



Exploring canyons in glassy energy landscapes using metadynamics

Amruthesh Thirumalaiswamy^a, Robert A. Riggleman^{a,1,2}, and John C. Crocker^{a,1,2}

Edited by Daan Frenkel, University of Cambridge, Cambridge, United Kingdom; received June 18, 2022; accepted September 8, 2022

The complex physics of glass-forming systems is controlled by the structure of the low-energy portions of their potential energy landscapes. Here we report that a modified metadynamics algorithm efficiently explores and samples low-energy regions of such high-dimensional landscapes. In the energy landscape for a model foam, our algorithm finds and descends meandering canyons in the landscape, which contain dense clusters of energy minima along their floors. Similar canyon structures in the energy landscapes of two model glass formers—hard sphere fluids and the Kob–Andersen glass—allow us to reach high densities and low energies, respectively. In the hard sphere system, fluid configurations are found to form continuous regions that cover the canyon floors up to densities well above the jamming transition. For the Kob–Andersen glass former, our technique samples low-energy states with modest computational effort, with the lowest energies found approaching the predicted Kauzmann limit.

glasses | metadynamics | energy landscape | hard spheres

Many outstanding problems in understanding glasses have been related to their potential energy landscapes (1), hypersurfaces describing a system's total potential energy spanning the high-dimensional configuration space formed by all the particles' spatial coordinates (2). These landscapes have a complex geometry, as indicated by an energy-dependent distribution of basin hypervolumes (3, 4), highly tortuous steepest descent paths (5, 6), and fractal clustering of local minima (6). Rapid quenching from a very high temperature (equivalent to following a steepest descent path from a random configuration) reports the landscape's local minima, or inherent structures (ISes) (2), each weighted by their associated catchment basin hypervolume (3, 7). The ensemble of such quenched configurations (8) having the largest basin hypervolumes are distinctly different from the lowest-energy ISes (9, 10) that control glass transitions (11–13). Navigating such high-dimensional spaces and mapping the arrangement of glassy states remains a major challenge. Further, the heterogeneous nature of glassy dynamics (14) makes the use of collective descriptors of system dynamics ineffective. Methods like eigenvector-following (15) and techniques for exhaustively enumerating ISes (16) allow the spatial arrangement of glassy states to be explored but tend to be computationally expensive (17, 18) in large systems. Swap Monte Carlo (9, 10, 19) and similar methods allow the canonical sampling of glassy states but jump around configuration space, obscuring the states' arrangement, and these methods are ineffective in bonded systems. Meanwhile, optimizers such as basin-hopping (20) efficiently find the lowest states but operate on a modified landscape without barriers, overlooking interesting characteristics of the landscape in the process.

Here we modify a high-dimensional metadynamics (21) algorithm by Yip and coworkers (22) and use it to discern the arrangement of ISes in glassy energy landscapes, calling our approach the metadynamics-inspired multifractal sampling explorer (MIMSE). When this algorithm is applied to a model of foams or soft glassy materials (SGMs) (6), it finds meandering canyon-like structures in the landscape and descends into them to find many ISes clustered on the canyons' floors. Analyzing small ensembles of biased trajectories reveals that these canyons have well-defined widths and become narrower as they meander to progressively lower energies. To apply this approach to a popular model of dense hard sphere (HS) fluids (9), we examine the corresponding extended energy landscape for compressible or soft spheres. The algorithm again finds and descends similar canyons, finding HS configurations (where the soft sphere potential energy is zero) that cover the canyon floor. These zero energy configurations appear to form a continuous, connected domain, which can be effectively found up to volume fractions as high as 0.68. However, locating these HS states from random points in configuration space becomes exponentially costly at higher densities. Last, when applied to a simple model of an atomic glass former, the Kob–Andersen (KA) model (23), our algorithm again finds and rapidly descends canyons, reaching very low energies with modest computational effort.

Significance

The inability of glasses to find their lowest-energy configurations is often attributed to their potential energy landscapes: rugged, barrier-filled surfaces in high-dimensional space that seem incredibly difficult to navigate. We report that for three glassy systems, to the contrary, the landscape is pervaded by canyon-like metabasins. While the earlier-noted rugged surfaces line the canyon floors, the canyons' walls can be followed directly to lower-energy states. This remarkable numerical discovery enables our algorithmic approach to find very low energy glassy states and will stimulate accelerated simulation approaches in glasses and perhaps related fields such as deep learning.

Author affiliations: ^aDepartment of Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA 19014

Author contributions: A.T., R.A.R., and J.C.C. designed research; A.T. performed research; and A.T., R.A.R., and J.C.C. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹R.A.R. and J.C.C. contributed equally to this work.

²To whom correspondence may be addressed. Email: rrig@seas.upenn.edu or jcrocker@seas.upenn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2210535119/-DCSupplemental>.

Published October 18, 2022.

Exploring many canyons yields a distribution of canyon floor energies the lowest of which approach the predicted Kauzmann (24) energy limit—the energy at the Kauzmann temperature, the notional point for an ideal glass transition. Surprisingly, the low-dimensional canyon floors that contain the glassy configurations were found to have the same effective fractal dimension in all three systems' landscapes.

The MIMSE Algorithm

MIMSE is an athermal, metadynamics-based approach (21) directly applied to the $3N$ -dimensional potential energy landscape of N particles and uses sequentially added bias potentials to overcome energy barriers (22). We use the fast inertial relaxation engine (FIRE) (25) to relax the system configuration on the biased energy landscape; tests using steepest descent gave similar results but at higher computational cost. The bias potential employed is a $3N$ -dimensional, bell-shaped, smooth quartic function with the form

$$U(\mathbf{r}) = \begin{cases} U_0 \left(1 - \left(\frac{\|\mathbf{r} - \mathbf{r}_m\|}{U_\sigma} \right)^2 \right)^2, & \text{if } \|\mathbf{r} - \mathbf{r}_m\| < U_\sigma \\ 0, & \text{otherwise.} \end{cases} \quad [1]$$

Here $\mathbf{r}$ represents the $3N$ -dimensional system configuration, and $\mathbf{r}_m$ is the IS minimum location around which the bias is centered. The bias parameters U_σ and U_0 represent the $3N$ -dimensional radial extent and energetic height of the bias, respectively. To ensure that the center of mass of the system remains at rest, we modify the corresponding bias forces to ensure the total force on the system is zero. To ensure stable configuration dynamics, we use models with energy landscapes that are continuous and differentiable.

Algorithmic flow starts at a local minimum of the landscape, sampled typically from the quenched ensemble, described earlier. Starting at such a minimum, a bias is added centered on its location, forming a local maximum on the biased landscape. To move away from this unstable maximum, the system is given a small, random $3N$ -dimensional displacement (corrected so as to preserve the position of the system's center of mass). This configuration is then relaxed to the nearest minimum on the biased landscape using FIRE, and the process repeated by adding a new bias at each subsequent minimum.

In conventional metadynamics, the bias radius is smaller than the separation between neighboring ISes, leading to basin filling and exhaustive enumeration of minima (21, 22). Here we use larger bias radii, such that when a bias is added, many ISes within U_σ of $\mathbf{r}_m$ are covered by that bias, effectively forcing the system over their associated nearby energy barriers. Values of U_σ and U_0 that lead to efficient landscape descent in energy are found by manual tuning using a divide and conquer-like approach (*Materials and Methods*). For computational efficiency, a $3N$ -dimensional neighbor list is maintained to track the different biases affecting the system at its current $3N$ -dimensional position. In practice, we find that reasonably old biases can often be retired without difficulty, keeping computational and memory costs manageable. Additional details of the algorithm are provided in *Materials and Methods*.

As the algorithm proceeds, we distinguish among the minima encountered and store and analyze the subset which lie outside the bias potentials, which also correspond to physical ISes of the unbiased energy landscape. This approach consistently yields ISes in glassy landscapes.

SGM Energy Landscape

We first demonstrate our approach on a model for foam or SGMs (26–28), whose energy landscape was characterized in an earlier study (6). The system consists of a highly polydisperse collection of $N \sim 350$ compressible soft spheres interacting via purely repulsive harmonic interactions when they overlap. We choose a volume fraction, $\phi = 0.75$, about 0.03 higher than the threshold for solidity or jamming, ϕ_j^{SGM} , in this model (6). The distribution of radii (R) resembles that of bubbles in a ripening foam, described by a Weibull distribution with polydispersity $\Delta = (\langle R^2 \rangle - \langle R \rangle^2)^{1/2} / \langle R \rangle \simeq 0.64$; see *SI Appendix* for model details. To create initial configurations of this model, we select random points in configuration space and relax them to their nearest energy minimum using FIRE. This ensemble of quenched ISes has a roughly Gaussian distribution of energies (Fig. 1 *A*, *Left*).

Our key finding is that for certain bias parameters, MIMSE rapidly descends down the potential energy landscapes of glassy systems. Typical results for the SGM model landscape are shown in Fig. 1, reaching energies that are much lower than the energies of quenched configurations after only a few hundred biases. These proceed to lower energy logarithmically with additional effort, reaching energies which, compared to the quenched energy distribution, are $\mathcal{Z}_U \approx -5$ SDs below the mean of the initial quenched




Fig. 1. Exploring a soft-glassy material landscape with MIMSE finds low-energy ISes efficiently. (A) (*Left*) The probability distribution of energies for an ensemble of 1,000 different quenched initial configurations, formed by FIRE relaxation of random configurations. Applying MIMSE, the energy of the configurations decreases rapidly, depending on bias radius, U_σ . (*Right*) Rescaled energy relative to the mean and SD of the quenched energies U_q , also known as the z score, $\mathcal{Z}_U$. (B) The descending energy trajectory appears roughly logarithmic in time. (*Inset*) The number of biases that are required to descend a given energy range (dashed lines in B) shows a strong dependence on U_σ , which can be approximated by a power law over a narrow range.

ISes (Fig. 1, *Right*). Notably, such significant descent in energy is only found to occur in a window of bias radii, $2 \lesssim \mathcal{U}_\sigma \lesssim 5$, and the number of biases required to descend by a given energy shows a strong dependence on the bias radius (Fig. 1 *B*, *Inset*). Larger biases descend the landscape far faster than smaller ones, up to a maximum radius. The energetic height of the bias $\mathcal{U}_0$ appears less critical; descent down the energy landscape is observed in a wide window, $10^{-3} \lesssim \mathcal{U}_0/\mathcal{U}_\sigma \lesssim 10^3$.

Our first task is to understand what features of the SGM landscape enable our algorithm to descend to such low energies and to explain the narrow window of bias radii where it occurs. To this end, we formed ensembles of 1,100 FIRE trajectories, each descending the same bias around the same starting IS minimum, but initialized with different isotropically random $3N$ -dimensional displacements, and examined the results as a function of bias radius. We find that after initially moving radially away from the center in random directions due to the strong bias force, these paths then drift in the angular ($3N - 1$) dimensions according to the gradient of the underlying landscape. Finally, these paths terminate in 1,100 different energy minima near the edge of the bias.

To descend the energy landscape, the FIRE minimizer must be following gradients in the underlying landscape; perhaps the window of bias radii is due to suitable gradients only being present on corresponding length scales. To test this idea, we first compute the average drift of the ensemble of trajectories, such as would be caused by an underlying gradient. Specifically, for each trajectory we first compute a $3N$ -dimensional unit vector $\hat{u}$ pointing from the initial IS to the new energy minimum. Next, we compute the Euclidean length of the average of these unit vectors, $\|\langle \hat{u} \rangle\|$, akin to the center of mass of the new minima, as a function of bias radius, $\mathcal{U}_\sigma$. The results for three starting ISes at different depths in the landscape are shown in Fig. 2*A*. While the average drift shows a broad peak for the $\mathcal{U}_\sigma$ where energy descent is observed, a similar amount of drift is seen for much smaller and larger radii, failing to explain the lack of energy descent in those cases.

To gather clues about descent and barrier crossing in the underlying landscape, we applied a hierarchical single-linkage clustering algorithm (29, 30) to the unit vector ensembles analyzed above. The clustering algorithm characterizes clusters as sets of points such that each point in a cluster has at least one other constituent point within a given distance threshold (*Materials and Methods*). The idea behind this analysis of clustering is that if a subset of the trajectories are crossing a barrier, they will form a separate cluster from the others. For small and larger $\mathcal{U}_\sigma$ values, we find multiple clusters indicating the exploration of IS clusters separated by barriers. For intermediate values of $\mathcal{U}_\sigma$, we report that these ISes form a single extended cluster. Remarkably, we now find a one-to-one correspondence between the bias radii that give efficient energy descent and the $\hat{u}$ vectors that form a single cluster, for ISes at the three different energies considered (Fig. 2*B*). Moreover, it may be noted that the maximum $\mathcal{U}_\sigma$ that forms a single cluster gets progressively smaller at lower-energy domains in the landscape, decreasing from 7 to 6 to 5 units. This observation and the correspondence between bias radius and energy descent is confirmed by another result: bias radii in the range 5 to 7 rapidly descend the landscape at first but then get stuck, plateauing at intermediate energies (*SI Appendix*, Fig. S1). Analysis of the same vectors using complete linkage clustering fails to resolve these clusters well, suggesting that they are extended and intertwined, rather than compact and well separated.

These findings indicate how the algorithm works: for the appropriate bias size, it is simply following clusters of ISes like a trail of breadcrumbs to low-energy portions of the landscape. Finding




Fig. 2. Trajectory ensembles reveal landscape gradients and IS clustering. Ensembles of 1,100 FIRE trajectories descending the same hyperspherical bias potential after different, isotropically random initial displacements move radially before being affected by the underlying landscape and terminating near the bias edge (schematics shown in *Insets*). The final relative displacements are converted to unit vectors, $\hat{u}$. (A) The norm of the average of the 1,100 unit vectors $\|\langle \hat{u} \rangle\|$ reveals the average drift due to an energy gradient, for three different depths in the landscape, $\mathcal{Z}_U \approx -2.8$ (squares), -3.8 (circles), and -4.7 (triangles), showing a broad peak around $\mathcal{U}_\sigma \approx 3$. (B) Performing single-linkage hierarchical cluster analysis (29) on the same $\hat{u}$ ensembles yields a single cluster only for $\mathcal{U}_\sigma \gtrsim 2$, with an upper limit that depends weakly on $\mathcal{Z}_U$. Landscape energy descent only occurs for bias radii and energies where a single cluster is found, as explained in SGM Energy Landscape.

multiple clusters with too large biases indicates the crossing of large-scale barriers into other domains, losing the trail. Finding many clusters with too small biases indicates the crossing of small-scale barriers but leaves the algorithm trapped, exploring within a dense IS cluster at nearly the same location. Last, a fractal scaling analysis, discussed in *KA Glass Former* and *SI Appendix*, of the sampled configurations suggests that the set of low-energy ISes occupy a subspace with a low effective dimensionality of ≈ 2.5 .

The above analyses can be summarized schematically in Fig. 3*A*. The landscape consists of meandering canyons or tubes [as suggested by earlier studies (5, 6)], with the ISes forming extended dense clusters along the canyon floors. These dense clusters act as traps for MIMSE when the bias radius is too small. For optimal bias radii, as shown in Fig. 3*A*, the biased FIRE path bounces off the canyon walls (finding no ISes there), evolves under the influence of the landscape gradient, and ends at a new minimum farther down the canyon. Sometimes this path extends out of the biased region completely, ending in a physical IS. Too large biases push the configuration into adjacent canyons, apparently preventing energy descent. Since the largest $\mathcal{U}_\sigma$ that descends the landscape decreases at lower energies, it appears that the canyon tapers as it descends to lower energies. Given that the canyon is high-dimensional, the reduction in its hypervolume due to even a slight decrease in width is enormous.




Fig. 3. Schematic aerial view of the energy landscape showing canyons (white region) with tortuous floors (dashed line) descending to lower energy (left to right). (A) For the SGM system, ISes (black dots) are clustered along the canyon floor, and MIMSE skips from IS to IS (solid red path), when biases (large red circles) are sequentially applied (left to right). Red open circles represent biased minima; unbiased minima (black arrows) are stored as output. (B) For the HS glass, the canyon floor is covered by a slender connected domain having zero energy (blue), as if flooded by a stream of water, and MIMSE skips between configurations where the bias and the zero energy domain intersect (small red circles). Some ISes in A may have zero energy (blue) but may be either blocked by other ISes or too rare to sample.

HS Fluid

To determine if the energy landscapes of glass formers also contain canyon structures that resemble those in the SGM system, we apply MIMSE to two additional well-studied models. The first is a model of HS fluids (9, 10, 31), having a size distribution given by $P(R) \sim R^{-3}$, $R \in [R_{\min}, R_{\max}]$, $R_{\min}/R_{\max} = 0.4492$, and $N = 1,000$ particles. Such a polydisperse system has been shown to consistently avoid crystallization, instead undergoing kinetic arrest into an HS glass as the volume fraction is increased (9, 10).

Importantly, the HS potential segments configuration space into two domains, a physical one with zero energy and a forbidden domain with one or more particle overlaps. To form a continuous energy landscape, we consider a soft-sphere extension to the HS model, with harmonic repulsion; see *SI Appendix* for details. This model thus resembles the SGM model but has a different particle radius distribution and volume fraction. We then consider any configurations found with energy per particle less than a small tolerance ($U/N < U_{\text{tol}} = 10^{-16}$) to be physical HS configurations (9). For such states with zero energy, we use a different order parameter to follow the progress of our algorithm across the landscape, corresponding to the mean coordination or

contact number $\langle z \rangle$. Since pair contacts containing unconstrained or rattler particles can be trivially relaxed, $\langle z \rangle$ is a measure of the force bearing contacts. Its value is confirmed to be insensitive to the choice of U_{tol} (*SI Appendix*, Fig. S2).

We perform MIMSE relaxation experiments varying the volume fraction initially over a small range, $0.66 < \phi < 0.67$. To provide initial ISes, we FIRE relax random configurations. Since this range is above the jamming or random close-packed volume fraction for this model, $\phi_j^{\text{HS}} \simeq 0.65$, this initial quench step consistently yields ISes having finite energy, residing on the soft sphere portion of the landscape. Applied to these quenched states, we find MIMSE again rapidly descends the energy landscape (Fig. 4A) when the bias radius is in a small range ($3 \lesssim U_{\sigma} \lesssim 5$), indicating that there are canyons similar to those in the SGM landscape.

For the lower range of volume fractions we study, $\phi \lesssim 0.667$, MIMSE is able to consistently reach states with zero energy, corresponding to HS configurations (Fig. 4A). The first HS configurations found are isostatic (or nearly so), with $\langle z \rangle \approx 6$. Further application of biases briefly yields a mixture of HS and finite energy states, before reaching a part of the canyon where it samples a continuous series of HS states. Continuing MIMSE results in configurations with $\langle z \rangle$ dropping by roughly three orders of




Fig. 4. MIMSE descends both the soft sphere and HS portions of the HS glass landscape. (A) MIMSE efficiently descends down canyons, reducing the energy (Left) of the soft-sphere system, ending in a nearly isostatic, jammed HS state. Further application of biases (Right) shows that the average coordination number of these zero-energy configurations, $\langle z \rangle$, drops significantly, indicating structural relaxation, before reaching a steady state value. (B) Running MIMSE on a range of different starting volume fractions shows us that MIMSE requires a larger number of biases to reach the HS portion of the landscape, becoming very computationally expensive for $\phi \gtrsim 0.667$. The colors represent independent runs starting from four different random initializations in configuration space. The dashed line is an exponential eye guide. The arrow-headed lines denote lower limits derived from incomplete, long simulations.

magnitude after $\sim 10^3$ biases (Fig. 4A). Notably, when sampling the zero-energy domain, MIMSE finds a new, unbiased zero-energy state after every bias addition and no finite energy ISes. The $g(r)$ of these configurations on the $\langle z \rangle$ plateau resembles the jammed case (SI Appendix, Fig. S2B), suggesting that the algorithm is sampling a region of the landscape corresponding to high pressures, adjacent to the jamming line (9, 32).

The behavior of the MIMSE algorithm while exploring such HS states provides insights about their distribution in configuration space, summarized in Fig. 3B. Starting from an initial quenched configuration, MIMSE descends a finite energy canyon until it first reaches a portion of the canyon floor containing small puddles of zero energy, before reaching a region where nearly all ISes are covered by a slender, connected domain of HS configurations. It is as if the canyon floor were flooded by a stream or level set having zero energy. In this case, for every added bias, the canyon walls funnel all FIRE paths to points on the intersection of the edge of the bias and the edge of a zero energy domain (Fig. 3B). FIRE then halts when the total energy per particle first drops below U_{tol} . The fact that the hyperspherical edge of every bias intersects a zero energy domain leads us to hypothesize that the domain is continuous (at least on the bias length scale used to search the configuration space).

As ϕ increases, MIMSE must work harder to reach HS states. Starting from a random quenched configuration, the number of biases required to reach the HS states (proportional to the contour distance traveled along the canyon floor) increases exponentially with ϕ (Fig. 4B). Notably, some initial configurations (and canyons) appear deeper than others, with a broad dispersion in the amount of effort to reach the HS states. Indeed, for $\phi \gtrsim 0.667$, MIMSE does not consistently reach HS states with a reasonable computational effort. It is as if the entire landscape has been raised and the stream drained away to a deeper, unreachably distant stretch of the canyon. In the opposite limit, the jamming volume fraction ϕ_J^{HS} corresponds to lowering the landscape so far as to flood the entire canyon, such that the HS domain covers even the highest, quenched ISes on the landscape. Overall, the finite energy portions of the canyons resemble that seen in the SGM case (Fig. 3A). Moreover, the number of biases between subsequent ISes on the soft sphere portion of the landscape increases with ϕ , and the number of FIRE steps required per bias increases by roughly 10-fold for $\phi \gtrsim 0.666$. This suggests a ϕ -dependent change in the ruggedness of the landscape.

Our results thus far finding canyons leading deeper into the landscape poses the question of how far (or deep in ϕ) do the HS domain streams in these canyons go? Theoretically, the entropy of the equilibrium HS fluid is expected to nearly vanish (become subextensive) at an HS glass transition corresponding to ϕ_K^{eq} . An extrapolation of simulation data for this system in ref. 9 suggests that this occurs at $\phi \approx 0.672$, only slightly higher than that reached by MIMSE above. Thus, it is natural to conjecture that the canyons we are exploring might cease to contain HS states (or run dry) for some higher value of ϕ .

To reach deeper portions of the canyons and find higher volume fraction HS states, we combined MIMSE with an adapted affine compression/relaxation scheme (33, 34); see SI Appendix for details. This method does consistently push the configuration farther down the canyon and reach higher ϕ states, again limited by computational effort. This approach typically yields HS states with volume fractions of $\phi \approx 0.670 \pm 0.001$. Remarkably, however, 2 of 30 runs reached dramatically higher volume fractions, $\phi = 0.681$ and 0.691 , the latter roughly 0.04 above ϕ_J^{HS} . Such dense configurations have previously been generated using swap Monte Carlo combined with compression methods (10). Notably,

unlike swap, the particle displacements used by our approach resemble physically allowed moves. Our findings thus suggest that a significant fraction of random quenched configurations (at least $2/30$) are connected to these ultradense states by physical trajectories.

To explore the landscape around the MIMSE compression generated states, we used them as initial states for MIMSE runs, dilating them by various amounts. Without dilation, MIMSE quickly jumps out of the HS domain, sometimes finding a few adjacent HS states, and then only sampling soft sphere ISes. This suggests that rather than a stream, the HS domains are puddles smaller than or comparable in size to the bias radius. When the configuration is dilated by $\Delta\phi \approx 1 - 2 \times 10^{-3}$, however, MIMSE again finds long streams of HS states (and familiar canyon walls) that are easy to navigate (finding a new HS IS after every bias addition) adjacent to all the configurations including the one at $\phi = 0.681$ but not the densest $\phi = 0.691$ configuration. The landscape looks qualitatively different surrounding the latter state. This intriguingly suggests a change in the landscape for $0.68 < \phi < 0.69$, but clearly, further study is required.

KA Glass Former

Last, to determine if the canyons we find are also a feature of the energy landscapes of atomic glasses, we consider the KA model. This consists of a binary mixture having a total of $N = 256$ particles with $x_A = 0.8$ at a total number density of $\rho_{total} = 1.2$. To ensure force continuity, we employ a quadratically smoothed, truncated form of the Lennard–Jones potential (35):

$$V(\mathbf{r}_{ij}) = \begin{cases} k\epsilon_{ij} \left(\left(\frac{\sigma_{ij}}{\|\mathbf{r}_{ij}\|} \right)^{12} - \left(\frac{\sigma_{ij}}{\|\mathbf{r}_{ij}\|} \right)^6 \right) \\ + \nu(\|\mathbf{r}_{ij}\|), & \text{if } \|\mathbf{r}_{ij}\| < \mathbf{r}_c \\ 0, & \text{otherwise,} \end{cases} \quad [2]$$

where $\epsilon_{AA} = 1.0$, $\epsilon_{BB} = 0.5$, and $\epsilon_{AB} = 1.5$; $\sigma_{AA} = 1.0$, $\sigma_{BB} = 0.88$, and $\sigma_{AB} = 0.8$; and $\nu(\|\mathbf{r}_{ij}\|)$ is a smoothing function (see SI Appendix for details). As before, we first choose random points in configuration space and FIRE relax them to their first energy minimum to form a quenched ensemble of ISes (Fig. 5A, Left).

When an optimal bias radius is used, we find that MIMSE efficiently descends to low-energy portions of the KA landscape (Fig. 5A). As in the SGM case (Fig. 1) the algorithm also proceeds to lower energy roughly logarithmically over the ISes sampled, despite the KA landscape's presumed differences from the earlier case. As before, the maximum bias radius is much smaller than the contour length of the configuration space path, consistent with a meandering canyon-like structure. Descent, however, occurs in a much narrower window of bias radius, $1.0 \lesssim U_\sigma \lesssim 1.5$. The KA system also requires far more biases for each new IS sampled. For the optimal bias parameters used, the distribution of the number of biases per new IS is heavy-tailed with a median of ≈ 9 and $\langle n \rangle$ of ≈ 27 , versus $\langle n \rangle \approx 1.4$ in the SGM case. In addition to requiring more biases on average for each new IS, the algorithm must occasionally apply a large number of biases ($n > 10^3$) to generate a new IS. Examination of the trajectories reveals the configuration path doubling back upon itself multiple times before crossing energy barriers. We interpret these differences from the SGM system as the landscape of the KA system being more rugged, requiring the filling of subbasins (having a broad range of hypervolumes) with biases to cross the energy barriers necessary to find each new IS. We typically terminate the MIMSE run after a fixed amount of computational effort.




Fig. 5. MIMSE descends the landscape of the KA glass former. (A) (Left) The probability distribution of energies for an ensemble of 1,000 quenched initial configurations, formed by FIRE relaxation of random configurations. (Right) We find that MIMSE successfully reaches low-energy configurations for bias radii, $U_\sigma \approx 1$. Both runs used U_0 of 20. (B) Plotting a distribution of minimum energies from $\sim 1,000$ runs shows that it closely approaches the Kauzmann limit (dashed line) (24). (C) The CDFs (36) of the pairwise separations of the ISes for all three systems show that they form low-dimensional clusters with very similar effective dimension: $D_f \approx 2.4$. Simulation parameters were $U_0 = 3, U_\sigma = 3$ (SGM); $\phi = 0.669, U_0 = 0.85, U_\sigma = 3.98$ (HS system with $U > 0$); and $U_0 = 20, U_\sigma = 1.25$ (KA). The nonuniform sampling of the ISes was corrected as explained in SI Appendix.

The computational effort required to sample low-energy states compares favorably to conventional methods. Earlier studies using simulated annealing of the KA model (11) provide a relationship between IS energy and fictive equilibrium temperature. Characteristic temperatures such as the mode-coupling and Kauzmann temperatures can thus be mapped to their corresponding IS energies (24). Benchmarking our algorithm against a Langevin dynamics implementation of simulated annealing finds an ≈ 100 -fold improvement in computational speed to reach configurations near the mode-coupling energy ($U/N \approx -7.00$).

Repeated MIMSE relaxations from a number of different starting configurations yield a roughly Gaussian distribution of minimum energies (Fig. 5B). Since these runs are performed for a fixed computational effort, this distribution is kinetically defined by our algorithm and not a meaningful feature of the landscape itself. A structural analysis of these low-energy configurations using common neighbor analysis (37, 38) confirms that they are amorphous. An earlier study using basin-hopping minimization (20) yielded a sample of similar energy values for amorphous structures in the same system (39), suggesting that MIMSE is acting like existing optimizers. The empirically predicted value of the Kauzmann energy, U_K (24), is ~ 3 SDs below the mean of our distribution, $U_K/N \approx -7.08$. Sampling several hundred metabasins yields ISes very close to but not below the Kauzmann energy, consistent with this empirical prediction (Fig. 5B).

For comparison with conventional methods, a lengthy Langevin dynamics simulation equilibrated at around the mode-coupling temperature (24) yielded a much narrower distribution of IS energies when quenched, $U/N \approx -7.00 \pm 0.008$ and no states anywhere near the predicted Kauzmann limit.

The fact that the algorithm can fill subbasins in the canyon floor implies that the low-energy configurations occupy a low-dimensional subspace embedded in the higher-dimensional configuration space. To estimate the effective dimensionality of the canyon floor, we compute the fractal dimension of the ISes we find, correcting for their nonuniform sampling using methods described in *Materials and Methods* and SI Appendix. Specifically, we compute the correlation dimension (36) of a subensemble of ISes, i.e., the scaling exponent of the cumulative distribution function (CDF) of the $3N$ -dimensional Euclidean distances between all pairs of ISes. If the ISes form a fractal, the correlation dimension would report the fractal dimension, D_f , via $CDF \sim \|\Delta \mathbf{r}\|^{D_f}$. The roughly power law scaling seen (Fig. 5C) is consistent with a fractal dimension, $D_f \approx 2.4$ (36). The slight curvature of the plot suggests multifractal rather than true self-similar scaling. Unexpectedly, the observed multifractal geometry of the KA canyon floor appears to be very similar to that observed for canyon floors in the SGM system and the soft sphere ($U > 0$) portions of the HS system (Fig. 5C). Further analysis of the SGM system confirms a crossover to a slightly higher effective dimension at shorter length scales (SI Appendix).

The canyons we report have ISes distributed along their floors distributed within low-dimensional subspaces. While the system is $3N$ -dimensional, only a subset of particles move significantly between consecutive ISes, characteristic of a system with dynamical heterogeneity. The low effective dimensionality of the canyon floor, in the range 2 to 3, may indicate that only a few soft (or unstable modes) dominate the transition between adjacent ISes. If these configuration space directions become uncorrelated, we would expect the effective dimension of the canyon to be that of a random walk $D_f = 2$. The slightly higher dimensionality observed on still longer length scales could indicate caging of canyon path by very large energy basins.

Discussion

We have found that our metadynamics algorithm provides interesting insights into the large-scale structure of three glassy energy landscapes, in particular the canyon-like subspaces that both contain the glassy configurations and provide direct routes leading to the landscapes' low-energy states. While the floors of these canyons contain dense clusters of ISes and resemble the well-known rugged and barrier-filled landscapes of glasses, our work reveals that these glassy domains are surrounded by large, high-dimensional canyon walls effectively devoid of minima and barriers, and the canyons contain energy gradients that lead directly to lower energies in the landscape. Such simply connected routes between glassy states and the lowest-energy states pose a conundrum—how is it that simulated annealing struggles to find very low energy states? The answer is presumably a free energy barrier, as suggested by our finding that the canyons in the SGM landscape become narrower at lower energies in the landscape. In the HS case, we are tempted to associate the streams of HS states we found with the results of ref. 5 for small packings of HSes; they found most HS states comprised long, narrow threads that terminate in high-density cores (which themselves contain very few states). Our finding of narrow streams in the landscape up to at least $\phi = 0.681$, however, seems incompatible with the prediction that the HS entropy becomes subextensive at

a lower value. In the KA glass former, the algorithm is able to consistently reach low-energy portions of the canyon, including energies approaching the predicted Kauzmann limit. This finding confirms that by virtue of its ability to follow the canyon walls to lower energy, MIMSE can act like a global optimizer similar to basin hopping (20). Finally, the presence of similar canyon-like, low-dimensional subspaces with very similar multifractal geometry in these seemingly different systems might explain the range of qualitatively similar physical phenomenon and kinetics in different glassy systems (13, 32, 40).

Future work will apply MIMSE to other glassy landscapes, including bonded model systems. While informative, MIMSE does not obey detailed balance, nor does it return samples corresponding to any canonical ensemble. We anticipate that future work will fruitfully hybridize MIMSE with other methods, such as swap (10, 19) or ghost particle (41) Monte Carlo, and parallel tempering (42) to enable barrier crossing while obtaining canonical sampling. Recently, the loss landscape of deep neural networks has been shown to have similar features (43) to that of a soft glassy matter system; perhaps MIMSE will enable useful exploration of such landscapes as well.

Materials and Methods

A detailed description of the algorithmic procedure is provided below. Code has been deposited by the authors in GitHub, <https://github.com/rar-ensemble/MIMSE> (44).

- 1) A random energetic minima sampled via FIRE from a random configuration (quenched ensemble) is selected as the starting point.
- 2) A bias potential is added on to the system centered around the selected minimum in $3N$ -dimensional configuration space. (Note that to ensure the bias potential is center of mass preserving during evaluation of the bias, we subtract the mass weighted sum of forces on all particles.)
- 3) The system is given a small, positional displacement in a random direction in the configuration space, scaled such that at least one particle exceeds the system force tolerance, F_{tol} . This displacement is adjusted to preserve the center of mass location. The system is then quenched on the biased energy landscape ($U_{tot} = U + U_{i,bias}$, where i signifies the bias number).
- 4) Biases are added in succession until a new unbiased minimum is reached. This is verified by ensuring that $U_{total,bias}/N < U_{tol}$.
- 5) The above process is repeated until a desired amount of the energy landscape is sampled.

To undertake the above process with utmost efficiency and accuracy for our systems of interest with highly rugged energy landscapes, we use a finite-ranged bias potential, namely, a symmetric quartic function, as described in Eq. 1. Parameters $\mathcal{U}_\sigma$ and $\mathcal{U}_0$ represent the $3N$ -dimensional Euclidean extent and energetic height of the bias, respectively. Throughout our simulation, the center of mass of the system remains at rest due to the bias force exerting no force on the center of mass. This is enforced by subtracting the mass weighted total force from individual bias force calculations. Despite this correction, our bias remains isotropic in the $3N - 1$ angular dimensions.

We study and select pairs of parameters that efficiently sample lower-energy states effectively (and do so with a small number of biases per every new unbiased minima). We use a divide and conquer-like search procedure to determine an optimal pair of parameters for our algorithm. Our initial guesses for both parameters are based on the relevant length scales and the potential energy scales of the landscape. For example, in the SGM system, we search for

an optimal $\mathcal{U}_0$ between $\sim \epsilon/2$ and $\sim N^2 \epsilon/2$. Meanwhile, for $\mathcal{U}_\sigma$, we search between $\sim \langle R \rangle$ and $\sim \sqrt{N} \langle R \rangle$. In the range of $\mathcal{U}_\sigma$, values were scanned over logarithmically spaced intervals. At each $\mathcal{U}_\sigma$ value, a range of $\mathcal{U}_0$ values was considered. Larger $\mathcal{U}_0$ values, much greater than the energy scale of the system, make the determination of minima numerically inaccurate. Values over the range that gave consistent exploration low-energy ISes was considered. A blind systematic scan over a larger range of $\mathcal{U}_\sigma$ and $\mathcal{U}_0$ values gives a similar result. The selection of parameters is further verified by studying the gradient and clusters formed by trajectories emerging from a starting IS using different $\mathcal{U}_\sigma$ values (Fig. 2).

To keep track of the biases, we use a modified extended $3N$ -dimensional neighbor list that keeps track of all biases in high-dimensional space around the system configuration within a cutoff distance $\mathcal{U}_{c1}$. Further, we have a long-range secondary cutoff length $\mathcal{U}_{c2}$ which defines our neighbor search boundary; biases are retired when the system is farther than $\mathcal{U}_{c2}$ from a particular bias. Thus, in effect we maintain neighbors (biases) within $\mathcal{U}_\sigma + \mathcal{U}_{c1}$ and update the neighbor list by searching for biases within $\mathcal{U}_{c2}$ from the system position. The update is done when the system moves a Euclidean displacement more than $\mathcal{U}_{c1}$, whose value is chosen for computational efficiency. Parameter $\mathcal{U}_{c2}$ is chosen so that the ultimate trajectory of the system is not affected. It may be noted that both of these parameters are system and bias size dependent and need to be tuned for each system to balance the efficiency of the method with not retiring biases prematurely.

Our metadynamics-based algorithm differs from that developed by Yip in three key ways: 1) We use a smooth, truncated bias potential (Eq. 1) instead of the traditional Gaussian bias, that enables the use of a high-dimensional neighbor list. This allows us to significantly reduce the computational overhead associated with computing the total bias potential and also facilitates the retirement of biases beyond a certain distance. 2) Further, we make sure that the center of mass of the system remains at rest throughout the simulation. We enforce this on the bias potential and the random displacement after bias addition, which appears to lead to significantly more efficient descent down the landscape. 3) Last, we use an optimized minimizer, FIRE (25), that helps boost the computational efficiency while exploring such glassy high-dimensional landscapes.

These modifications and additions allow us to overcome the challenges posed by our high-dimensional system. It must be noted that each system responds differently to the biases in a way that depends on the characteristic energy and length scales of that particular system.

MIMSE efficiently samples low-energy ISes in three different systems, ISes sampled in each occupying a low-dimensional subspace. We characterize the dimensionality of this subspace by computing the correlation dimension (36); it corresponds to the logarithmic slope of the CDF of the $3N$ -dimensional Euclidean distances between all pairs of points in the ensemble as function of those distances. However, since the ensemble of ISes or $3N$ -dimensional points sampled by MIMSE are not uniform, we introduce a decluttering technique as follows. We first determine the minimum $3N$ -dimensional distance threshold required to form one connected single-linkage-hierarchical cluster. We use this length scale l_{max} to declutter our ensemble by discarding points that are spaced closer than l_{max} (see *SI Appendix* for more details). We used the subensemble of remaining points (which are now uniformly sampled on large length scales) to determine the effective dimensionality.

Data, Materials, and Software Availability. Code and sample input files have been deposited in GitHub (<https://github.com/rar-ensemble/MIMSE>) (44).

ACKNOWLEDGMENTS. We are grateful for useful conversations with Andrea Liu and Talid Sinno. This work was supported by NSF-Division of Material Research 1609525 and 1720530 and computational resources provided by XSEDE (Extreme Science and Engineering Discovery Environment) through TG-DMR150034.

1. D. J. Wales, J. P. K. Doye, M. A. Miller, P. N. Mortenson, T. R. Walsh, *Energy Landscapes: From Clusters to Biomolecules* (John Wiley & Sons, Ltd., 2007), vol. 115, pp. 1–111.
2. F. H. Stillinger, T. A. Weber, Hidden structure in liquids. *Phys. Rev. A* **25**, 978–989 (1982).
3. C. P. Massen, J. P. Doye, Power-law distributions for the areas of the basins of attraction on a potential energy landscape. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **75**, 037101 (2007).

4. S. Martiniani, K. J. Schrenk, J. D. Stevenson, D. J. Wales, D. Frenkel, Turning intractable counting into sampling: Computing the configurational entropy of three-dimensional jammed packings. *Phys. Rev. E* **93**, 012906 (2016).
5. S. S. Ashwin, J. Blawdziewicz, C. S. O'Hern, M. D. Shattuck, Calculations of the structure of basin volumes for mechanically stable packings. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **85**, 061307 (2012).

6. H. J. Hwang, R. A. Riggleman, J. C. Crocker, Understanding soft glassy materials using an energy landscape approach. *Nat. Mater.* **15**, 1031–1036 (2016).
7. N. Xu, D. Frenkel, A. J. Liu, Direct determination of the size of basins of attraction of jammed solids. *Phys. Rev. Lett.* **106**, 245502 (2011).
8. A. J. Liu, S. R. Nagel, The jamming transition and the marginally jammed solid. *Annu. Rev. Condens. Matter Phys.* **1**, 347–369 (2010).
9. L. Berthier, D. Coslovich, A. Ninarello, M. Ozawa, Equilibrium sampling of hard spheres up to the jamming density and beyond. *Phys. Rev. Lett.* **116**, 238002 (2016).
10. M. Ozawa, L. Berthier, D. Coslovich, Exploring the jamming transition over a wide range of critical densities. *SciPost Phys.* **3**, 1–30 (2017).
11. S. Sastry, P. G. Debenedetti, F. H. Stillinger, Signatures of distinct dynamical regimes in the energy landscape of a glass-forming liquid. *Nature* **393**, 554–557 (1998).
12. P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, F. Zamponi, Fractal free energy landscapes in structural glasses. *Nat. Commun.* **5**, 3725 (2014).
13. C. P. Royall, F. Turci, S. Tatsumi, J. Russo, J. Robinson, The race to the bottom: Approaching the ideal glass? *J. Phys. Condens. Matter* **30**, 363001 (2018).
14. L. Berthier, Dynamic heterogeneity in amorphous materials. *Physics (College Park Md.)* **4**, 42 (2011).
15. C. J. Cerjan, W. H. Miller, On finding transition states. *J. Chem. Phys.* **75**, 2800–2806 (1981).
16. D. J. Wales, Exploring energy landscapes. *Annu. Rev. Phys. Chem.* **69**, 401–425 (2018).
17. L. J. Munro, D. J. Wales, Defect migration in crystalline silicon. *Phys. Rev. B Condens. Matter Mater. Phys.* **59**, 3969–3980 (1999).
18. T. F. Middleton, J. Hernández-Rojas, P. N. Mortenson, D. J. Wales, Crystals of binary Lennard-Jones solids. *Phys. Rev. B Condens. Matter Mater. Phys.* **64**, 184201 (2001).
19. A. D. S. Parmar, B. Guiselin, L. Berthier, Stable glassy configurations of the Kob-Andersen model using swap Monte Carlo. *J. Chem. Phys.* **153**, 134505 (2020).
20. D. J. Wales, J. P. K. Doye, Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *J. Phys. Chem. A* **101**, 5111–5116 (1997).
21. A. Laio, M. Parrinello, Escaping free-energy minima. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 12562–12566 (2002).
22. A. Kushima *et al.*, Computing the viscosity of supercooled liquids. *J. Chem. Phys.* **130**, 224504 (2009).
23. W. Kob, H. C. Andersen, Testing mode-coupling theory for a supercooled binary Lennard-Jones mixture I: The van Hove correlation function. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **51**, 4626–4641 (1995).
24. C. A. Angell, B. E. Richards, V. Velikov, Simple glass-forming liquids: Their definition, fragilities, and landscape excitation profiles. *J. Phys. Condens. Matter* **11**, A75–A94 (1999).
25. E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, P. Gumbsch, Structural relaxation made simple. *Phys. Rev. Lett.* **97**, 170201 (2006).
26. P. Sollich, F. Lequeux, P. Hébraud, M. E. Cates, Rheology of soft glassy materials. *Phys. Rev. Lett.* **78**, 2020–2023 (1997).
27. P. Sollich, Rheological constitutive equation for a model of soft glassy materials. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **58**, 738–759 (1998).
28. P. Sollich, M. E. Cates, Thermodynamic interpretation of soft glassy rheology models. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **85**, 031127 (2012).
29. J. C. Gower, G. J. S. Ross, Minimum spanning trees and single linkage cluster analysis. *J. R. Stat. Soc. Ser. C Appl. Stat.* **18**, 54–64 (1969).
30. A. Christy, G. M. Gandhi, S. Vaithyasubramanian, Cluster based outlier detection algorithm for healthcare data. *Procedia Comput. Sci.* **50**, 209–215 (2015).
31. L. Berthier *et al.*, Configurational entropy measurements in extremely supercooled liquids that break the glass ceiling. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 11356–11361 (2017).
32. M. D. Ediger, M. Gruebele, V. Lubchenko, P. G. Wolynes, Glass dynamics deep in the energy landscape. *J. Phys. Chem. B* **125**, 9052–9068 (2021).
33. N. Xu, J. Bławdziewicz, C. S. O'Hern, Random close packing revisited: Ways to pack frictionless disks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **71**, 061306 (2005).
34. K. W. Desmond, E. R. Weeks, Random close packing of disks and spheres in confined geometries. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **80**, 051305 (2009).
35. S. D. Stoddard, J. Ford, Numerical experiments on the stochastic behavior of a Lennard-Jones gas system. *Phys. Rev. A* **8**, 1504–1512 (1973).
36. P. Grassberger, I. Procaccia, Characterization of strange attractors. *Phys. Rev. Lett.* **50**, 346–349 (1983).
37. A. Stukowski, Structure identification methods for atomistic simulations of crystalline materials. *Model. Simul. Mater. Sci. Eng.* **20**, 045021 (2012).
38. A. Stukowski, Visualization and analysis of atomistic simulation data with OVITO—the open visualization tool. *Model. Simul. Mater. Sci. Eng.* **18**, 015012 (2009).
39. T. F. Middleton, D. J. Wales, Energy landscapes of some model glass formers. *Phys. Rev. B Condens. Matter Mater. Phys.* **64**, 024205 (2001).
40. P. G. Debenedetti, F. H. Stillinger, Supercooled liquids and the glass transition. *Nature* **410**, 259–267 (2001).
41. F. Turci, C. P. Royall, T. Speck, Nonequilibrium phase transition in an atomistic glassformer: The connection to thermodynamics. *Phys. Rev. X* **7**, 031028 (2017).
42. U. H. Hansmann, Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.* **281**, 140–150 (1997).
43. G. Chen, C. K. Qu, P. Gong, Anomalous diffusion dynamics of learning in deep neural networks. *Neural Netw.* **149**, 18–28 (2022).
44. A. Thirumalaiswamy, R. A. Riggleman, J. C. Crocker, MIMSE algorithm modules. GitHub. <https://github.com/rar-ensemble/MIMSE>. Deposited 6 October 2022.