *Science* **349**, 1253751 (2015).
 - [5] M. G. Noro and D. Frenkel, Extended corresponding-states behavior for particles with variable range attractions, *J. Chem. Phys.* **113**, 2941 (2000).
 - [6] G. Meng, N. Arkus, M. P. Brenner, and V. N. Manoharan, The free-energy landscape of clusters of attractive hard spheres, *Science* **327**, 560 (2010).
 - [7] R. W. Perry, M. C. Holmes-Cerfon, M. P. Brenner, and V. N. Manoharan, Two-Dimensional Clusters of Colloidal Spheres: Ground States, Excited States, and Structural Rearrangements, *Phys. Rev. Lett.* **114**, 228301 (2015).
 - [8] Y. Kallus and M. Holmes-Cerfon, Free energy of singular sticky-sphere clusters, *Phys. Rev. E* **95**, 022130 (2017).
 - [9] R. J. Baxter, Percus-Yevick equation for hard spheres with surface adhesion, *J. Chem. Phys.* **49**, 2770 (1968).
 - [10] M. Holmes-Cerfon, S. J. Gortler, and M. P. Brenner, A geometrical approach to computing free-energy landscapes from short-ranged potentials, *Proc. Natl. Acad. Sci.* **110**, E5 (2013).
 - [11] N. Arkus, V. Manoharan, and M. Brenner, Deriving finite sphere packings, *SIAM J. Discrete Math.* **25**, 1860 (2011).
 - [12] R. S. Hoy, J. Harwayne-Gidansky, and C. S. O'Hern, Structure of finite sphere packings via exact enumeration: Implications for colloidal crystal nucleation, *Phys. Rev. E* **85**, 051403 (2012).
 - [13] R. S. Hoy, Structure and dynamics of model colloidal clusters with short-range attractions, *Phys. Rev. E* **91**, 012303 (2015).
 - [14] M. Holmes-Cerfon, Enumerating rigid sphere packings, *SIAM Rev.* **58**, 229 (2016).
 - [15] C. P. Royall, S. R. Williams, T. Ohtsuka, and H. Tanaka, Direct observation of a local structural mechanism for dynamic arrest, *Nat. Mater.* **7**, 556 (2008).
 - [16] A. Malins, S. R. Williams, J. Eggers, H. Tanaka, and C. P. Royall, Geometric frustration in small colloidal clusters, *J. Phys.: Condens. Matter* **21**, 425103 (2009).
 - [17] D. J. Wales, A microscopic basis for the global appearance of energy landscapes, *Science* **293**, 2067 (2001).
 - [18] L. Trombach, R. S. Hoy, D. J. Wales, and P. Schwerdtfeger, From sticky-hard-sphere to Lennard-Jones-type clusters, *Phys. Rev. E* **97**, 043309 (2018).
 - [19] L. Trombach and P. Schwerdtfeger, Gregory-Newton problem for kissing sticky spheres, *Phys. Rev. E* **98**, 033311 (2018).
 - [20] J. W. R. Morgan, D. Mehta, and D. J. Wales, Properties of kinetic transition networks for atomic clusters and glassy solids, *Phys. Chem. Chem. Phys.* **19**, 25498 (2017).
 - [21] K. Bezdek, Contact numbers for congruent sphere packings in Euclidean 3-space, *Discrete Comput. Geometry* **48**, 298 (2012).
 - [22] K. Bezdek and M. Khan, Contact numbers for sphere packings, *New Trends in Intuitive Geometry*, edited by G. Ambrus, I. Bárány, K. Böröczky, G. Fejes Tóth, and J. Pach, Bolyai Society

- Mathematical Studies Vol 27 (Springer, Berlin, Heidelberg, 2018).
- [23] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, New York, 1992).
- [24] D. E. King, Dlib-ml: A machine learning toolkit, *J. Mach. Learn. Res.* **10**, 1755 (2009).
- [25] A. Trubiano, SHS continuation, <https://github.com/onehalfatsquared/SHScontinuation> (unpublished).
- [26] W. Kabsch, A solution for the best rotation to relate two sets of vectors, *Acta Crystallogr. Sect. A* **32**, 922 (1976).
- [27] M. E. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation* (Oxford University Press, New York, 2010).