

Preparation history can outweigh descriptor definition in frictionless bidisperse packings below jamming

Quoc-Bao Nguyen

Hanoi University of Civil Engineering

Hoang-Quan Nguyen

University of Transport and Communications

Thai-Son Vu

`sonvt2@huce.edu.vn`

Hanoi University of Civil Engineering

Research Article

Keywords: Granular packing, Preparation protocol, Bidisperse spheres, Contact network, Continuum percolation, Homogenization

Posted Date: October 28th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11152620/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Preparation history can outweigh descriptor definition in frictionless bidisperse packings below jamming

Quoc-Bao Nguyen¹, Hoang-Quan Nguyen², Thai-Son Vu^{1,*}

¹ Faculty of Bridge and Road Engineering, Hanoi University of Civil Engineering, 55 Giai Phong, Bach Mai ward, 11618 Ha Noi, Viet Nam

² University of Transport and Communications, No. 3 Cau Giay Street, Lang Ward, Hanoi, Viet Nam

* Corresponding author: sonvt2@huce.edu.vn

ORCID: Quoc-Bao Nguyen 0000-0003-4406-2476; Hoang-Quan Nguyen 0000-0003-2348-4063; Thai-Son Vu 0009-0005-2491-6018.

Abstract

Descriptions of granular packings below jamming treat morphology, including descriptors built on small surface gaps, as a function of state variables such as size ratio, fine fraction and crowding. Jammed packings depend on preparation, but protocol-dependent percolation has been studied along two-dimensional compression paths near jamming, not on matched configurations at fixed density far below jamming. Here we hold every state variable and the particle number fixed in simulated bidisperse packings of frictionless spheres, change only the assembly route (two numerical algorithms), and compare the resulting shift of a fine-network connectivity descriptor with that caused by its own gap threshold. At size ratio 4, crowding 0.70 and a 5% gap threshold, the route moves the connectivity onset by 0.334 in fine fraction, whereas a tenfold change of the threshold moves the growth-prepared onset by 0.086. Preparation also exceeds the growth-prepared threshold shift at size ratios 6 and 8 (bounds only). At size ratios 4 and 8, the largest route difference lies in gaps below 2% of the fine diameter. A contact-scale history coordinate improves out-of-sample prediction of connectivity within a condition but does not transfer across crowding; at one matched condition, the routes also differ in homogenized shear stiffness. At size ratio 4 the preparation shift is about four times the growth-prepared threshold shift but no larger than the shake-equilibrated one under the same tenfold change, so both sensitivities should be checked in every preparation family before a contact descriptor enters a state-variable model.

Keywords: Granular packing; Preparation protocol; Bidisperse spheres; Contact network; Continuum percolation; Homogenization

Nomenclature

Symbols used in more than one section. Quantities without a unit are dimensionless.

Symbol	Meaning	Unit	Defined in
λ	size ratio of coarse to fine diameter	—	Sec. 2
ξ	fine volume fraction	—	Eq. (1)
φ	solid volume fraction	—	Eq. (1)
φ_m	maximum packing fraction	—	Sec. 2
s	crowding, φ divided by the maximum packing fraction	—	Eq. (1)
ε	distance from maximum packing, $1 - s$	—	Sec. 5
d_c, d_f	coarse and fine sphere diameters	length	Sec. 2
L	edge of the periodic cubic cell	coarse diameters	Sec. 2
N_c, N_f	numbers of coarse and fine spheres	—	Eq. (2)
t	connectivity gap threshold, fraction of the fine diameter	—	Eq. (3)
r_{ij}	minimum-image center distance of a fine pair	length	Eq. (3)
C_{max}	largest connected fine cluster	—	Eq. (3)
p_∞^f	largest-cluster fraction of the fine network	—	Eq. (3)
χ_{cl}	mean finite-cluster size	particles	Eq. (4)
h_{ij}	surface gap of fine pair i, j	length	Eq. (3)
δ	contact cutoff, fraction of the fine diameter	—	Eq. (5)
C_δ	contact fraction of near fine pairs	—	Eq. (5)
$C_{\delta,G}, C_{\delta,S}$	contact fraction of the growth-prepared and shake-equilibrated references	—	Eq. (6)
χ_h	history coordinate	—	Eq. (6)
Ψ	history-aware connectivity relation	—	Sec. 6
g_{ff}	fine–fine radial distribution function	—	Sec. 2
ξ_g	connectivity onset in fine fraction	—	Sec. 2
$\xi_{g,G}, \xi_{g,S}$	onsets of the growth-prepared and shake-equilibrated families	—	Sec. 3
$\Delta\xi_g$	history shift of the onset	—	Sec. 4
S_t	threshold sensitivity of the onset	—	Eq. (7)
S_p	preparation sensitivity of the onset	—	Sec. 3
R	ratio of preparation to threshold sensitivity	—	Eq. (7)
M	morphology descriptor	—	Eq. (8)
H	closure map from state variables and descriptor to effective stiffness	—	Eq. (8)
C_{eff}	effective stiffness	matrix stiffness	Eq. (8)
Φ_f^∞	fine-network capacity	—	Eq. (9)
μ_{hom}	homogenized shear modulus	matrix shear modulus	Sec. 7
μ_m	matrix shear modulus	stress	Sec. 7

1. Introduction

Structural and constitutive descriptions of granular packings rest on a working assumption, rarely tested directly, that fixing the macroscopic state fixes the packing morphology and any descriptor computed from it. For a bidisperse assembly the state is taken as the size ratio, fine fraction and crowding relative to maximum packing, and a morphology descriptor M is written $M(\lambda, \xi, s)$. Partial-coordination models of dense 2D bidisperse packings likewise use size ratio and fine fraction alone [1], and effective-property closures of particulate composites take the descriptor as a function of the same arguments [2–4]. Soil mechanics abandoned the corresponding assumption long ago: void ratio and stress do not determine the mechanical response of a sand, so a fabric tensor is carried as an additional argument [5]. We ask whether gap-defined contact descriptors require the same step, and if so, how large the missing argument is.

Fine-network connectivity is a natural test case: two fine spheres are linked when their surface gap falls below a small fraction of the fine diameter, and the largest-cluster fraction reports whether a system-spanning fine network exists. This continuum-percolation construction [6, 7] is used for transport in particulate composites [8] and is a candidate input for morphology-informed effective-property closures [3, 4]. It carries two methodological choices whose consequences have not been examined together: the value of the gap threshold and the assembly protocol. At fixed state variables, the question is whether the assembly route changes the descriptor more than a change of its own gap threshold does. We compare two algorithmic preparation families: growth-prepared (G) packings, grown from randomly seeded spheres, and shake-equilibrated (S) packings, obtained from G by shake–relax cycles at fixed solid fraction that drain the near-contact pairs left by growth.

The protocol is known to matter for jammed packings. The jamming density of soft and hard spheres depends systematically on preparation history [9–11], and granular packings retain long-lived memory of how they were built [12, 13]. Near-contact pair statistics are a known fingerprint of the packing state [14], and a protocol-independent characterization of random close packing has to be constructed rather than measured [15]. Structure is affected as well. Lubricated, vibrated and directly compressed assemblies of the same grains differ in coordination number, rattler fraction and fabric at comparable solid fraction [16]. Stress-controlled routes give the same frictional spheres states of equal volume fraction whose mean coordination numbers differ by up to one contact per particle [17]. Protocol-dependent structure is most pronounced in the loose regime [18], which contains all our states (crowding 0.60 to 0.70, 0.30 to 0.40 below maximum packing).

In jammed binary packings, whether the fine species carries load changes discontinuously with size ratio and composition. Prasad et al. [19] locate a subjamming transition at a critical size asymmetry near 5.75, above which the small spheres become load-bearing. Petit et al. [20] find a further discontinuous transition inside the jammed regime, emerging near size ratio 4.5, whose structural and pressure scaling has since been mapped [21, 22]. Hara et al. [23] obtain two jammed phases separated by a first-order line in the pressure–composition plane. That work and ours ask the same question about the fine network, and the answers are not in competition: those transitions are properties of the jammed state, whereas our states sit far below maximum packing. Our size ratios of 4, 6 and 8 straddle both critical asymmetries, and at all three the preparation shift exceeds the growth-prepared onset shift under a tenfold threshold change (Sec. 4), one reason we do not interpret the contrast as a signature of either asymmetry.

Keim et al. [24] and Paulsen and Keim [25] review how far-from-equilibrium disordered matter stores readable signatures of its history; the systems studied are almost always driven, with a

training amplitude written in and read out mechanically or dynamically [26, 27]. Our packings are quiescent, their only history is the assembly route, our readout is geometric and static, and shake-relax is an erasure trajectory rather than a training protocol. The history coordinate (Sec. 2) therefore measures how much of the assembly-route contact population survives, not a stored input that a suitable readout could recover.

Shen et al. [28] show, for frictionless bidisperse disks in two dimensions, that a contact percolation transition occurs well below jamming under quasi-static overdamped compression, insensitive to the compression step and the damping coefficient over the ranges tested. Kovalcinova et al. [29] find, for frictional polydisperse disks compressed between walls, that the contact-percolation area fraction rises from 0.776 to 0.798 as the compression speed is reduced to zero, where it coincides with jamming. At a given area fraction, slower compression makes percolation less likely, and they attribute the sub-jamming transition of [28] to overdamped dynamics. A protocol-dependent percolation transition at a common packing fraction is therefore established in two dimensions, along compression paths near jamming. To our knowledge, static three-dimensional configurations of identical particles at identical packing fraction that differ only in assembly route have not been compared far below jamming. Neither study, nor work on other percolation transitions in these systems (the connectivity of attractive assemblies [30], the force network of jammed ones [31]), asks how a protocol effect compares with that of the descriptor’s own arbitrary parameter.

We measure both sensitivities with all state variables and the particle number identical between G and S, so that no residual density difference can contribute to the contrast, and map both onsets directly at size ratio 4. At the reference threshold, the onset shift between the families (the history shift) exceeds the growth-prepared onset shift under a tenfold change of the gap threshold at every size ratio tested, as a lower bound at size ratios 6 and 8. At size ratio 4 and crowding 0.70 the history shift is about four times larger. The contrast replicates at a third size ratio, 6, held out of every fit (Secs. 4 and 6). At size ratios 8, 6 and 4, the largest structural difference is localized in gaps below 2% of the fine diameter, a window narrow enough that a generous threshold is largely blind to it. This localization is reconciled in Sec. 5 with the known protocol universality of near-contact pair correlations at jamming [32–34].

We then reduce the categorical protocol label to the history coordinate, a continuous contact-scale quantity measured on the configuration; adding it improves out-of-sample prediction of the connectivity within each condition. The reduction does not transfer across crowding. As a practical consequence (Sec. 7), an effective-stiffness closure of the form $C_{eff}=H(\lambda,\xi,s,M(\lambda,\xi,s))$ is not uniquely specified; the protocol-aware form carries the history coordinate as a fourth argument of M . At one matched condition, configurations that differ only in assembly route also differ in linear-elastic response at finite stiffness contrast.

The direct comparison of the two onset shifts is the test we propose for any contact descriptor; its normalized form per e-fold of the threshold, the ratio R (Sec. 3), allows comparison across size ratios and descriptors. For the connectivity examined here, preparation is the larger effect against the growth-prepared threshold sensitivity at every size ratio, whereas at size ratio 4, where the shake-equilibrated onset is mapped, R with that family as denominator falls below one per e-fold over the 5–10% band. Both sensitivities therefore have to be checked in every preparation family before a contact descriptor enters a state-variable model.

2. Numerical and morphological framework

Two numerical preparation routes produce packings with identical size ratio, fine fraction, solid fraction, crowding and particle number; below we define both routes and the quantities used to compare them. Later sections compare two shifts of the onset of a connectivity descriptor: the shift between the two routes (the preparation sensitivity) and the shift produced by the descriptor's own gap threshold (the threshold sensitivity).

The reference periodic cubic cell has edge $L = 4 d_c$ and contains N_c coarse and N_f fine spheres at size ratio $\lambda = d_c/d_f$ and fine volume fraction ξ . The spheres are frictionless, the configurations carry no load, and the relaxation step described below imposes only the non-overlap constraint. The crowding state is $s = \phi/\phi_m$, with ϕ the solid volume fraction. The maximum packing fraction $\phi_m(\lambda, \xi)$ is evaluated from the packing model as implemented by Vu et al. [4], which rests on the binary and polydisperse close-packing results of Hopkins et al. [35] and Farr and Groot [36]. This choice of model fixes only the meaning of s .

We treat s as a normalization for reporting rather than as a physically privileged coordinate, in part because the loose end of the accessible density range is protocol-bounded, not geometrically fixed [37]. The denominator of s is also protocol-sensitive for polydisperse hard spheres [38], which is one reason we do not treat the λ -ordering of the growth onsets as a purely geometric statement. Neither this sensitivity nor the choice of packing model affects the within- λ protocol comparisons, which are made at identical ϕ .

System sizes are $N_c = 16\text{--}48$ coarse spheres in the reference cell, with N_f up to $\approx 9,400$ fine spheres; the smaller coarse counts occur at the fine-rich end of the composition scan. In the finite-size endpoint sequence described below, N_c is 73 and 125. With single-particle volumes $v_c = \pi d_c^3/6$ and $v_f = \pi d_f^3/6$, and the volumes entering ξ taken as the totals over each species, the fine-fraction and crowding variables and the integer particle counts are

$$\xi = \frac{v_f}{v_f + v_c}, \quad s = \frac{\phi}{\phi_m} \quad (1)$$

$$N_c = \text{round}\left[\frac{(1-\xi)\phi L^3}{v_c}\right], \quad N_f = \text{round}\left[\frac{\xi\phi L^3}{v_f}\right] \quad (2)$$

Because each count is rounded to an integer independently, the realized ξ and ϕ differ from the pre-rounding targets, by a median of 0.0009 (maximum 0.0057) in ξ and 0.0015 (maximum 0.0041) in ϕ . All reported values of the state variables of Eq. (1) are computed from the integer counts actually generated.

Preparation. The growth-prepared endpoint (G) is produced by seeding $N_c + N_f$ spheres uniformly at random at reduced radii, with the particle counts fixed by Eq. (2). All radii are then grown stepwise to nominal size. After each growth step, a customized collective-rearrangement scheme removes overlaps, following the general strategy of classic packing generators [39, 40]. A periodic k-d-tree neighbor search under the minimum-image convention finds every overlapping pair, and each such pair is displaced symmetrically along its line of centers by a damped fraction (0.55) of the overlap. The sweep is iterated until the maximum relative overlap falls below 10^{-4} , for at most 3000 sweeps. The achieved ϕ is always verified against the target, and a realization that fails to reach the target ϕ within this budget is discarded. The verification is needed because an insufficient budget stalls the growth at an apparent, artificial packing ceiling.

Preparation history is then varied continuously by shake–relax cycles. Each Cartesian coordinate of every particle receives an independent Gaussian perturbation with standard deviation 0.10 of that particle’s radius, or 0.05 in the halved-amplitude variant. Positions are wrapped periodically, and the same relaxation is applied. The shake–relax cycle plays the role of tapping, with one difference. Tapping compacts a granular pile slowly and non-exponentially toward a vibration-intensity-dependent steady state [41, 42]. Our relaxation instead conserves ϕ exactly, so the shake–relax trajectory explores structure under repeated agitation at fixed density, without compacting.

Every relaxed state must pass two acceptance checks: the maximum residual overlap stays below the 10^{-4} acceptance, and the solid volume fraction is conserved exactly. The near-contact pair population is monitored only as a diagnostic, and no state is accepted or rejected on its basis. This population decreases between the two endpoints on every trajectory (Sec. 6).

Intermediate states along the trajectory (5–200 cycles) provide a continuous family of histories between the G endpoint and an operational shake-equilibrated reference endpoint (S). We use “shake-equilibrated” throughout as a label for the endpoint of the shake–relax operator applied for 200 cycles, not as a claim of thermodynamic equilibrium. At $\lambda = 4$ this 200-cycle reference is stable to 400 cycles; at $\lambda = 8$ its stability has not been demonstrated (Sec. 8). The relaxation operator is itself contact-generating: it separates overlapping pairs until they just touch. Sharper annealing toward an equilibrium hard-sphere reference at the same composition and density, for instance by swap Monte Carlo [43], would give a better-defined lower endpoint for the history coordinate χ_h of Eq. (6) and is the next step (Sec. 8).

Connectivity and diagnostics. The connectivity graph has the fine particles as vertices. An undirected edge joins i and j if and only if the gap rule of Eq. (3) holds under the periodic minimum-image distance. Eq. (3) states this rule and the connectivity measure built on the resulting graph, the largest-cluster fraction P_∞^f [44]; in it, r_{ij} is the minimum-image center distance and C_{max} the vertex set of the largest connected component.

$$h_{ij} = r_{ij} - d_f; \quad edge(i,j) \Leftrightarrow h_{ij} \leq t d_f; \quad P_\infty^f = \frac{|C_{max}|}{N_f} \quad (3)$$

The largest-cluster fraction, our measure of fine-network connectivity, is thus the share of fine particles that belong to the largest cluster, and it depends on the gap threshold t through the gap rule. Two distinct length parameters appear throughout: the connectivity gap threshold t and the contact cutoff δ . The threshold t is entered as a fraction of d_f , so the reference value 0.05 means 5% of d_f ; δ is defined in Eq. (5).

Three independent diagnostics are used. (i) The face-span status records whether the largest cluster has members within $2r_f$ (one fine diameter) of both opposite faces along any axis. It is a conservative proxy for periodic wrapping, not a topological test of it, and rigorous wrapping detection is not attempted. (ii) The susceptibility-like mean finite-cluster size χ_{cl} is given by Eq. (4), where n_k counts the clusters holding k fine particles.

$$\chi_{cl} = \frac{\sum' k^2 n_k}{\sum' k n_k} \quad (4)$$

In Eq. (4), the primed sums exclude the largest cluster when it is unique and retain all clusters when two or more tie for largest. The ratio is set to zero when no finite cluster remains, which happens

only in the fully connected limit. (iii) The third diagnostic is the mean graph degree. Fine–fine radial distribution functions $g_{ff}(r)$ probe pair statistics.

The continuum-percolation threshold of penetrable and near-contact sphere systems is known to high precision [45]. As an order-of-magnitude consistency check rather than a criterion, the per-family mean graph degrees reported in Sec. 4 bracket the mean degrees at which a continuum percolation transition is commonly found.

History coordinate. Let C_δ , Eq. (5), be the contact fraction of the near-pair universe: all fine–fine pairs $i < j$ with gap below $0.5 d_f$, under the minimum-image convention.

$$C_\delta = \frac{\#\{(i,j): h_{ij} / d_f < \delta\}}{\#\{(i,j): h_{ij} / d_f < 0.5\}} \quad (5)$$

We take $\delta = 0.02$, i.e. 2% of the fine diameter, as the reference cutoff. Here δ is only the contact cutoff, not the small-to-large radius ratio (the inverse of λ) that Petit and Sperl [21, 22] denote by δ . The history coordinate χ_h normalizes the contact fraction between the growth-prepared (G) and shake-equilibrated (S) references of the same (λ, ξ, s) condition, Eq. (6).

$$\chi_h = \frac{C_\delta - C_{\delta,S}}{C_{\delta,G} - C_{\delta,S}} \quad (6)$$

Here $C_{\delta,G}$ and $C_{\delta,S}$ are the contact fractions of the G and S references, so $\chi_h = 1$ at the G reference (full memory) and 0 at the S reference. The normalization is well posed: in every condition studied, $C_{\delta,G} - C_{\delta,S} \geq 0.18$.

Empirically, the G reference is reproducible across seeds: $C_{\delta,G} = 0.352$ and 0.349 for the $\lambda = 8$ and $\lambda = 4$ reference conditions, with seed-to-seed standard deviations of 0.0011 ($\lambda = 8$) and 0.0015 ($\lambda = 4$). At the seven compositions of the $\lambda = 4$ scan in Sec. 4 where both references were computed, $C_{\delta,G}$ stays within 0.349 – 0.361 for $\xi = 0.35$ – 0.70 at fixed crowding (Fig. A1a). Over this range the growth protocol therefore deposits a nearly fixed fraction of near pairs at contact, independent of the fine fraction.

The G reference is not independent of crowding: at $\lambda = 8$ it drifts $0.318 \rightarrow 0.329 \rightarrow 0.353$ for $s = 0.60, 0.65$ and 0.70 (Fig. A1b). χ_h is therefore normalized per condition. It is not clipped: stochastic fluctuations may place individual states outside the nominal $[0, 1]$ endpoint interval, and the halved-amplitude trajectory at $\lambda = 4$ reaches a history coordinate of -0.15 at 200 cycles.

The threshold scan uses three seeds per composition. Five full history trajectories were run for the $\lambda = 4$ condition and three for $\lambda = 8$ at $s = 0.70$; one of the $\lambda = 8$ trajectories was sampled densely over the first 30 cycles. The trajectory pool also holds trajectories at two further crowding levels, $s = 0.60$ (two seeds) and $s = 0.65$, and one halved-amplitude trajectory at $\lambda = 4$. It further holds a finite-size endpoint sequence at the $\lambda = 4$ condition with cell sizes $L/d_c = 4, 5$ and 6 , spanning $N_f = 1281$ – 4322 , with five seeds at the smallest size and two at each of the other two.

An independent $\lambda = 6$ replication campaign comprises a two-seed grown composition scan at $\xi = 0.15$ – 0.35 with $N_f = 1741$ – 4505 , two full trajectories at $\xi = 0.30$, and one shake-equilibrated state at $\xi = 0.35$. The whole campaign used an algorithm, tolerances, gap definition and χ_h normalization identical to those of the $\lambda = 4$ and $\lambda = 8$ runs.

The history coordinate is robust to its own defining cutoff. Recomputing χ_h with δ anywhere in 0.5–5% preserves the ordering of all states (rank correlation ≥ 0.96) and leaves the held-out prediction error of the test in Sec. 6 on a flat plateau (Fig. A2a). The reference $\delta = 2\%$ lies on that plateau.

Uncertainties follow one convention. A mean over seeds is written with $\pm$ and a standard deviation, and a bootstrapped interval as a range in parentheses. A bare range of a measured quantity is called a spread. Library versions, the random seeds of every realization and the neighbor-search implementation are recorded in the reproduction guide of the code repository listed under Data availability.

Onset definition. All percolation onsets reported below are operational finite-cell connectivity onsets, not thermodynamic thresholds. For each preparation family and gap threshold, ξ_g is the $P_\infty^f = 0.5$ crossing of the seed-mean connectivity curve, obtained by linear interpolation between adjacent scanned compositions. Onset confidence intervals resample the seeds with replacement independently at each composition (2000 replicates, percentile method). For the $\lambda = 4$ onset map, the bootstrap instead propagates a replicate-derived scatter model (Sec. 4).

Relaxation tolerance and the contact population. Residual overlaps within the 10^{-4} acceptance of the overlap-removal sweep fall in the contact bin of Eq. (5). The reference cutoff $\delta = 0.02$ integrates a window two hundred times wider. Appendix B shows that the growth-endpoint quantities do not follow the acceptance over five decades, so the contact population is resolved at finite precision and does not depend on where the solver stopped.

3. Threshold sensitivity of the onset

To separate the arbitrariness of the descriptor from the effect of preparation history, we first vary the gap threshold alone. Sec. 4 then changes only the assembly route and compares the two onset shifts. In the conditions tested, the threshold sensitivity of the growth-prepared onset is small and falls with size ratio from $\lambda = 4$ to $\lambda = 8$. The shake-equilibrated onset is mapped only at $\lambda = 4$, so its threshold sensitivity is reported with the $\lambda = 4$ onset map in Sec. 4.

Figure 1 shows P_∞^f against ξ for the growth-prepared family at $\lambda = 8$, $s = 0.70$, for five thresholds spanning a factor of ten. The onset is abrupt: P_∞^f rises from 0.02 to 0.91 between $\xi = 0.15$ and 0.20. The interpolated onset ξ_g moves only from 0.178 ($t = 2\%$) to 0.174 ($t = 20\%$) after rounding.

We define the threshold sensitivity as the finite-difference log-slope of the onset over the full measured range of t . Eq. (7) gives it together with the ratio R , which compares it with the preparation sensitivity defined below.

$$S_t = \frac{|\xi_g(t_2) - \xi_g(t_1)|}{\ln(t_2 / t_1)}, \quad R = \frac{S_p}{S_t} \quad (7)$$

S_t is normalized per e-fold of t , so quoting per decade would divide every value of R by $\ln 10$. Over the full range at $\lambda = 8$, the unrounded onsets differ by 0.0047, so $S_t = 0.0047 / \ln 10 \approx 0.0020$. A paired bootstrap, which resamples the seeds at each fine fraction with the same draw at every threshold, gives this shift in ξ_g a 95% interval of 0.0036–0.0060. The interval excludes zero, so S_t is resolved at $\lambda = 8$.

The preparation sensitivity $S_P = |\xi_{g,G} - \xi_{g,S}|$ is the onset shift between the two preparation families at the reference threshold $t = 5\%$; Sec. 4 measures it, and R is quoted at that threshold, with the growth-prepared threshold sensitivity as denominator, unless stated otherwise. R above one means that the assembly route moves the onset more than an e-fold change of the threshold does.

At $\lambda = 4$ a tenfold change of the threshold moves the growth-prepared onset by 0.086, from 0.354 at $t = 2\%$ to 0.268 at $t = 20\%$. There $S_t \approx 0.037$, about eighteen times its value at $\lambda = 8$; across the three size ratios S_t falls monotonically with λ , to 0.0217 at $\lambda = 6$ and ≈ 0.0020 at $\lambda = 8$. Table 1 collects the two sensitivities and their ratio at the three size ratios.




Figure 1. Largest-cluster fraction of the fine network against fine fraction for five gap thresholds (growth-prepared family, $\lambda = 8$, $s = 0.70$; error bars: realization scatter over three seeds per fine fraction). The dotted line marks the operational crossing level used throughout.

4. The protocol split at identical state variables

Sec. 3 changed only the gap threshold, over a factor of ten. To isolate the contribution of the assembly route, we now hold every state variable fixed, change only the preparation, and ask how far the fine-network connectivity moves. The preparation effect is larger than the threshold effect on the growth-prepared onset, stable in cell size at $\lambda = 4$, and present again at $\lambda = 6$; at $\lambda = 4$ the threshold moves the shake-equilibrated onset by a comparable amount.

Figure 2 compares the two endpoint preparations at identical size, fine fraction, crowding and seed lineage. At $\lambda = 8$, $\xi = 0.20$, $s = 0.699$, growth-prepared states give $P_\infty^f = 0.916$ and 0.906 (two seeds; a third grown realization gives 0.915). After 200 shake-relax cycles, the same particles at the same ϕ give 0.075 and 0.017 . At $\lambda = 4$ ($\xi = 0.35$, two seeds), the corresponding values are 0.714 and 0.738 against 0.021 and 0.036 . Across these four matched pairs at $\lambda = 8$ and $\lambda = 4$ the split is 0.69 to 0.89 in absolute largest-cluster fraction.




Figure 2. Largest-cluster fraction of the two preparation endpoints at identical state variables and system size, for six matched pairs, all at $s = 0.70$: two seeds at $\lambda = 8$ ($\zeta = 0.20$), two at $\lambda = 6$ ($\zeta = 0.30$) and two at $\lambda = 4$ ($\zeta = 0.35$). Annotations give the ratio of the two endpoint values.

At $\lambda = 4$ the split is not a small-cell artifact. Repeating the endpoint comparison at three cell sizes (Fig. 3) leaves the growth-prepared endpoint size-stable, with $P_\infty^f = 0.69 \pm 0.07$, 0.71 ± 0.01 and 0.66 ± 0.05 for $N_f = 1281$, 2501 and 4322. At all three sizes the per-family mean graph degree stays within 2.13–2.15 for G and 1.20–1.21 for S. The shake-equilibrated endpoint decreases with size ($0.033 \rightarrow 0.018 \rightarrow 0.014$), and its spanning probability drops to zero at the largest size, consistent with the expected finite-size behavior of a subcritical largest-cluster fraction. Consequently the endpoint contrast $\Delta P_\infty^f = 0.64$ – 0.69 is size-stable while the connectivity ratio of the two endpoints grows from $21\times$ to $47\times$ (Fig. 3b).




Figure 3. Finite-size behavior of the endpoint comparison at $\lambda = 4$, $\zeta = 0.35$, $s = 0.70$ (two seeds per size, five at $N_f = 1281$). (a) Largest-cluster fraction of both family endpoints versus particle number. (b) Endpoint contrast and connectivity ratio versus particle number.

At $\lambda = 4$ we scanned the shake-equilibrated family in fine fraction until its transition was crossed, so both onsets are measured rather than bounded (Fig. 4). Most scanned fine fractions carry one or two configurations. At those with a single configuration, a bootstrap that resamples only the observed replicates would propagate no uncertainty. We therefore quote a model-based interval obtained by bootstrap propagation of a replicate-derived scatter model. The model gives every fine fraction the scatter measured where replicates exist. That scatter is a per-configuration standard

deviation $\sigma = 0.10\sqrt{P_\infty^f(1 - P_\infty^f)}$, which is 0.05 at the crossing level; each bootstrap replicate perturbs the seed mean at a fine fraction with n configurations by Gaussian noise of standard deviation $\sigma/\sqrt{n}$. Replicates in which the shake-equilibrated curve never crosses 0.5 within the scan (about 0.5%) are discarded.

On that basis the growth-prepared onset is $\xi_{g,G} = 0.327$ (model-based 95% CI 0.323 to 0.331) and the shake-equilibrated onset is $\xi_{g,S} = 0.661$ (0.639 to 0.681). The shake-equilibrated family percolates at twice the fine fraction of the growth-prepared family, a directly measured history shift $\Delta\xi_g = 0.334$ (0.311 to 0.354). These intervals depend on the assumed scatter. When σ is increased by a factor of 1.5, which matches the scatter observed at the growth-prepared endpoint, the interval of the history shift widens to 0.274–0.361. If the replicates without a shake-equilibrated crossing are then kept, the upper limit of the shake-equilibrated onset lies beyond the scanned range.

The purely empirical alternative, resampling only the observed replicates, narrows the shake-equilibrated and shift intervals to 0.655–0.666 and 0.325–0.342 but slightly widens the growth-prepared interval to 0.323–0.332. The text quotes the model-based intervals. The empirical intervals are those plotted in Fig. 4 and, at $\lambda = 4$ and 6, in Fig. 5a; at $\lambda = 8$, Fig. 5a plots the paired-bootstrap interval of the growth-prepared onset.

At $\lambda = 4$ and $s = 0.70$, the history shift of 0.334, evaluated at the reference threshold $t = 5\%$, is 3.9 times the 0.086 by which a tenfold change of the gap threshold, from $t = 2\%$ to 20% , moves the growth-prepared onset (Sec. 3). This ratio is the $\lambda = 4$ value of R divided by $\ln 10$; per decade, R is 3.9 (8.9 per e-fold) at $\lambda = 4$, > 1.46 (> 3.3) at $\lambda = 6$ and > 3.9 (> 9) at $\lambda = 8$.

Of the four thresholds scanned at $\lambda = 4$ ($t = 2\%, 5\%, 10\%$ and 20%), the shake-equilibrated onset is mapped only at 5% and 10% , and between them it moves from 0.661 to 0.382, a shift of 0.279. With this shift as the denominator, $R = 0.83$ per e-fold over the 5–10% band; in raw terms the preparation shift is 1.2 times this threshold shift. The threshold sensitivity of the descriptor therefore itself depends on the preparation route.

The shake-equilibrated scan at $\lambda = 8$ lies outside the present study (its $\xi \geq 0.25$ states equilibrate an order of magnitude more slowly), so the shake-equilibrated onset there is reported as a lower bound. At $t = 5\%$ the growth-prepared onset at $\lambda = 8$ is $\xi_{g,G} = 0.177$, and the shake-equilibrated family remains far below percolation at $\xi = 0.20$ (Fig. 2). Hence $\xi_{g,S} > 0.20$ and $S_P > 0.023$. The ratio $R = S_P/S_t$ at $\lambda = 8$ then exceeds 9 when the growth-prepared threshold sensitivity is taken at its one-sided 95% paired-bootstrap upper bound; with its point value, the bound is ≈ 11 . Both are lower bounds, since $\xi_{g,S}$ is not directly mapped there.

The direct comparison of the two onset shifts on the same configurations is the primary result, and R , Eq. (7), is its normalized form for comparing size ratios and descriptors; only at $\lambda = 4$ are both S_t and the history shift $\Delta\xi_g$ measured directly.

R depends on the threshold at which its numerator, the history shift, is evaluated, while the denominator remains the full-range growth-prepared S_t of Eq. (7). At $\lambda = 4$ that dependence is strong. Evaluating both onsets over the measured threshold range gives $\xi_{g,G} = 0.354, 0.327, 0.292$ and 0.268 at $t = 2\%, 5\%, 10\%$ and 20% . The shake-equilibrated onset leaves the scanned range of ξ at both ends. At $t = 2\%$ that family percolates nowhere in the scan, so $\xi_{g,S} > 0.70$ and $R > 9.2$. At

$t = 5\%$ and 10% , where it is mapped, $R = 8.9$ and 2.4 . At $t = 20\%$ it percolates at every fine fraction scanned, so $\xi_{g,S} < 0.35$ and, if that onset still lies above the growth-prepared one, $R < 2.2$. Over the same tenfold change the shake-equilibrated onset therefore moves by more than 0.35 , at least as much as the history shift of 0.334 .

R falls toward $t = 20\%$ because the descriptor loses resolution there. At both outer thresholds it no longer resolves ξ for the shake-equilibrated family, whose largest-cluster fraction varies across the full scan by 0.026 at $t = 2\%$ and by 0.048 at $t = 20\%$, against 0.587 and 0.708 at the two intermediate thresholds; at $t = 20\%$ it also barely separates the two families (Appendix C). Within the band where the descriptor resolves ξ for both families, R with the growth-prepared denominator falls from 8.9 to 2.4 , and for that denominator preparation remains the larger of the two sensitivities throughout the band.




Figure 4. Directly mapped percolation onsets of both preparation families at $\lambda = 4$, $s = 0.70$ and the reference gap threshold $t = 5\%$. Points: individual seeds; lines: seed means; bands: empirical 95% intervals of the interpolated onsets, from resampling the observed seeds at each fine fraction.

A replication campaign at the intermediate size ratio $\lambda = 6$ reproduces the difference between the two endpoints but maps only the growth-prepared onset. We ran it after the $\lambda = 4$ and $\lambda = 8$ analyses were frozen, under identical definitions, and deliberately excluded $\lambda = 6$ from the predictive-validation pool of Sec. 6. A two-seed composition scan of the growth-prepared family locates $\xi_{g,G}(\lambda = 6) = 0.277$ by the same pre-defined interpolation (Fig. 5). Its interval, 0.276 – 0.278 , comes from resampling the two seeds at each fine fraction. That interval is not a population-level uncertainty. The growth-prepared onsets are therefore ordered $\xi_{g,G} = 0.177 < 0.277 < 0.327$ for $\lambda = 8, 6, 4$.

At $\xi = 0.30$, just above the growth-prepared onset, the connectivity ratio of the two endpoints is $27\times$ and $63\times$ for the two seeds (Fig. 2). Per-family mean graph degrees of 2.52 (G) and 1.23 (S) bracket the transition. The shake-equilibrated state at $\xi = 0.35$ remains subcritical after 200 cycles ($P_\infty^f = 0.046$). The shake-equilibrated onset at $\lambda = 6$ therefore lies above this fine fraction, $\xi_{g,S}(\lambda = 6) > 0.35$. With the unrounded growth-prepared onset 0.277 and S_t of Sec. 3, this gives the weak bound $R(\lambda = 6) > 3.3$.

At the reference threshold and per e-fold of the threshold, the preparation shift therefore exceeds the growth-prepared threshold sensitivity at all three size ratios, whereas at $\lambda = 4$ it is comparable

to the shake-equilibrated threshold sensitivity over the 5–10% band. With every state variable fixed, preparation alone shifts the directly mapped onset at $\lambda = 4$ by 0.334 in fine fraction.




Figure 5. Size-ratio systematics across $\lambda = 4, 6, 8$. (a) Growth-prepared onsets (measured) and shake-equilibrated onsets (measured at $\lambda = 4$; lower bounds, marked by arrows, at $\lambda = 6$ and 8). Error bars: empirical 95% intervals from resampling the seeds at $\lambda = 4$ and 6 ; paired-bootstrap 95% CI at $\lambda = 8$. (b) Threshold sensitivity S_t of the growth-prepared onset versus λ , over $t = 2\%$ to 20% (log scale).

Table 1. Threshold and preparation sensitivities of the fine-network percolation onset at the three size ratios, all at $s = 0.70$. $\xi_{g,G}$ and $\xi_{g,S}$ are the growth-prepared and shake-equilibrated onsets in fine fraction at a gap threshold of 5% . $S_p = |\xi_{g,S} - \xi_{g,G}|$ is the preparation sensitivity at $t = 5\%$, S_t the threshold sensitivity of Eq. (7) of the growth-prepared family over $t = 2\text{--}20\%$, normalized per e -fold of t , and $R = S_p/S_t$, likewise per e -fold. Parenthetical ranges are 95% intervals: at $\lambda = 4$ the model-based propagated intervals of Sec. 4, at $\lambda = 6$ the interval from resampling the two seeds of each fine fraction. At $\lambda = 6$ and $\lambda = 8$ the shake-equilibrated family was not mapped, so its onset and every quantity derived from it are lower bounds. At $\lambda = 8$, a paired bootstrap over seeds resolves the onset shift behind S_t , and R is the one-sided 95% lower bound from that bootstrap.

λ	$\xi_{g,G}$	$\xi_{g,S}$	S_p	S_t	R
4	0.327 (0.323–0.331)	0.661 (0.639–0.681)	0.334 (0.311–0.354)	0.0374	8.9
6	0.277 (0.276–0.278)	> 0.35	> 0.073	0.0217	> 3.3
8	0.177	> 0.20	> 0.023	0.0020	> 9

5. Where the difference sits: the contact region of the pair-gap distribution

To locate the difference within the pair statistics, we examine the near-pair gap distributions of the two preparation endpoints directly (Fig. 6), at $\lambda = 8$, $\xi = 0.20$, $s = 0.70$, and check the contact-scale contrast at $\lambda = 6$ and $\lambda = 4$. We find that the largest difference lies at contact, within the first 2% of the fine diameter, and that the two preparations differ by less at larger separations.

The growth-prepared state carries a pronounced contact peak: 35% of near fine–fine pairs sit below 2% of d_f . The growth history left these pairs at or near contact even though the packing sits far below maximum packing, and any finite threshold connects them (Fig. 6). Shake–relax drains the contact peak progressively: after 30 cycles the sub-2% fraction falls from 35% to 22%.




Figure 6. Near-pair surface-gap distributions of the two preparation endpoints ($\lambda = 8$, $\xi = 0.20$, $s = 0.70$): probability density of the fine-fine gap h divided by d_f , over pairs with gap below $0.5 d_f$. Solid blue bars: growth-prepared; hatched orange bars: shake-equilibrated. The dotted line marks the contact cutoff $\delta = 0.02$ and the dashed line the gap threshold of 5% of d_f .

Figure 7 compares $g_{ff}(r)$ between the two endpoints of the first matched pair (seed 1), and Table 2 lists the contact-bin values for both matched pairs at this condition. With bins of 0.025 fine diameters, as in Fig. 7, the contact-bin ratio of the growth state to its descendant is 2.1 and 2.6 in the first bin above contact, between one and 1.025 fine diameters, which is the bin Table 2 reports. The tallest bin of Fig. 7 lies just below one fine diameter and holds the pairs left at contact within the relaxation acceptance; the ratio in that bin is 2.4 for the first pair (30.3 against 12.4). Both ratios depend on the bin width, as they must for so narrow a peak. A bin-free count of fine-fine pairs with gap below $0.02 d_f$ gives pair-count ratios of 2.4 and 2.9. The excess does not survive the first bin: the ratio falls to 1.1 over the next 0.025 and below unity after that.

Table 2. Values of the fine-fine radial distribution function in the first bin above contact (between one and 1.025 fine diameters) for the two matched pairs at $\lambda = 8$, $\xi = 0.20$, $s = 0.70$. G and S denote the growth-prepared and shake-equilibrated endpoints. The last column counts pairs with gap below $0.02 d_f$ and needs no binning.

Pair	Contact bin, G	Contact bin, S	Ratio, bin 0.025	Ratio, bin 0.05	Pair-count ratio
1	18.3	8.7	2.1	1.7	2.4
2	17.6	6.8	2.6	2.0	2.9

Outside the contact region the two preparations also differ, but by less. For separations beyond 1.5 fine diameters, the absolute difference in g_{ff} never exceeds 0.54, with a mean of 0.25, on values of order 1.5–2.5. For the first matched pair, the contact-bin signal is $\Delta g / \langle g \rangle \approx 0.71$, against at most 0.36 and typically 0.13 at medium range. The contact-scale difference is therefore larger by a factor of two to six. This largest single difference, inside the first $0.02 d_f$ of the gap distribution, is also the one that the gap-thresholded graph integrates directly.

The $\lambda = 6$ pairs at $\xi = 0.30$, $s = 0.70$ (Sec. 4) reproduce the same contact-scale difference out of sample. There the contact-fraction ratio is 2.5: $C_{\delta,G} = 0.343$ against $C_{\delta,S} = 0.138$, with $C_{\delta,G}$ again close to the ≈ 0.35 family value. On the same bin convention the g_{ff} contact-bin ratio is 2.2, with a pair-count ratio of 2.4 and absolute far-field differences that stay below 0.26.

The localization holds as well at $\lambda = 4$ ($\xi = 0.35$, $s = 0.70$, five seeds), where both onsets are mapped. In the first bin above contact the fine–fine radial distribution function averages 6.68 for G against 3.60 for S, and the pair-count ratio is 2.23. In G, 35% of near pairs lie below 2% of the fine diameter, against 14% in S, and in all five seeds the largest difference between the two gap distributions falls in their first bin, within 1% of the fine diameter.

Our data show that, at fixed state variables far below jamming, the amplitude of the near-contact structure depends on the assembly route. The difference in amplitude is large enough to move a network observable across its transition. The near-contact structure itself is well characterized for soft-sphere packings approaching unjamming. The first peak narrows toward a delta function, with a power-law decay on its outer side [46], and Donev et al. [14] catalog the pair-correlation fingerprints.

Near-contact pair correlations are known to be universal at jamming: Charbonneau et al. [32] show that the scaling relations for particles that bear small forces and for particles nearly in contact are universal across preparation protocols and spatial dimensions. The associated small-gap exponent is constrained theoretically [33] and measured to high precision once localized buckling excitations are removed [34]. The amplitude difference reported here does not conflict with that body of work, for two reasons. First, the universality is a critical statement about the jamming limit. Every state analyzed here sits at crowding $s \leq 0.70$, i.e. at a distance $\varepsilon = 1 - s \geq 0.30$ from maximum packing, where the asymptotic scaling of the jamming limit is not expected to hold.

Second, our protocol dependence is in the amplitude of the near-contact population, not in any exponent; we do not measure a protocol-dependent small-gap exponent. Our connectivity descriptor is governed by an integral of the pair-gap distribution over a window two orders of magnitude wider than the relaxation acceptance, rather than by the asymptotic form of that distribution at zero gap.

The amplitude is the fraction C_δ of near pairs inside a fixed gap window, and the contact-fraction ratio of the two families is 2.2 to 2.6 across the seeds of the $\lambda = 8$ condition. It is smaller than the pair-count ratios of Table 2 (2.4 and 2.9) because the two families also differ in their number of near pairs. At $\lambda = 8$, $\xi = 0.20$, $s = 0.70$, the largest difference between the two preparations lies within the first 2% of the fine diameter, and the largest-cluster fractions of the two endpoints differ by ≈ 0.84 to 0.89 (Fig. 2).




Figure 7. Fine–fine radial distribution function g_{ff} of the growth-prepared and shake-equilibrated endpoints ($\lambda = 8$, $\xi = 0.20$, $s = 0.70$; seed 1, bins of 0.025 fine diameters). The inset magnifies the contact region; the tallest bin, just below one fine diameter, collects the pairs at contact.

6. The history coordinate

Section 5 placed the largest difference between the two families in the contact region of the pair-gap distribution. The history coordinate is the contact fraction of that region, normalized per condition between the growth-prepared and shake-equilibrated references. The measured coordinate is tested here as a replacement for the categorical protocol label. It orders and predicts connectivity within each of the two conditions tested, and a mapping fitted at one crowding level fails at another.

We follow the structure continuously along the erasure trajectory. Figure 8 tracks χ_h and P_∞^f along thirteen shake-relax trajectories in eight series. χ_h decays from 1 toward 0 with a common shape across conditions and seeds. At $\lambda = 4$ the decay is monotonic to within ± 0.05 . On the densely sampled $\lambda = 8$, $s = 0.70$ trajectory, the history coordinate rises by 0.37 between cycles 60 and 100 while the largest-cluster fraction rises from 0.67 to 0.86; no other trajectory reproduces this excursion. P_∞^f falls at condition-dependent rates. The $\lambda = 4$ condition, which sits just above its percolation onset, loses its network within ten cycles, whereas the $\lambda = 8$ condition decays gradually over a hundred.




Figure 8. History coordinate χ_h (left) and largest-cluster fraction (right) along shake-relax trajectories. Eight series cover thirteen trajectories: a five-seed mean at $\lambda = 4$, three individual seeds at $\lambda = 8$ and $s = 0.70$, and two individual seeds at $\lambda = 6$. The other two series are a two-seed mean at $\lambda = 8$ and $s = 0.60$ and a single trajectory at $\lambda = 8$ and $s = 0.65$. The shaded band is one standard deviation over the five $\lambda = 4$ seeds. The halved-amplitude trajectory at $\lambda = 4$ appears in Fig. 9, not here.

Figure 9 replots the same data as P_∞^f against χ_h . Within each condition the relation is monotone and consistent across seeds, with one exception: a transient excursion in the first $s = 0.60$ realization that the second does not reproduce. The pooled Spearman rank correlation over the sampled states is $\rho = 0.87$. States on the same trajectory are not independent. We therefore attach no nominal p-value; the inferential statement is the grouped out-of-sample validation below.

Across conditions the curves of Fig. 9 do not coincide. Connectivity therefore takes the form $P_\infty^f = \Psi(\lambda, \xi, s, \chi_h)$, with a condition-dependent Ψ that is neither unique nor universal. Connectivity still changes along each trajectory while λ , ξ and s stay fixed, and within the predictive model class tested here χ_h tracks that change as an additional structural coordinate. Within each condition, the categorical protocol label is thereby replaced by a continuous, measurable contact-scale coordinate along which intermediate states can be located.

Because χ_h is normalized by the two family references of its own condition, it presupposes those references rather than eliminating them. The unnormalized contact fraction C_δ is the protocol-free functional of a configuration, and χ_h is its dimensionless, per-condition presentation.

Of two protocol-free alternatives tested against criteria fixed in advance (Table C1), a k-nearest-neighbor graph saturates in every state, and a quantile-adaptive threshold separates the families at $\lambda = 8$ but hardly at $\lambda = 4$; neither is a fixed gap-defined descriptor.




Figure 9. Largest-cluster fraction against the history coordinate for the trajectories of Fig. 8. Pale purple circles on dashed lines: the five $\lambda = 4$ seeds; black squares, dark gray triangles and light gray inverted triangles: the three $\lambda = 8$ seeds at $s = 0.70$; blue diamonds and light blue pentagons: the two $\lambda = 6$ seeds; green filled X markers and brown plus signs on dotted lines: $\lambda = 8$ at $s = 0.60$ and 0.65 ; large orange crosses without lines: the halved-amplitude trajectory at $\lambda = 4$.

Three further tests sharpen the status of χ_h : out-of-sample prediction within a condition, transfer from one crowding level to another, and a halved shake amplitude. Appendix C reports further supporting checks, among them the cluster statistics along the trajectories (Figs. A3 and A4) and the interaction of the gap threshold with history (Fig. A2b).

The first test is out-of-sample prediction with grouped cross-validation. The cross-validation pool contains exactly eight trajectories, five at $(\lambda = 4, \xi = 0.35, s = 0.70)$ and three at $(\lambda = 8, \xi = 0.20, s = 0.70)$. Each fold holds out one entire trajectory, never individual states, so there are eight folds. In each fold we refit two models on the remaining trajectories of the same condition.

The state-only baseline, $P_\infty^f = F(\lambda, \xi, s)$, is the mean P_∞^f of the training states of that condition. By construction, this baseline is constant along each trajectory, because λ , ξ and s are exactly fixed there (Fig. A5). It is the null reference rather than an arbitrary restriction: if connectivity were a function of the state variables alone, it could not evolve along a fixed-state trajectory. The history-aware model is a per-condition monotone fit $P_\infty^f = \Psi(\chi_h)$ with no tuned hyperparameters, obtained by isotonic regression (pool-adjacent-violators, linear interpolation between fitted values, clamped extrapolation). The bootstrap resamples entire held-out trajectories with 4000 percentile-method replicates.

Over the 48 held-out states, the state-only baseline gives RMSE 0.298 and the history-aware model gives RMSE 0.064, an error reduction of 78% with bootstrap 95% confidence interval 72–88% (Fig. A5). The 78% reduction does not depend on the smoother (isotonic and logistic fits agree to 0.001 in RMSE). The G and S references of the history coordinate are condition means over all seeds, including the held-out trajectory. Recomputing them in each fold from the training trajectories alone gives identical RMSEs, reduction and interval: within a fold the history

coordinate is then an increasing affine function of the contact fraction, and isotonic regression is invariant under such a map.

Two controls show that the reduction belongs to a contact-scale coordinate rather than to the act of adding a fourth argument. A random per-state coordinate, put through the identical procedure, gives a median reduction of -6.5% over 200 draws, with 5th to 95th percentile -20% to $+2\%$; it predicts worse than the state-only baseline. The shake-cycle count, the protocol label that the history coordinate is meant to replace, gives a 36% reduction whether entered linearly or logarithmically, less than half the reduction obtained with the history coordinate. Within a single condition, the raw contact fraction C_δ and the mean graph degree predict comparably. The first test shows that a contact-scale coordinate suffices, but not that χ_h is the only one that does.

In the second test, the history-aware model fitted at $s = 0.70$ does not transfer to $s = 0.60$. For the $s = 0.60$ trajectories, χ_h is computed with that condition's own family references. Applied without recalibration, the $s = 0.70$ model mispredicts these trajectories with RMSE 0.56, larger than their entire range of P_∞^f . It fails because at $s = 0.60$ even the full-memory states sit below their percolation onset ($P_\infty^f = 0.08\text{--}0.11$ in the growth-prepared state, two seeds).

In the third test, halving the shake amplitude to 0.05 of the particle radius at $\lambda = 4$ produces states that fall on the same $P_\infty^f\text{--}\chi_h$ branch at equal χ_h (large orange crosses in Fig. 9). In this trajectory, halving the amplitude leaves the connectivity reached at a given history coordinate unchanged.

The $\lambda = 6$ replication trajectories, computed with their own family references, decay along their own monotone $P_\infty^f\text{--}\chi_h$ branch within realization scatter (Fig. 8; blue diamonds and pentagons in Fig. 9). The qualitative structure of the history coordinate replicates at a size ratio never used in constructing it.

Combining the directly measured onsets with the trajectory data gives a first history-conditioned percolation map in the (ξ, χ_h) plane (Fig. 10a). The percolation boundary runs from $\xi_{g,G} = 0.327$ at full memory to $\xi_{g,S} = 0.661$ at the shake-equilibrated endpoint. Three arrested intermediate-history states, 25 cycles from the growth-prepared endpoint, appear as diamonds in Fig. 10a and constrain the interior. At $\chi_h \approx 0.09$ to 0.14 , the boundary lies between $\xi = 0.50$ (connectivity 0.08) and $\xi = 0.60$ (connectivity 0.61), at or slightly below the dashed linear guide between the two measured onsets.

Along the erasure trajectories, connectivity changes at fixed density under driving alone; the closest existing analog is the reversible–irreversible transition of periodically sheared suspensions [26]. There the control parameter is the driving amplitude, whereas here it is the number of erasure cycles at fixed amplitude, and the halved-amplitude test, which leaves the connectivity at a given history coordinate unchanged, distinguishes the two situations. We regard the work of Bi et al. [47] as the closest conceptual antecedent of our construction. They show that the transitions between unjammed, fragile and shear-jammed states are governed by the fraction of force-bearing grains, a scalar measured on the contact network, rather than by density alone. Our construction has the same form, a measured scalar rather than a control parameter, but differs from the force-bearing fraction in three respects. Our scalar is geometric, since our configurations carry no load; it orders a connectivity descriptor, not a mechanical state; and in our frictionless, quiescent system it cannot be generated by shear and is instead deposited by assembly and drained by agitation. The ratio R

of Eq. (7) has no counterpart in that work, where a force fraction has no threshold parameter to vary.

Like the force-bearing fraction of Bi et al., a measure of local structure extracted by supervised learning orders rearrangement propensity better than density alone [48]. The search for additional state variables in granular matter is an old program; it begins with the compactivity of Edwards and Oakeshott [49]. χ_h is a coordinate rather than a state variable of that kind: it is normalized per condition and makes no claim to ensemble status. In the conditions tested here, the history-aware model reduces the out-of-sample error of the state-only baseline by 78% within a condition, but fitted at $s = 0.70$ it does not transfer to $s = 0.60$.

7. A mechanical consequence for morphology-informed closures

Secs. 3–6 established the central result of this work, that preparation changes fine-network connectivity at fixed state variables. The mechanical consequence followed here supports that result rather than forming a separate line of evidence, and it is limited to one matched condition ($\lambda = 4$, $\xi = 0.35$, $s = 0.70$, stiffness contrasts 10 and 30) with three seeds. We ask what the preparation dependence implies for effective-property closures built on a morphology descriptor, and whether the two preparation families also differ mechanically. Homogenizing matched pairs of states that share every state variable, we find the growth state 6.83% stiffer in shear than its own shake-equilibrated descendant, at contrast 10 and averaged over three shear modes.

A morphology-informed effective-property closure [3, 4] takes the effective stiffness to be a function of the state variables and of a morphology descriptor M . In the protocol-blind form assumed in Sec. 1, the state variables themselves fix M . M may be the fine-network connectivity (measured by the largest-cluster fraction), a cross-scale coupling measure not developed here, or the fine-network capacity of Eq. (9) below. The capacity falls by a factor of 12 to 54 between the two preparation endpoints of the eight cross-validation trajectories at fixed (λ, ξ, s) , five of them at $\lambda = 4$ and three at $\lambda = 8$. Because M is history-dependent, such closures cannot be uniquely specified by composition and crowding alone. A history argument added to M plays the role that the fabric tensor [5] plays alongside void ratio and stress in soil mechanics; here it serves a connectivity descriptor and is measured on the configuration rather than postulated.

Eq. (8) is the protocol-aware form of the closure; it carries the history coordinate as a fourth argument of M . Throughout, C_{eff} denotes the effective stiffness of the linear-elastic homogenization problem, and H is a map left undetermined here.

$$C_{eff} = H(\lambda, \xi, s, M(\lambda, \xi, s, \chi_h)) \quad (8)$$

Eq. (8) states which arguments a closure of this kind requires once M depends on history; it is not itself a constitutive law. Both sides are expressed in units of the matrix stiffness. The phase properties are held fixed within each calculation and are left out of the argument list, so the relation holds between dimensionless quantities at a fixed phase contrast and Poisson ratio.

For a mechanical measurement (Fig. 10b), we voxelize each growth state and its own shake-equilibrated descendant at $\lambda = 4$, $\xi = 0.35$, $s = 0.70$, with the spheres as a single inclusion phase in the matrix. The two members of a pair carry identical particles and match in ϕ to four decimals. We homogenize both by an FFT scheme [50] at stiffness contrast 10, with the settings given in Appendix D. A single shear mode does not characterize a finite cell, so we solve all three; Table 3 lists the moduli seed by seed.

Averaged over the xy, xz and yz modes, the growth state is stiffer than its own descendant by $6.83 \pm 0.04\%$. The xy mode alone gives $6.39 \pm 1.16\%$. Averaging the three modes reduces the seed-to-seed scatter from ± 1.16 to ± 0.04 points, so the scatter of the single mode is almost entirely the anisotropy of that mode (Fig. 10b). The growth states spread more between modes than their descendants do (Table 3); at this condition, the mode anisotropy is itself a preparation signature. The single-mode excess was also computed at contrast 30, where it is $20.9 \pm 3.1\%$.

Table 3. Homogenized shear moduli of the matched growth-prepared (G) and shake-equilibrated (S) states at $\lambda = 4$, $\xi = 0.35$, $s = 0.70$ and contrast 10. Moduli are in units of the matrix shear modulus. Upper block: three-mode means on the 64^3 grid, with the mode spread taken as the range over the xy, xz and yz modes. Lower block: seed 1 in the xy mode on three grids. Excesses are computed before the moduli are rounded.

Seed	Modulus, G	Modulus, S	Excess (%)	Excess, xy (%)	Mode spread, G (%)	Mode spread, S (%)
Three-mode means, 64^3						
1	2.87614	2.69269	6.81	5.89	1.28	0.34
2	2.88275	2.69880	6.82	5.57	1.76	0.49
3	2.88353	2.69801	6.88	7.72	2.40	0.97
Mean			6.83 ± 0.04	6.39 ± 1.16		
Grid convergence, seed 1, xy mode						
64^3	2.8520	2.6933	5.89			
96^3	2.7812	2.6254	5.94			
128^3	2.7423	2.5901	5.88			

The excess is insensitive to grid resolution and to voxel registration, and it is not a difference in rasterized solid content (Appendix D). The voxel grid does not resolve the contact scale. A fine particle spans two voxels of radius on the 64^3 grid and three on the 96^3 grid. The gap window of 2% of the fine diameter, which Sec. 5 identifies at $\lambda = 8$ (replicated at $\lambda = 6$ and 4), is therefore a small fraction of a single voxel. Placing even one voxel across that window would need an $800 \times 800 \times 800$ grid, beyond the voxel method used here. The calculation therefore sees the coarse-grained arrangement of the fine phase that the contact-scale rearrangement leaves behind. Within that limit, the measurement shows that the structural difference between the two families has a linear-elastic signature at $\lambda = 4$, $\xi = 0.35$, $s = 0.70$ and contrasts 10 and 30. The measurement compares two microstructures and does not validate any closure.

As a practical corollary, a closure calibrated on numerically generated microstructures and compared with experiments carries an implicit assumption about which preparation family the real material resembles. Placing a real material on the χ_h axis requires physical measurements and a calibration of both reference endpoints. As a minimum, we recommend making that assumption explicit and, where possible, constraining it by contact-sensitive measurements.

The fine-network capacity, defined. Let C_{max} be the largest fine cluster of Eq. (3) at the reference threshold t . The *fine-network capacity* is the volume fraction of fine solid carried by that cluster; Eq. (9) gives it for a single configuration.

$$\Phi_f^\infty = \xi \varphi P_\infty^f \quad (9)$$

By Delesse’s principle, this volume fraction equals the mean areal fraction of fine solid belonging to the largest connected cluster on a random cross-section of the cell. The capacity differs from the fraction of fine particles that P_∞^f counts only by the factor $\xi\varphi$, so at fixed ξ and φ it is proportional to the connectivity and is not an independent second descriptor.

To find which descriptor the modulus follows, we homogenize every state along three of the five $\lambda = 4$ trajectories, eighteen states in all, and rank them by each descriptor. The Spearman coefficient is 0.91 for the contact fraction of Eq. (5), 0.59 for the fine-network connectivity and, by construction, also 0.59 for the fine-network capacity of Eq. (9); see Fig. 11. These values are a descriptive ranking over pooled states that are not independent within a trajectory, and they carry no p-value.

The densely sampled $\lambda = 8$ trajectory is held out of this comparison. At the resolution used, its fine particles span one voxel of radius, too coarse to carry a mechanical signal. Its modulus varies by 3.8% over a range in which the connectivity falls from 0.92 to 0.07, consistent with that limit.

We repeat the ranking inside each trajectory, where the six states share a seed and a fine fraction and differ only in shake history. The contact fraction ranks higher in all three trajectories, with a median coefficient of 0.94 against 0.77 for the connectivity (0.83, 0.94 and 1.00 against 0.77, 0.84 and 0.49).

The ordering is also visible without the statistic. Ten shake cycles drop the connectivity by a factor of ten while the modulus moves by 2%. Over the following one hundred and ninety cycles the connectivity stays almost unchanged, while the modulus keeps falling in step with the contact fraction.

Zhao et al. [51] (preprint) report that soft composite solids trained mechanically near a shear-jamming boundary cross over from granular-like to biopolymer-like strain stiffening, driven by the structural memory of their contact network and without any nonlinearity of the matrix. That work and ours agree that the history of an inclusion contact network, and not only the volume fraction of inclusions, affects the mechanical response of the composite in the regime each study examines. Their networks are trained *near* shear jamming and respond nonlinearly, whereas our states are quiescent, far below jamming ($\varepsilon = 1 - s \geq 0.30$) and probed in linear response.

On a one-parameter erasure family such as the three $\lambda = 4$ trajectories of Fig. 11, any coordinate that decays monotonically will rank the modulus well, and any descriptor that saturates early will not. Over the eighteen states of Fig. 11, the shake-cycle count, the protocol label that the history coordinate is meant to replace, ranks the modulus at 0.96, above the contact fraction itself. The ranking therefore separates a saturating descriptor from non-saturating ones, and it identifies the contact fraction, a contact-scale integral, as a better candidate input for a closure than the gap-thresholded connectivity. The case for a contact-scale input in a closure of the form of Eq. (8) rests not on this ranking but on the contact-region localization of Sec. 5 and on the out-of-sample test of Sec. 6, in which the cycle count gives an error reduction of 36% against 78% for the measured history coordinate. At one matched condition ($\lambda = 4$, contrast 10, three seeds), the growth state is 6.83% stiffer than its shake-equilibrated descendant. Over eighteen states, the modulus ranks with the contact fraction (Spearman 0.91) more closely than with the connectivity (0.59), though less closely than with the shake-cycle count.




Figure 10. (a) History-conditioned percolation map at $\lambda = 4$, $s = 0.70$: states in the (ξ, χ_h) plane, filled where percolated. Stars mark the two measured onsets, diamonds the arrested intermediate-history states; the dashed line is a guide, not a fitted boundary. (b) Excess, in percent, of the homogenized shear modulus, μ_{hom}/μ_m , of each growth state over its own shake-equilibrated descendant at identical state variables, for three independent seeds. The columns give inclusion–matrix contrast 10 in a single shear mode and averaged over three shear modes, both at 64^3 , then the single mode at 96^3 and at contrast 30. Markers are seeds and the rule is their mean.




Figure 11. Homogenized shear modulus of every state along three of the five $\lambda = 4$ preparation trajectories at $\xi = 0.35$ and $s = 0.70$, contrast 10 on a 64^3 grid, averaged over the two voxel registrations. (a) Against the contact fraction of Eq. (5). (b) Against the largest-cluster fraction, on the same vertical axis. Color and marker shape identify the seed, and lines join the states of one trajectory in order of shake cycles.

8. Statistical robustness and scope

The results are limited by sampling, by the preparation algorithms, by cell size and by numerical resolution; each limit is stated below with the work that would lift it. Within these limits the data support the protocol contrast at fixed state variables, and the growth-prepared endpoint is insensitive to the relaxation acceptance used to build it.

Out-of-sample evaluation is grouped by trajectory, and its confidence intervals resample whole trajectories. The history coordinate is insensitive to its own cutoff over a tenfold range (Sec. 2). States generated at half the shake amplitude at $\lambda = 4$ fall on the same branch as the full-amplitude states, but after 200 cycles their contact fraction is 0.105, below the 0.14 of the full-amplitude shake-equilibrated states. The shake-equilibrated reference therefore depends on the shake amplitude, which limits the scope of the results.

The preparation families are two numerical assembly routes, and neither endpoint is calibrated against a physical processing route. The growth-prepared family keeps particles in contact at 30% below maximum packing, a configuration that purely repulsive suspensions do not reach on their own. We therefore treat the history axis as a structural coordinate, not as a model of mixing, shear or rest, and we do not extend the results to physical preparation protocols. A third, physically motivated protocol such as Lubachevsky–Stillinger compression is the next step for the preparation comparison.

Shake–relax cycling does not satisfy detailed balance, so the shake-equilibrated state is an operational reference at 200 cycles rather than a thermodynamic equilibrium. It is stable to 400 cycles at $\lambda = 4$, whereas its stability at $\lambda = 8$ has not been demonstrated (Appendix C). An independent anchor for the zero of the history coordinate, for instance hard-sphere Monte Carlo at matched solid fraction, would remove the residual algorithmic content of that endpoint.

The onsets are finite-cell connectivity onsets, interpolated at a largest-cluster fraction of 0.5 on the fine-fraction grid. Their quoted intervals propagate seed scatter; at $\lambda = 4$ they are model-based, and their width depends on the assumed scatter (Sec. 4). The interpolation adds an uncertainty that we do not quantify. Denser sampling of the fine fraction across the transition would tighten both the quoted intervals and the interpolation uncertainty.

Our finite-size evidence covers only size ratio 4 and only the endpoint contrast, which is size-stable across three cell sizes (Sec. 4). We do not attempt critical scaling, because extracting exponents would need a five-size sequence with matched seed counts. Finite-size scaling of the susceptibility peak (Appendix C) would need the same sequence. We did not attempt the corresponding sequence at size ratio 8, since the larger fine population equilibrates an order of magnitude more slowly.

The mechanical result is a measured difference at one matched condition and finite contrast. It rests on three seeds, three shear modes, two stiffness contrasts, two grids (three for one pair) and two voxel registrations. It is bounded by resolution rather than statistics, because the voxel grid does not resolve the two-percent gap window (Sec. 7). Attributing the stiffness excess of the growth-prepared states to load paths through individual near contacts would need a contact-resolved calculation. The same limit applies to the modulus ranking of Sec. 7, which identifies the better candidate input for a closure but neither a load-path mechanism nor the reason that input ranks better.

Regenerating the growth-prepared endpoint over five decades of relaxation acceptance, from 10^{-3} to 10^{-8} , spreads its contact fraction by 0.06% at $\lambda = 4$ and by 1.7% at $\lambda = 8$ (Table B1 in Appendix B). Within these bounds the data support the central claim that two numerical assembly routes at fixed size ratio, fine fraction, crowding and system size produce fine-network connectivities that differ across the transition. The difference is ordered by the history coordinate, which is measured on pair gaps within the first few percent of the fine diameter.

9. Conclusions

Below jamming, the morphology of the frictionless bidisperse sphere packings simulated here is not determined by size ratio, fine fraction and crowding alone. Two assembly routes applied to the same particles at the same solid fraction produce fine networks with different connectivity onsets. At $\lambda = 4$ the onsets differ by a factor of two in fine fraction. At the reference threshold, this preparation shift is larger than the onset shift of the growth-prepared family under a tenfold change

of the descriptor’s own gap threshold. At $\lambda = 4$, however, the shake-equilibrated onset moves so strongly with the threshold that, per e-fold over the 5–10% band where it is mapped, its threshold sensitivity is comparable to the preparation shift. In the matched pairs at $\lambda = 8$, repeated at $\lambda = 6$ and 4, the difference is concentrated in the near-contact region of the pair-gap distribution, most sharply inside the first 2% of the fine diameter. A gap-thresholded graph integrates that window directly. Beyond it the pair statistics also differ, by a relative amount several times smaller. In these packings, descriptors built on generous thresholds therefore register only weak signatures of the difference.

Replacing the protocol label by a measured history coordinate succeeds only in part. Within one crowding condition the coordinate orders states and improves the out-of-sample prediction of connectivity, but it does not transfer between conditions. A mapping fitted at $s = 0.70$ mispredicts the $s = 0.60$ trajectories by more than their whole range of connectivity. Neither of the two protocol-free alternatives tested in Sec. 6 gives a fixed gap-defined descriptor free of history dependence.

The preparation dependence also has a mechanical consequence (Sec. 7). At $\lambda = 4$ and $s = 0.70$, one matched condition with three seeds, the homogenized shear modulus orders with the contact fraction of the near-contact window and only loosely with the connectivity built on it. The shake-cycle count orders it at least as well, so this ranking at one condition is not evidence independent of the structural results.

One limit applies to all of these results. Both preparation families are numerical algorithms, and neither endpoint is calibrated against a physical processing route. Where a real material sits between the two families is not established here.

The part of this work offered for use beyond these packings is methodological. A morphology descriptor is normally validated by showing that it correlates with the property it is meant to inform. A descriptor intended to enter a state-variable closure must, beyond that correlation, also be tested for its sensitivity to the preparation route and to its own defining parameter. At $\lambda = 4$ the threshold sensitivity differs between the two families, so both sensitivities have to be checked in every preparation family. Both sensitivities are onset shifts, measured in fine fraction on the same configurations. The ratio R of Eq. (7) is the normalized form of that comparison, per e-fold of the threshold, and allows size ratios and descriptors to be compared. As Sec. 4 reports, R depends on the threshold at which the preparation shift is evaluated and on the family that supplies the threshold sensitivity.

What we propose for another material or another descriptor is this comparison of two sensitivities and the matched-preparation design behind it. The particular onsets and moduli measured here do not carry over, and the outcome itself has to be tested again there. The evidence behind this outcome comes from one material, frictionless spheres, prepared by two numerical assembly routes. The comparison of the two sensitivities was made at crowding $s = 0.70$ and three size ratios; the lower crowding levels enter only the tests of the history coordinate. The particle configurations, the preparation and analysis code, the homogenization solver and the scripts that recompute the onsets, sensitivities and cross-validation are openly available in the repository listed under Data availability, so a reader can run the test on a descriptor of their own choosing.

Appendix A. Supporting figures

The five figures below support statements made in the body of the paper. They are characterization and robustness checks, or additional views of a result already shown in the body, and none of them carries a claim of its own.




Figure A1. The growth-reference contact fraction $C_{\delta,G}$: (a) against fine fraction at fixed crowding ($\lambda = 4$ scan; the square marks $\lambda = 8$ at the same s); (b) against crowding at fixed fine fraction.




Figure A2. (a) Dependence of χ_h on its contact cutoff δ : rank correlation with the reference definition and held-out prediction error, both against δ . (b) Threshold $\times$ history interaction at $\lambda = 8$, $s = 0.70$: largest-cluster fraction versus χ_h for five gap thresholds.




Figure A3. Susceptibility-like mean finite-cluster size along shake-relax trajectories at $\lambda = 8$, $\xi = 0.20$, for three crowding levels and multiple seeds.




Figure A4. Rank-size representation of the cluster populations at three history states ($\lambda = 8$, seed 1): the two preparation endpoints and a mid-trajectory state.




Figure A5. Grouped leave-one-trajectory-out validation: predicted versus simulated largest-cluster fraction for the state-only model (red squares) and the history-aware model (blue circles) over 48 held-out states. The dotted line marks perfect prediction.

Appendix B. Relaxation acceptance and the contact population

This appendix asks whether the near-contact population behind the history coordinate is a feature of the packings or a record of where the relaxation stopped. It combines a scale argument with a direct test in which the growth-prepared endpoint is regenerated over five decades of acceptance. In that test the measured quantities do not follow the acceptance.

Residual overlaps within the relaxation acceptance of 10^{-4} are counted in the contact bin of Eq. (5). The inner scale of the near-contact population, the smallest gap it resolves, is therefore set by the acceptance and not by the physics. In the growth-prepared endpoint, about a quarter of near fine-fine pairs lie within $2 \times 10^{-4} d_f$ of contact, twice the acceptance. The corresponding effective log-log slope of C_δ between that gap and the reference cutoff is $\ln(0.352/0.25)/\ln(0.02/(2 \times 10^{-4})) \approx 0.07$. This slope is far below the slope of order unity that a pair correlation finite at contact would give.

The low slope does not, by itself, indicate an artifact. A genuinely arrested contact network also places pairs at zero gap. In that case a tighter acceptance would resolve the same population more finely and leave C_δ at $\delta = 0.02$ and the connectivity P_∞^f unchanged.

Three features of our data are consistent with a physical origin of the near-contact population. The reference cutoff $\delta = 0.02$ integrates a window two orders of magnitude wider than the acceptance, so C_δ is not a measurement at the acceptance scale. The growth-prepared reference is reproducible to ± 0.0015 between seeds (five at $\lambda = 4$, three at $\lambda = 8$) and flat to within 0.012 across $\xi = 0.35$ –0.70. A quantity governed by a stochastic stopping criterion would be neither. Along a shake–relax trajectory the acceptance is held fixed by construction. A change in the acceptance therefore can produce neither the decay of C_δ nor the decay of P_∞^f . Such a change also cannot place the halved-amplitude states on the same branch as the full-amplitude states.

Whether the growth-prepared endpoint itself is independent of the acceptance is a separate question, which the scale separation alone does not settle. We tested it directly, since the near-contact population could otherwise be taken as a record of where the solver stopped. We regenerated the growth protocol at accepted relative overlaps from 10^{-3} to 10^{-8} , with three seeds at each and all else fixed. The runs cover the two principal conditions ($\lambda = 8$, $\xi = 0.20$ and $\lambda = 4$, $\xi = 0.35$, both at $s = 0.70$), and Table B1 lists the outcome.

Table B1. Growth-endpoint quantities against the relaxation acceptance at the two principal conditions, both at $s = 0.70$. Growth endpoints were regenerated for this test. Entries are means $\pm$ standard deviations over three seeds. The residual is the largest relative overlap left after relaxation, over the three seeds.

Acceptance	Largest residual	C_δ	Mean degree	P_∞^f
$\lambda = 4, \xi = 0.35$				
10^{-3}	9.8×10^{-4}	0.3487 ± 0.0033	2.131 ± 0.035	0.672 ± 0.076
10^{-4}	9.8×10^{-5}	0.3486 ± 0.0026	2.128 ± 0.039	0.662 ± 0.074
10^{-6}	9.8×10^{-7}	0.3485 ± 0.0027	2.128 ± 0.039	0.663 ± 0.074
10^{-8}	9.9×10^{-9}	0.3485 ± 0.0027	2.128 ± 0.039	0.663 ± 0.074
$\lambda = 8, \xi = 0.20$				
10^{-3}	9.0×10^{-4}	0.3476 ± 0.0008	2.946 ± 0.061	0.893 ± 0.024
10^{-4}	9.9×10^{-5}	0.3474 ± 0.0049	2.883 ± 0.112	0.884 ± 0.011
10^{-6}	1.0×10^{-6}	0.3535 ± 0.0076	2.918 ± 0.120	0.884 ± 0.017
10^{-8}	9.9×10^{-9}	0.3495 ± 0.0035	2.921 ± 0.120	0.882 ± 0.032

Over the full five decades the residual overlap follows the acceptance, so the control does act on the relaxation. The quantities of the growth-prepared endpoint do not follow it. The contact fraction spreads by 0.06% at $\lambda = 4$ and by 1.7% at $\lambda = 8$. At $\lambda = 8$ this spread lies inside the seed-to-seed scatter of the regenerated runs. That scatter is itself well above the ± 0.001 –0.0015 scatter among the cached references used to normalize χ_h . Two of the three $\lambda = 8$ runs at the tightest acceptance exhausted the relaxation budget before formally converging, at residual overlaps of 7×10^{-9} and 1×10^{-8} . Their measured quantities agree with those of the other runs.

The acceptance bounds the gap at 10^{-4} of the fine diameter or below, while the window the descriptor integrates is 2% of it, two hundred times wider, and this separation of scales is why the growth-prepared endpoint does not follow the acceptance. The history coordinate is therefore not

a measure of solver convergence. The contact population is a feature resolved at finite precision rather than a record of where the solver stopped.

Appendix C. Supporting tests of the history coordinate

This appendix collects the supporting tests for Sec. 6: the cluster statistics along the trajectories, a second realization of the $s = 0.60$ trajectory, a third crowding level, the interaction of the gap threshold with history, and two protocol-free alternatives to the gap-threshold graph (Table C1). It closes with the stability of the shake-equilibrated endpoint beyond 200 cycles.

The cluster statistics give an independent signature of the connectivity loss along each trajectory. The susceptibility-like mean finite-cluster size χ_{cl} is $O(1)$ in the fully connected growth-prepared state. It passes through a maximum mid-trajectory, reaching ≈ 370 at $\lambda = 8$, before decaying. That maximum is consistent with crossing a percolation transition [44], here traversed by changing history at fixed (λ, ξ, s) rather than by changing the fine fraction. Without finite-size scaling of the peak, we attach no critical significance to its height or position. Both scatter substantially between realizations (Fig. A3). On the topology side, mid-trajectory states develop a broad population of intermediate clusters absent from both endpoints (Fig. A4).

A second realization of the $s = 0.60$ trajectory decays monotonically. It identifies the transient excursion in the first realization (Sec. 6) as a single-realization fluctuation near the onset.

The crowding axis carries three levels. At $\xi = 0.20$, the connectivity of the growth-prepared endpoint rises from 0.09 ($s = 0.60$, mean of two seeds) through 0.76 ($s = 0.65$) to 0.91 ($s = 0.70$). The percolation onset of the growth-prepared family in crowding therefore lies between $s = 0.60$ and 0.65 at this fine fraction. This bracket is the first slice of a history-conditioned (ξ, s) percolation map.

The gap threshold interacts with history. Recomputing P_∞^f over the full threshold range at every history state shows that the protocol split itself is a contact-scale phenomenon. Figure A2b shows this at $\lambda = 8$; the numbers below cover all four $\lambda = 4$ and $\lambda = 8$ matched pairs of Sec. 4. At $t = 2\%$ of d_f , the endpoint split ranges from $15\times$ to $123\times$, with a median near $30\times$. At $t = 20\%$ it nearly vanishes ($1.0\text{--}1.1\times$), because shaking moves pairs out of the contact region but not out of the 10–20% gap band.

The growth-prepared onset still moves little with the threshold. Descriptors built on generous thresholds, however, are largely blind to preparation memory in these packings. This blindness is consistent with the contact-region localization of Sec. 5, and the threshold dependence of R reported in Sec. 4 quantifies it.

Protocol-free alternatives. Two descriptors that adapt to the configuration or ignore the gap scale were tested against two criteria fixed in advance: at equal (λ, ξ, s) , the growth-prepared and shake-equilibrated endpoints should give similar values, and the descriptor should still resolve the fine fraction. The quantile-adaptive descriptor links two fine particles when their surface gap lies at or below the 25th percentile of the near-pair gaps of that configuration, counted up to a gap of 0.5 fine diameters. By construction it connects the same fraction of near pairs under either protocol. The k -nearest-neighbor descriptor links each fine particle to its $k = 2$ or 3 nearest fine neighbors and has no gap scale. Table C1 compares both with the gap-threshold graph at $t = 5\%$. The kNN largest-cluster fraction lies at 0.99–1.00 in every state, so it resolves neither history nor fine

fraction. The quantile descriptor separates the two families by a factor of 3.2–4.2 at $\lambda = 8$ but only 0.91–2.22 at $\lambda = 4$, and at $\lambda = 4$ the growth-prepared value falls below the shake-equilibrated one in two of five seeds. At $\lambda = 4$ the quantile threshold therefore nearly removes the history dependence, but it is then no longer a fixed gap-defined descriptor.

Table C1. Largest-cluster fraction of the fine network for the gap-threshold graph and the two protocol-free descriptors, at the two preparation endpoints G (growth-prepared, 0 cycles) and S (shake-equilibrated, 200 cycles). $\lambda = 4$: $\xi = 0.35$, $s = 0.70$, five seeds; $\lambda = 8$: $\xi = 0.20$, $s = 0.70$, three trajectories. G and S entries are means $\pm$ standard deviations over seeds or trajectories; the G/S column gives the spread of the per-seed or per-trajectory ratios.

Descriptor	$\lambda = 4$, G	$\lambda = 4$, S	$\lambda = 4$, G/S	$\lambda = 8$, G	$\lambda = 8$, S	$\lambda = 8$, G/S
Gap threshold, $t = 5\%$	0.693 ± 0.061	0.033 ± 0.011	12.7–33.9	0.912 ± 0.005	0.055 ± 0.027	12.2–53.9
Quantile threshold, $q = 25$	0.103 ± 0.040	0.072 ± 0.025	0.91–2.22	0.583 ± 0.076	0.153 ± 0.036	3.24–4.24
kNN graph, $k = 2$	0.995 ± 0.006	0.995 ± 0.002	0.99–1.01	0.999 ± 0.001	0.997 ± 0.001	1.00
kNN graph, $k = 3$	1.000 ± 0.000	1.000 ± 0.000	1.00	1.000 ± 0.000	1.000 ± 0.000	1.00

S endpoint beyond 200 cycles. At $\lambda = 4$, all thirteen shake-equilibrated states at 200 cycles were carried on to 300 and 400 cycles with the same shake amplitude and relaxation tolerance and new random seeds. They comprise the five trajectory endpoints at $\xi = 0.35$ and eight states of the composition scan. At 200, 300 and 400 cycles, the shake-equilibrated onset at $t = 5\%$ is 0.661, 0.638 and 0.648, a fluctuation of about 0.02 without a directed drift, and the contact fraction averaged over the seven compositions is 0.1404, 0.1432 and 0.1434. Along the five $\lambda = 4$ trajectories, the change of the contact fraction between 100 and 200 cycles is 1–9% of its total change. At $\lambda = 8$, convergence is not demonstrated. The densely sampled trajectory still has a largest-cluster fraction of 0.86 at cycle 100 and falls to 0.07 at cycle 200, and its history coordinate jumps by 0.37 between cycles 60 and 100 (Sec. 6). The shake-equilibrated endpoint is therefore an operational reference at 200 cycles, stable to 400 cycles at $\lambda = 4$.

Appendix D. FFT homogenization details

Linear-elastic FFT homogenization runs on the voxelized microstructure with periodic boundary conditions and strain-controlled macroscopic shear. Both phases have Poisson ratio 0.3, the inclusion–matrix shear-modulus contrast is 10 or 30, and the iteration is the fixed-point scheme of Moulinec and Suquet [50]. Grids are 64^3 and 96^3 , with a 128^3 grid for one matched pair. The iteration stops when the homogenized modulus is stationary to a relative 10^{-9} over two successive iterations. The choice of reference modulus, the convergence checks, the residual diagnostics and three checks on the stiffness excess are given below.

For the calculations of Sec. 7, the reference shear modulus of the fixed-point scheme is set to 0.6 times the inclusion modulus. This choice satisfies the convergence condition of the basic scheme and reaches the fixed point about eight times faster than a reference equal to the inclusion modulus itself. Reference moduli of 5.5 and 6.5 in matrix units, which differ by 18%, agree on the converged value to 0.04%. Every solve reported in Sec. 7 converged in 72 to 76 iterations.

Two residuals are monitored on the converged field. The first, the Hill–Mandel gap, compares the volume average of the local work with the product of the average stress and average strain; it stays at 2.3×10^{-3} . The second, the equilibrium residual, is the Fourier-space contraction of the wavevector with the stress divided by its own norm; it lies between 0.105 and 0.116. This second figure is a property of the diagnostic, not of the iteration. All of it comes from the three Nyquist planes of the transform, where the wavenumber assigned to the highest resolved mode does not

represent a derivative. Excluding those planes leaves an equilibrium residual below 4×10^{-8} and a Hill–Mandel gap below 3×10^{-9} on the same converged fields. The governing equations are therefore satisfied to the tolerance of the iteration wherever the transform resolves a gradient. We report the unfiltered figures as well because they are what the standard diagnostic returns.

Three checks bound the stiffness excess reported in Sec. 7. First, the excess is insensitive to resolution, although the moduli themselves are not. Averaged over the three seeds, the xy excess moves from 6.41% on the 64^3 grid (the mean of the two voxel registrations) to 6.56% on the 96^3 grid, while the moduli shift by 2.5%. Seed by seed, it moves from 5.82% to 5.94%, from 5.84% to 6.07% and from 7.58% to 7.67%. The pair carried to 128^3 (seed 1) holds its excess at 5.88%, within 0.06 points of its 96^3 value. Second, the excess is insensitive to voxel registration; a half-voxel displacement moves the modulus of a given state by 0.28% in the median and 0.68% at worst. Third, the excess is not a difference in rasterized solid content; the voxel fractions of the six states lie within 0.095% of one another.

Computation (data repository). The verification script computes Eq. (9) for any cached configuration. It obtains the fine cluster C_{max} by union–find on the Eq. (3) graph, using a periodic k-d-tree pair search, the minimum-image convention and the reference threshold $t = 0.05$. The expression $\Phi_f^\infty = \xi \varphi |C_{max}|/N_f$ uses the realized ξ and φ of that configuration. No parameter is fitted. Companion scripts write one row per configuration with Φ_f^∞ , together with P_∞^f , χ_{cl} , the mean graph degree and C_δ for reference. They also report the homogenized shear modulus of each state and regenerate Figs. 10b and 11, all from the same cached states.

Declarations

Funding

No funding was received for conducting this study.

Data availability

The data and code that support the findings of this study are openly available in the Zenodo repository at <https://doi.org/10.5281/zenodo.22936224>, together with a reproduction guide. The deposit contains the particle configurations, the preparation and analysis code, the FFT homogenization solver and a verification script. The version cited here adds 26 configurations, the 300- and 400-cycle continuations of the thirteen shake-equilibrated states at size ratio 4, and the scripts that recompute the results added in this revision. The onsets, sensitivities, cross-validation and Figs. 10b and 11 can be regenerated from the package; the source image of Fig. 10a is not included.

Code availability

The preparation and analysis code, the FFT homogenization solver, the verification script and the scripts added in this revision are openly available in the same Zenodo repository, <https://doi.org/10.5281/zenodo.22936224>.

Author contributions

Quoc-Bao Nguyen: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft. Hoang-Quan Nguyen: Formal analysis, Validation,

Visualization, Writing – review & editing. Thai-Son Vu: Conceptualization, Methodology, Validation, Resources, Writing – review & editing, Supervision, Project administration. All authors have read and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Ethics approval

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Use of generative AI

A generative AI assistant (Anthropic Claude) was used to assist with language editing and selected computational tasks. The authors formulated the research questions and mathematical arguments, reviewed and verified the references and numerical results, and take full responsibility for the content of this work.

References

1. Wang, J., Shi, C., Ling, D.: A two-stage modification of the Dodds model for partial coordination number and contact proportion in dense 2D bidisperse granular packings with large size ratios. *Granul. Matter* 28, 16 (2026). <https://doi.org/10.1007/s10035-025-01610-y>
2. Torquato, S.: *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer, New York (2002). <https://doi.org/10.1007/978-1-4757-6355-3>
3. Marcadon, V., Hervé, E., Zaoui, A.: Micromechanical modeling of packing and size effects in particulate composites. *Int. J. Solids Struct.* 44, 8213–8228 (2007). <https://doi.org/10.1016/j.ijsolstr.2007.06.008>
4. Vu, T.-S., Tran, B.-V., Nguyen, H.-Q., Chateau, X.: A refined Morphological Representative Pattern approach to the behavior of polydisperse highly-filled inclusion-matrix composites. *Int. J. Solids Struct.* 270, 112253 (2023). <https://doi.org/10.1016/j.ijsolstr.2023.112253>
5. Li, X.S., Dafalias, Y.F.: Anisotropic critical state theory: role of fabric. *J. Eng. Mech.* 138, 263–275 (2012). [https://doi.org/10.1061/\(ASCE\)EM.1943-7889.0000324](https://doi.org/10.1061/(ASCE)EM.1943-7889.0000324)
6. Pike, G.E., Seager, C.H.: Percolation and conductivity: a computer study. I. *Phys. Rev. B* 10, 1421–1434 (1974). <https://doi.org/10.1103/PhysRevB.10.1421>
7. Balberg, I., Anderson, C.H., Alexander, S., Wagner, N.: Excluded volume and its relation to the onset of percolation. *Phys. Rev. B* 30, 3933–3943 (1984). <https://doi.org/10.1103/PhysRevB.30.3933>

8. Ambrosetti, G., Grimaldi, C., Balberg, I., Maeder, T., Danani, A., Ryser, P.: Solution of the tunneling-percolation problem in the nanocomposite regime. *Phys. Rev. B* 81, 155434 (2010). <https://doi.org/10.1103/PhysRevB.81.155434>
9. Torquato, S., Truskett, T.M., Debenedetti, P.G.: Is random close packing of spheres well defined? *Phys. Rev. Lett.* 84, 2064–2067 (2000). <https://doi.org/10.1103/PhysRevLett.84.2064>
10. Chaudhuri, P., Berthier, L., Sastry, S.: Jamming transitions in amorphous packings of frictionless spheres occur over a continuous range of volume fractions. *Phys. Rev. Lett.* 104, 165701 (2010). <https://doi.org/10.1103/PhysRevLett.104.165701>
11. Bertrand, T., Behringer, R.P., Chakraborty, B., O’Hern, C.S., Shattuck, M.D.: Protocol dependence of the jamming transition. *Phys. Rev. E* 93, 012901 (2016). <https://doi.org/10.1103/PhysRevE.93.012901>
12. Kumar, N., Luding, S.: Memory of jamming: multiscale models for soft and granular matter. *Granul. Matter* 18, 58 (2016). <https://doi.org/10.1007/s10035-016-0624-2>
13. Dai, B.-B., Li, T.-Q., Deng, L.-J., Yuan, W.-H., Yang, J.: Memory of fabric anisotropy in the static packing of granular materials. *Granul. Matter* 25, 70 (2023). <https://doi.org/10.1007/s10035-023-01360-9>
14. Donev, A., Torquato, S., Stillinger, F.H.: Pair correlation function characteristics of nearly jammed disordered and ordered hard-sphere packings. *Phys. Rev. E* 71, 011105 (2005). <https://doi.org/10.1103/PhysRevE.71.011105>
15. Xu, N., Blawdziewicz, J., O’Hern, C.S.: Random close packing revisited: ways to pack frictionless disks. *Phys. Rev. E* 71, 061306 (2005). <https://doi.org/10.1103/PhysRevE.71.061306>
16. Agnolin, I., Roux, J.-N.: Internal states of model isotropic granular packings. I. Assembling process, geometry, and contact networks. *Phys. Rev. E* 76, 061302 (2007). <https://doi.org/10.1103/PhysRevE.76.061302>
17. Santos, A.P., Srivastava, I., Silbert, L.E., Lechman, J.B., Grest, G.S.: Protocol-dependent frictional granular jamming simulations: cyclical, compression, and expansion. *Front. Soft Matter* 3, 1326756 (2024). <https://doi.org/10.3389/frsfm.2023.1326756>
18. Silbert, L.E.: Jamming of frictional spheres and random loose packing. *Soft Matter* 6, 2918–2924 (2010). <https://doi.org/10.1039/C001973A>
19. Prasad, I., Santangelo, C., Grason, G.: Subjamming transition in binary sphere mixtures. *Phys. Rev. E* 96, 052905 (2017). <https://doi.org/10.1103/PhysRevE.96.052905>
20. Petit, J.C., Kumar, N., Luding, S., Sperl, M.: Additional transition line in jammed asymmetric bidisperse granular packings. *Phys. Rev. Lett.* 125, 215501 (2020). <https://doi.org/10.1103/PhysRevLett.125.215501>
21. Petit, J.C., Sperl, M.: Structural transitions in jammed asymmetric bidisperse granular packings. *Granul. Matter* 25, 43 (2023). <https://doi.org/10.1007/s10035-023-01329-8>
22. Petit, J.C., Sperl, M.: Pressure model and scaling laws in jammed bidisperse granular packings. *Granul. Matter* 27, 23 (2025). <https://doi.org/10.1007/s10035-024-01500-9>

23. Hara, Y., Mizuno, H., Ikeda, A.: Phase transition in the binary mixture of jammed particles with large size dispersity. *Phys. Rev. Res.* 3, 023091 (2021). <https://doi.org/10.1103/PhysRevResearch.3.023091>
24. Keim, N.C., Paulsen, J.D., Zeravcic, Z., Sastry, S., Nagel, S.R.: Memory formation in matter. *Rev. Mod. Phys.* 91, 035002 (2019). <https://doi.org/10.1103/RevModPhys.91.035002>
25. Paulsen, J.D., Keim, N.C.: Mechanical memories in solids, from disorder to design. *Annu. Rev. Condens. Matter Phys.* 16, 61–81 (2025). <https://doi.org/10.1146/annurev-conmatphys-032822-035544>
26. Cort  , L., Chaikin, P.M., Gollub, J.P., Pine, D.J.: Random organization in periodically driven systems. *Nat. Phys.* 4, 420–424 (2008). <https://doi.org/10.1038/nphys891>
27. Paulsen, J.D., Keim, N.C., Nagel, S.R.: Multiple transient memories in experiments on sheared non-Brownian suspensions. *Phys. Rev. Lett.* 113, 068301 (2014). <https://doi.org/10.1103/PhysRevLett.113.068301>
28. Shen, T., O’Hern, C.S., Shattuck, M.D.: Contact percolation transition in athermal particulate systems. *Phys. Rev. E* 85, 011308 (2012). <https://doi.org/10.1103/PhysRevE.85.011308>
29. Kovalcinova, L., Goulet, A., Kondic, L.: Percolation and jamming transitions in particulate systems with and without cohesion. *Phys. Rev. E* 92, 032204 (2015). <https://doi.org/10.1103/PhysRevE.92.032204>
30. Lois, G., Blawdziewicz, J., O’Hern, C.S.: Jamming transition and new percolation universality classes in particulate systems with attraction. *Phys. Rev. Lett.* 100, 028001 (2008). <https://doi.org/10.1103/PhysRevLett.100.028001>
31. Pathak, S.N., Esposito, V., Coniglio, A., Pica Ciamarra, M.: Force percolation transition of jammed granular systems. *Phys. Rev. E* 96, 042901 (2017). <https://doi.org/10.1103/PhysRevE.96.042901>
32. Charbonneau, P., Corwin, E.I., Parisi, G., Zamponi, F.: Universal microstructure and mechanical stability of jammed packings. *Phys. Rev. Lett.* 109, 205501 (2012). <https://doi.org/10.1103/PhysRevLett.109.205501>
33. Wyart, M.: Marginal stability constrains force and pair distributions at random close packing. *Phys. Rev. Lett.* 109, 125502 (2012). <https://doi.org/10.1103/PhysRevLett.109.125502>
34. Charbonneau, P., Corwin, E.I., Parisi, G., Zamponi, F.: Jamming criticality revealed by removing localized buckling excitations. *Phys. Rev. Lett.* 114, 125504 (2015). <https://doi.org/10.1103/PhysRevLett.114.125504>
35. Hopkins, A.B., Stillinger, F.H., Torquato, S.: Densest binary sphere packings. *Phys. Rev. E* 85, 021130 (2012). <https://doi.org/10.1103/PhysRevE.85.021130>
36. Farr, R.S., Groot, R.D.: Close packing density of polydisperse hard spheres. *J. Chem. Phys.* 131, 244104 (2009). <https://doi.org/10.1063/1.3276799>
37. Pica Ciamarra, M., Coniglio, A.: Random very loose packings. *Phys. Rev. Lett.* 101, 128001 (2008). <https://doi.org/10.1103/PhysRevLett.101.128001>

38. Baranau, V., Tallarek, U.: Random-close packing limits for monodisperse and polydisperse hard spheres. *Soft Matter* 10, 3826–3841 (2014). <https://doi.org/10.1039/C3SM52959B>
39. Jodrey, W.S., Tory, E.M.: Computer simulation of close random packing of equal spheres. *Phys. Rev. A* 32, 2347–2351 (1985). <https://doi.org/10.1103/PhysRevA.32.2347>
40. Lubachevsky, B.D., Stillinger, F.H.: Geometric properties of random disk packings. *J. Stat. Phys.* 60, 561–583 (1990). <https://doi.org/10.1007/BF01025983>
41. Knight, J.B., Fandrich, C.G., Lau, C.N., Jaeger, H.M., Nagel, S.R.: Density relaxation in a vibrated granular material. *Phys. Rev. E* 51, 3957–3963 (1995). <https://doi.org/10.1103/PhysRevE.51.3957>
42. Nowak, E.R., Knight, J.B., Ben-Naim, E., Jaeger, H.M., Nagel, S.R.: Density fluctuations in vibrated granular materials. *Phys. Rev. E* 57, 1971–1982 (1998). <https://doi.org/10.1103/PhysRevE.57.1971>
43. Ninarello, A., Berthier, L., Coslovich, D.: Models and algorithms for the next generation of glass transition studies. *Phys. Rev. X* 7, 021039 (2017). <https://doi.org/10.1103/PhysRevX.7.021039>
44. Stauffer, D., Aharony, A.: *Introduction to Percolation Theory*, 2nd edn. Taylor and Francis, London (1994).
45. Rintoul, M.D., Torquato, S.: Precise determination of the critical threshold and exponents in a three-dimensional continuum percolation model. *J. Phys. A: Math. Gen.* 30, L585–L592 (1997). <https://doi.org/10.1088/0305-4470/30/16/005>
46. Silbert, L.E., Liu, A.J., Nagel, S.R.: Structural signatures of the unjamming transition at zero temperature. *Phys. Rev. E* 73, 041304 (2006). <https://doi.org/10.1103/PhysRevE.73.041304>
47. Bi, D., Zhang, J., Chakraborty, B., Behringer, R.P.: Jamming by shear. *Nature* 480, 355–358 (2011). <https://doi.org/10.1038/nature10667>
48. Cubuk, E.D., Schoenholz, S.S., Rieser, J.M., Malone, B.D., Rottler, J., Durian, D.J., Kaxiras, E., Liu, A.J.: Identifying structural flow defects in disordered solids using machine-learning methods. *Phys. Rev. Lett.* 114, 108001 (2015). <https://doi.org/10.1103/PhysRevLett.114.108001>
49. Edwards, S.F., Oakeshott, R.B.S.: Theory of powders. *Physica A* 157, 1080–1090 (1989). [https://doi.org/10.1016/0378-4371\(89\)90034-4](https://doi.org/10.1016/0378-4371(89)90034-4)
50. Moulinec, H., Suquet, P.: A numerical method for computing the overall response of nonlinear composites with complex microstructure. *Comput. Methods Appl. Mech. Eng.* 157, 69–94 (1998). [https://doi.org/10.1016/S0045-7825\(97\)00218-1](https://doi.org/10.1016/S0045-7825(97)00218-1)
51. Zhao, Y., Pan, D., Pang, Y., Barés, J., Xu, C., Liu, C., Hu, H., Jin, Y., Xu, Q.: Programming strain-stiffening in soft composites via structural memory near jamming. Preprint at arXiv:2604.20437 (2026). <https://arxiv.org/abs/2604.20437>