communications materials

Article

https://doi.org/10.1038/s43246-024-00583-4

Designing athermal disordered solids with automatic differentiation

Check for updates

Mengjie Zu & Carl P. Goodrich

The ability to control forces between sub-micron-scale building blocks offers significant potential for designing new materials through self-assembly. Traditionally, this involves identifying a crystal structure with a desired property and then designing building-block interactions so that it assembles spontaneously. However, this paradigm fails for structurally disordered solids, which lack a well-defined structure. Here, we show that disordered solids can still be treated from an inverse self-assembly perspective by bypassing structure and directly targeting material properties. Using the Poisson's ratio as a primary example, we demonstrate how differentiable programming links interaction parameters with emergent behavior, enabling iterative training to achieve the desired Poisson's ratio. We also tune other properties, including pressure and local 8-fold structural order, and can even control multiple properties simultaneously. This robust, transferable, and scalable approach can handle a wide variety of systems and properties, demonstrating the utility of disordered solids as a practical avenue for self-assembly platforms.

To what extent can the material properties of disordered solids be controlled? Recent results show that the properties of random spring networks close to the isostatic point can be radically and precisely tuned through slight adjustments to the network topology[1–5]. More specifically, the Poisson's ratio, ν, of a randomly generated network can be tuned to either the upper or lower bounds by removing only ~1% of the springs—the choice of which springs to remove determines the final value of ν[1]. This is possible because each spring's contribution to the bulk modulus is independent of its contribution to the shear modulus, meaning that their ratio, which determines ν, can be tuned by the choice of removed springs. However, in a real material, it is usually not possible to make precise, targeted alterations to structure. In fact, the notion of removing springs from a network simply does not translate to materials made up of particulate building blocks, ranging from atoms or molecules to larger objects like proteins or colloids. Unlike a spring network, one cannot simply remove the interaction between two particular particles, let alone do so in a scalable way.

Nevertheless, this paper proposes a strategy for tuning the properties of particulate-based disordered solids. This strategy is akin to an inverse self-assembly approach, except that interaction parameters are adjusted to tune properties rather than structure. The necessary connection between interaction parameters and emergent properties is made by exploiting a class of numerical techniques called Automatic Differentiation (AD), enabling the exploration of high-dimensional and complicated design spaces. Once this connection is made, the system is "trained" similar to how one trains a neural network.

We demonstrate this strategy by considering the simple case of athermal sticky spheres. Specifically, we consider a two-dimensional system of N particles divided evenly into $n_{\rm sp}$ species, where the diameter of each species, σ_α, and the binding energy between each pair of species, $B_{\alpha\beta} = B_{\beta\alpha}$, can be continuously varied. Motivated by DNA-coated colloids, the particles interact via a Morse potential with the short-ranged repulsive part replaced with a finite soft repulsion (see Fig. 1a and Methods)[6,7]. The system is prepared at zero temperature following the protocol developed in the study of the jamming of soft spheres[8,9], where particles are placed randomly (corresponding to infinite temperature) and then quenched to the nearest local energy minimum.

The final material properties, such as the Poisson's ratio, clearly depend on the $n_{\rm sp} + n_{\rm sp}(n_{\rm sp}+1)/2$ values of σ_α and $B_{\alpha\beta}$, but we are left with two critically important questions. First, can these parameters be adjusted in order to accurately and precisely control material properties? For example, how close can ν be tuned to a particular target ν^*? Second, can multiple properties be controlled simultaneously and independently, enabling highly nontrivial design? We will show that, even for our exceedingly simple model, the answers to both questions are a definitive yes.

This result is obtained by directly connecting the $n_{\rm sp} + n_{\rm sp}(n_{\rm sp}+1)/2$ model parameters to changes in material properties. More precisely, we construct an objective function $\mathcal{L}$, e.g. $\mathcal{L} = (\nu - \nu^*)^2$, that indicates how far a property (e.g. ν) is from a target (e.g. ν^*), and use Automatic Differentiation (AD)[10–12] to calculate the gradient $\nabla_\theta \mathcal{L}$ of $\mathcal{L}$ with respect to the parameters $\theta = \{\sigma_\alpha\} \cup \{B_{\alpha\beta}\}$. $\nabla_\theta \mathcal{L}$ indicates how changes to the parameters affects the objective, and there are numerous gradient-descent-based

Institute of Science and Technology Austria, Am Campus 1, 3400 Klosterneuburg, Austria. e-mail: carl.goodrich@ist.ac.at

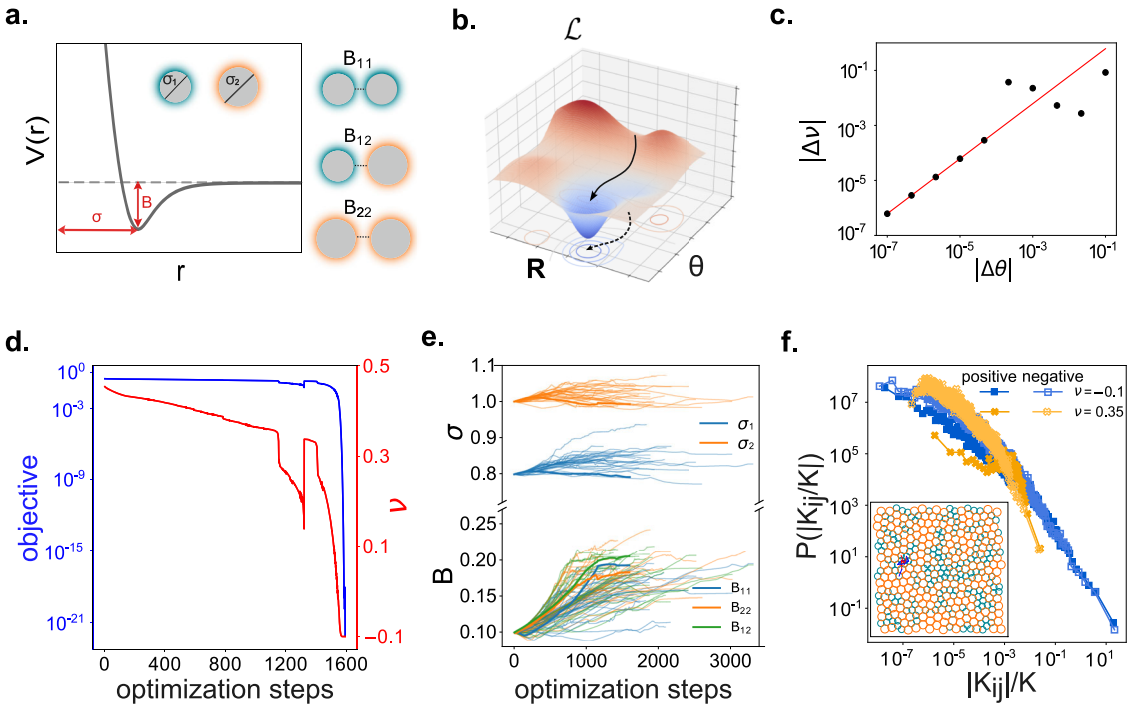

Fig. 1 | Tuning the Poisson's ratio of a single configuration. a A schematic representation of the interactions between particles. The particles interact via the harmonic-Morse potential, providing two important design parameters, their diameters σ and binding energy between them B, as discussed in the main text. The dashed gray line denotes zero. **b** Schematic of the landscape of the objective function $\mathcal{L}$ in configuration space **R** and parameter space θ. Colored regions with contour lines represent the basins of attraction of distinct minima. The black arrow represents the configuration during the optimization process. **c** The gradient $\nabla_\theta \nu_b$ accurately predicts changes in the Poisson's ratio due to small changes in the model parameters. The red line shows the prediction $|\nabla_\theta \nu_b \cdot \Delta\theta|$, where $\Delta\theta$ is the change in parameters in an arbitrarily chosen direction, and the points show the measured change in the Poisson's ratio $|\nu_b(\Delta\theta) - \nu_b(\Delta\theta = 0)|$ measured at $\Delta\theta$. Data is shown for a bidisperse system of 368 particles and binding energies 0.1, but the result is generic. **d** The evolution of the objective $\mathcal{L} = (\nu_b - \nu^*)^2$ and Poisson's ratio ν for an example optimization with target $\nu^* = -0.1$. The optimization is successful and converges in slightly less than 1600 optimization steps. Note: the spike around 1300 optimization steps corresponds to a discontinuous change in the position of the local minimum due to a structural rearrangement. **e** The evolution of the parameters for 35 successful optimizations. The thick lines correspond to the example shown in (**d**). The number of optimization steps is variable because the process terminates after convergence. **f** The distribution of how each pair of particles contributes to the bulk modulus after training with $\nu^* = -0.1$ (blue) and $\nu^* = 0.35$ (orange). When $\nu^* < 0.2$ (see Supplementary Fig. 3 for more comprehensive data), the positive and negative tails of this distribution broaden significantly. The positive and negative values are shown separately to demonstrate that, in this long tail, $P(K_{ij}) \approx P(-K_{ij})$. The inset shows the packing after training with $\nu^* = -0.1$ and highlights the 10 bonds with the largest $|K_{ij}|$, with blue (red) lines indicating positive (negative) K_{ij}. This shows that the broad tails are spatially localized.

algorithms[13] for using this information to minimize $\mathcal{L}$, thus tuning the property of interest (Fig. 1b). See "Methods" for more details.

At zero temperature, the N particles arrange themselves into one of many possible "configurations," or local minima of the potential energy landscape. In this paper, we will manipulate material properties first on the level of individual configurations, meaning we will pick a random configuration b and tune the parameters until, e.g., $\nu_b = \nu^*$. We find that this works surprisingly well all the way to the perfectly auxetic limit $\nu^* \to -1$. Interestingly, the Poisson's ratio of multiple configurations can also be tuned precisely and simultaneously, although this becomes more challenging as more configurations are optimized at the same time. We then consider the ensemble level, and again find a surprising ability to manipulate the average Poisson's ratio $\langle\nu\rangle$ within the range $0.2 < \langle\nu\rangle < 0.7$. Importantly, this approach is highly general, and we use it to tune other properties and even multiple properties simultaneously.

Results

Training individual configurations

To begin, we select a configuration b by placing N particles randomly in a 2d periodic box and minimizing the energy to the nearest minimum. We then define the objective $\mathcal{L} = (\nu_b - \nu^*)^2$, where ν_b is the measured Poisson's ratio and ν^* is the desired target chosen from between -1 and 1, which are the theoretical bounds for isotropic systems in two dimensions. Figure 1c confirms that our AD-based calculation of $\nabla_\theta \nu_b$ accurately predicts the change in ν_b over finite changes in the parameters θ. This is shown for a particular representative example where $N = 368$ particles are evenly divided into $n_{sp} = 2$ species with initial diameters of 0.8 and 1.0, a number density of $\rho = 1.6$, and with constant binding energies $B_{\alpha\beta} = 0.1$. Before training, we measure $\nu_b \approx 0.453$. Next, we train the system using standard gradient-descent-based algorithms (see Methods) to iteritively adjust θ before recalculating ν_b and $\nabla_\theta \nu_b$. At each step, we reminimize the energy with respect to the particle positions, allowing us to track the configuration as parameters change. Figure 1d shows how $\mathcal{L}$ and ν change during 1600 iterations. Despite only having 5 parameters, the Poisson's ratio is tuned exactly to the target, meaning that the Poisson's ratio at this configuration is successfully and accurately controlled.

Figure 1e shows how the parameters evolve during training for 35 successfully trained configurations. Clearly, the final optimized parameters depend strongly on the final configuration, which depends both on the initial configuration and the details of the optimization protocol. In all cases, the systems remain bidisperse and thus disordered, but there is no obvious trend in the relative particle diameters nor in the three binding energies.

For ν^* less than approximately 0.2 or greater than approximately 0.7, the final Poisson's ratio seems to be dominated by a small localized region. To see this, we calculate K_{ij} and G_{ij}, which are the contributions of the interaction between particles i and j to the bulk and shear moduli, respectively, so that $K = \sum_{ij} K_{ij}$ and $G = \sum_{ij} G_{ij}$[1,2]. We observe a typical probability distribution for $P(K_{ij})$ and $P(G_{ij})$ when $0.2 < \nu < 0.7$ (orange data in Fig. 1f), with no noticeable spatial correlations in the largest positive and negative

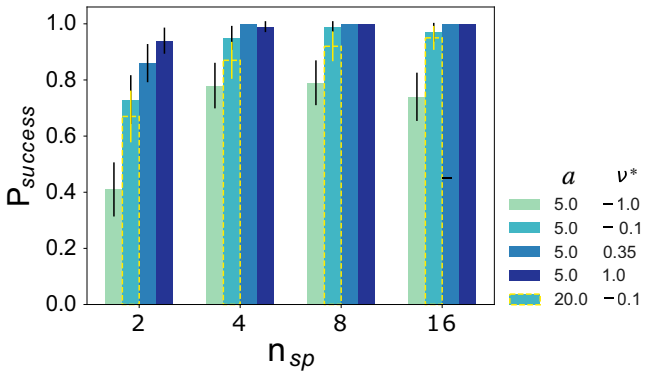

Fig. 2 | The ability to successfully train the Poisson's ratio at a given configuration depends on the number of species, n_{sp}, and the target Poisson's ratio, ν^*. The probability of successful training, $P_{success}$, is close to 1 except for when $n_{sp} < 4$ or when $\nu^* < -0.1$. Results are for a medium-ranged attractive potential ($a = 5$ in Eq. (1)). Using a shorter-ranged potential ($a = 20$, yellow dashed line) decreases the ability to successfully train.

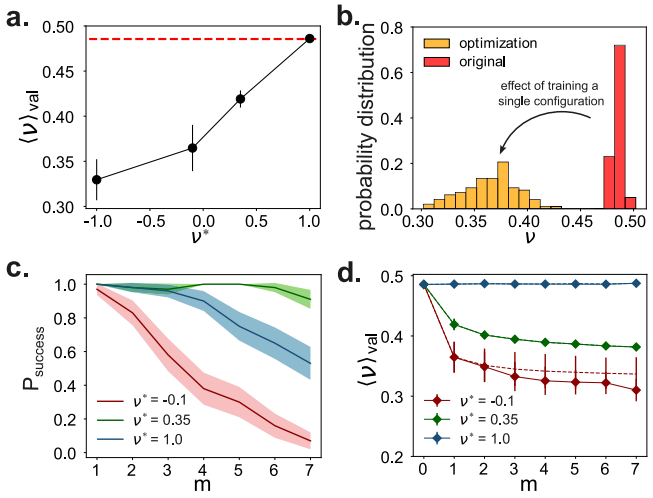

Fig. 3 | Training the Poisson's ratio at a single configuration, ν_b, tunes the ensemble average Poisson's ratio, $\langle\nu\rangle$. a $\langle\nu\rangle$ averaged over 100 randomly chosen configurations using the final, learned parameters obtained from the $n_{sp} = 16$ data from Fig. 2. We call this the "validation data" because, importantly, the Poisson's ratio from the configurations used for training were not included in these averages. The horizontal red line indicates the Poisson's ratio calculated with the initial parameters, so the difference between the points and the horizontal line indicates the amount of change in $\langle\nu\rangle$. **b** The full distribution of the validation data for the initial parameters (red) and the parameters obtained with $\nu^* = -0.1$ (orange). **c** $P_{success}$ when targeting the Poisson's ratio of multiple configurations simultaneously. $P_{success}$ decreases more rapidly for more aggressive ν^*. **d** Training multiple configurations increases the effect on $\langle\nu\rangle$, though with diminishing returns above $m \approx 3$. $m = 0$ means that the initial untrained parameters are used. The solid lines use only parameters after successful optimization, while the dashed lines use parameters after all optimizations.

values. However, for $\nu < 0.2$, the positive and negative tails of $P(K_{ij})$ widen considerably (blue data in Fig. 1f) but equally so they largely offset each other, and we find that the largest values are spatially localized (inset). Apparently, our training procedure has found a way to harness and exploit local non-affine response, which is known to influence the Poisson's ratio[14,15]. We hypothesize that the broad tails in $P(K_{ij})$ enable more dramatic tuning of K, allowing K/G to become small and thus ν to become negative. Conversely, for $\nu > 0.7$, the tails of $P(G_{ij})$ broaden, enabling K/G to become large. See Supplementary Fig. 3, 4 for more comprehensive data. While we do not fully understand the origin of these extreme and localized regions, we note that the range of ν where they occur coincides with our ability to train the ensemble average Poisson's ratio, see below.

Figure 2 shows the probability $P_{success}$ that we are able to successfully train a configuration of $N = 368$ particles for various target Poisson's ratios and different numbers of species. Training is considered successful if the objective $\mathcal{L}$ decreases below $\mathcal{L}_{thresh} = 10^{-6}$ within 10^4 optimization steps. Not surprisingly, $P_{success}$ increases with the number of species since that gives access to additional parameters, and with more conservative targets that are closer to the initial Poisson's ratio of approximately 0.5. We also find that the range of the attractive interaction matters, with shorter interactions making it slightly harder to train.

These results show that we can accurately and precisely control ν_b over a wide range. Interestingly, Fig. 3a–b shows that the parameters learned through this process influence the entire ensemble of configurations even though only one configurations was considered during training. For example, the parameters obtained from training a single configuration to $\nu^* = -1$ results in $\langle\nu\rangle = 0.33 \pm 0.02$, well below the initial untrained value of $\langle\nu\rangle = 0.49$. This is a first indication that we can manipulate ensemble averages.

This effect can be amplified by training on multiple configurations simultaneously. Specifically, we next choose a set of m configurations, with m ranging from 2 to 7, and define the objective $\mathcal{L} = \sum_{b'}(\nu_{b'} - \nu^*)^2$, where here the sum is over the m chosen configurations b'. Figure 3c shows that training becomes harder (less likely to succeed) as m increases, especially for aggressive targets ($\nu^* < 0$). Despite this, however, increasing m does increase the effect on $\langle\nu\rangle$, as shown by Fig. 3d.

Training ensemble-average quantities

We now directly consider the ensemble average Poisson's ratio, $\langle\nu\rangle$, by defining the objective $\mathcal{L} = (\langle\nu\rangle - \nu^*)^2$. Note that this is not the same as the $m \to \infty$ limit of the objective from the previous section because here we are only concerned with the mean of the Poisson's ratio, not its value at every individual configuration. The other conceptual difference is that we can only ever estimate $\langle\nu\rangle$ by averaging over a finite number of (randomly chosen) configurations, meaning that all calculations of $\mathcal{L}$ and $\nabla_\theta\mathcal{L}$ are necessarily stochastic. This is analogous to training in many machine learning models, where data is "batched" and gradients are highly noisy. Importantly, as demonstrated in Supplementary Fig. 1a, average gradients of $\langle\nu\rangle$ are predictive in the same way as in Fig. 1c, allowing the use of stochastic gradient descent optimization.

Figure 4a shows $\mathcal{L}$ during optimization with a target of $\nu^* = 0.35$ using $n_{sp} = 2$ species. At each step, we average over only 8 systems of size $N = 368$, leading to large noise in $\mathcal{L}$. This leads to a systematic overestimation of $\mathcal{L}$ whenever the noise is larger than $\langle\nu\rangle - \nu^*$. Therefore, in order to ascertain how well the training has done, we also perform "validation runs," where we use the final parameters obtained after 5000 optimization steps and calculate $\langle\nu\rangle$ over a fresh set of 2000 configurations. As indicated by the green circle, $\mathcal{L}_{val} \approx 10^{-6}$, meaning that we have trained $\langle\nu\rangle$ to ν^* with an accuracy of 0.001. The purple star in Fig. 4a shows an alternative validation where $\langle\nu\rangle$ is calculated from 100 configurations of $N = 2048$ particles each, demonstrating that even though our results are trained using small systems, the solution nevertheless applies to much larger systems.

Figure 4b shows the 5 parameters at each optimization step (thick lines). The decrease in parameter fluctuations is caused by our variable learning rate (see Supplementary Discussion D), but the lack of systematic trends over the final 500 optimization steps suggests that the optimization has converged. Note that all five parameters change non-trivially over the course of the optimization, and it is unclear if some of the parameters are "more important" than others. Another way to address this is to look at the magnitude of the gradients; Supplementary Fig. 1b hints that the Poisson's ratio may be more sensitive to the design of the smaller particle species.

The thin lines in Fig. 4b show the same data for 29 other training runs that also target $\nu^* = 0.35$, and all achieve similarly small $\mathcal{L}_{val}$. Clearly the final parameters are not unique, meaning that the solution is degenerate, but

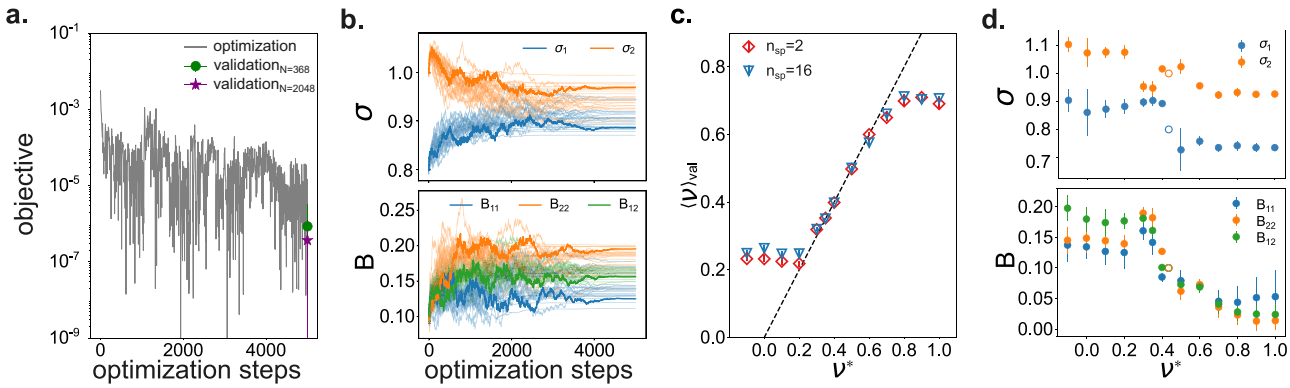

Fig. 4 | Tuning the ensemble. a The objective (gray line) while optimizing $\langle\nu\rangle$ with a target of $\nu^* = 0.35$ using $n_{sp} = 2$ species. Only 8 systems are averaged at each step, making the objective very noisy. Nevertheless, proper validation averages of $\langle\nu\rangle$ using the parameters obtained after 5000 optimization steps (green and purple points) show that the Poisson's ratio has indeed been tuned with very high accuracy. **b** The five parameters evolve in non-trivial ways over the course of optimization. The thick lines correspond to (**a**), while the thin lines show the results of 29 other optimization runs with the same ν^* that differ only in the random configurations that are sampled at each step. All 30 optimization attempts successfully lowered the objective below 10^{-5} and lead to similar, but not identical, trends in the final optimized parameters. **c** The final validated $\langle\nu\rangle$ after optimization as a function of the target ν^*. We can accurately and precisely tune the Poisson's ratio over the range $0.2 < \langle\nu\rangle < 0.7$. Interestingly, increasing the number of species does not improve our ability to train, suggesting that these limits might represent a fundamental barrier for sticky spheres. **d** The final parameters after optimization show clear trends with ν^*, indicating a general design strategy. The open circles indicate the initial parameters and Poisson's ratio before training. The averages are estimated over 10 to 35 successful optimizations.

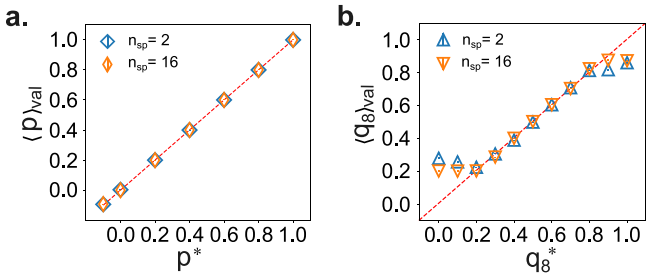

Fig. 5 | Tuning the ensemble-averaged pressure and a structural order parameter q_8. a, **b** The validated $\langle p\rangle$ and $\langle q_8\rangle$ after optimization as functions of the target p^* and q_8^* respectively. We demonstrate accurate and precise control of pressure throughout the test range, while fine-tuning q_8 within the range of $0.2 \leq q_8 \leq 0.9$.

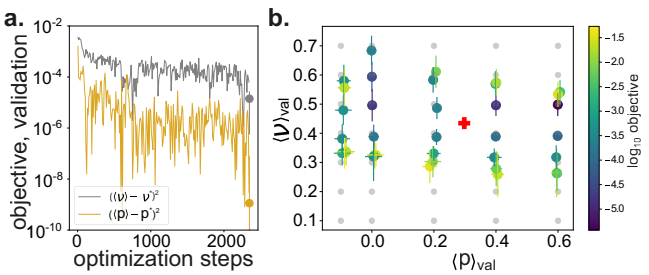

Fig. 6 | Tuning the ensemble-averaged Poisson's ratio and pressure simultaneously. a The two components of the objective during optimization with $\nu^* = 0.5$ and $p^* = 0.4$. **b** By changing ν^* and p^*, we can tune the pair $(\langle\nu\rangle, \langle p\rangle)$ over a non-trivial, 2 dimensional region.

there are still clear trends in the parameters (e.g. small particles bind more strongly to large particles than to other small particles). These trends translate to general design principles with direct experimental implications. To further motivate the experimental relevance of these results, Supplementary Fig. 8 shows that $\langle\nu\rangle$ remains within 0.01 of ν^* even when adding up to 5% polydispersity to the particle diameters and 5% heterogeneity to the binding energies. Note that incorporating polydispersity and/or heterogeneity directly in the training process would enable high accuracy tuning regardless of magnitude.

Figure 4c shows the final validated $\langle\nu\rangle_{val}$ as a function of the target ν^*, and shows that we are able to successfully tune the Poisson's ratio between roughly 0.2 and 0.7. The existence of such bounds in our ability to tune $\langle\nu\rangle$ is expected, especially for low ν^* as creating low- and negative-Poisson's ratio materials is notoriously challenging[16,17]. We expect that more complicated systems, for example with non-spherical particles, could expand these bounds. Notably, this range of roughly 0.2 to 0.7 coincides exactly with the onset of broad and localized bond-level response observed in Fig. 1f. While training the Poisson's ratio of an individual configuration outside of this range is possible, it appears to rely on configuration-specific structural motifs that do not generalize to the full ensemble. This is also consistent with the observation in Fig. 1e that the final parameters vary dramatically from one configuration to another.

Unlike for individual configurations, we find that considering more species does not increase the range over which we can tune the system (blue triangles in Fig. 4c). Intuitively, increasing the number of species should increase tunability, but there are no guarantees of this. It is unclear if improved training methodology could lead to better results with more species or if these bounds are imposed by physical constraints. Finally, Fig. 4d shows the final parameters after 5000 optimization steps for different ν^*, suggesting a general design strategy for tuning $\langle\nu\rangle$ over the range $0.2 < \langle\nu\rangle < 0.7$.

Training multiple ensemble-average quantities simultaneously

So far, we have focused on tuning the Poisson's ratio, but Fig. 5 shows that we can tune other quantities in exactly the same way, specifically the pressure and a structural order parameter q_8, which measures small amounts of local eightfold symmetry in neighbor orientations (see Supplementary Discussion F). Now we ask whether we can tune multiple average properties simultaneously. Figure 6a shows the simultaneous optimization of the Poisson's ratio ($\nu^* = 0.5$) and pressure ($p^* = 0.4$), for which we use a combined objective function $\mathcal{L} = (\langle\nu\rangle - \nu^*)^2 + (\langle p\rangle - p^*)^2$. After 2400 optimization steps, we observe a final validated loss of $\mathcal{L}_{val} = 6 \times 10^{-4}$, demonstrating success.

Not all combinations of $(\langle p\rangle, \langle\nu\rangle)$ can be obtained simultaneously, and Fig. 6b shows the final $(\langle p\rangle_{val}, \langle\nu\rangle_{val})$ obtained from a series of systematic trials with different (p^*, ν^*) (indicated by light gray dots). This data shows a well-defined two-dimensional region of tunability, within which simultaneous optimization can be obtained.

Discussion

We have shown that the properties of athermal disordered solids can be inverse designed using Automatic Differentiation (AD) to connect properties of interest to particle-particle interactions. This can be compared with typical self-assembly paradigms, where one manipulates interactions so that, for example, a desired crystal structure becomes the thermodynamic ground state[18–23]. Instead, our approach targets the statistics of the metastable states. This has the distinct advantage that the assembly protocol (athermal relaxation in our case) is baked into the calculation of these statistics, and so we do not need to worry about kinetic traps or other barriers to assembly.

Importantly, this AD-based approach works "on top of" existing numerical simulations; in principle, gradients can be propagated through any simulation or numerical calculation. Therefore, if a particular level of detail, system size, etc. is necessary for a given application, AD provides a scalable pathway for optimizing behavior. Recent advances in differentiable programming ecosystems, including the Python packages JAX and JAX-MD[24] used in this work, make it possible to incorporate AD into many existing calculations. One notable challenge is optimizing with respect to relative particle concentrations. This is because in finite systems, particle concentrations are rounded to determine a whole number of particles of a given type, and so gradients with respect to concentrations will necessarily vanish. Important other considerations include the predictiveness of gradients, strategies for avoiding large memory loads, and the importance of initial parameter guesses.

We have chosen a relatively simple system to demonstrate how AD can be used to train disordered solids, but this approach is extremely general. First, we have focused on a particular model of sticky spheres with tunable species-level interactions that is motivated by DNA-coated colloids[6,7]. As discussed above, however, it is straightforward to employ much more detailed or realistic energetic models. Next, we consider zero temperature dynamics, which is a common choice when studying, for example, jammed solids or any case where interaction energies are much larger than thermal fluctuations. It also makes a configuration well defined, and allows us to use so-called implicit differentiation to increase the numerical stability of gradients (see Methods). Nevertheless, it should be possible to extend our approach to finite-temperature simulations. Furthermore, our primary focus is on the Poisson's ratio, ν, because, as a unitless measure of elasticity, it cannot be adjusted by trivially changing energy scales. It can also be calculated exactly and efficiently at zero temperature. Finally, our focus on two-dimensional systems is purely for numerical convenience and, as shown in the Supplementary Discussion H, we can train three-dimensional systems just as well.

For this system, we have shown that the ensemble average Poisson's ratio, $\langle\nu\rangle$, can be tuned to anywhere in the range $0.2 < \langle\nu\rangle < 0.7$. Furthermore, we can tune the pressure and local structural order in a similar way. While other systems exhibit a variable Poisson's ratio – for example, jammed packings of soft spheres display $\langle\nu\rangle \to 1$ near the unjamming transition[9,25], our work demonstrates a new level of systematic and robust control. In addition, we have shown that directly connecting objectives to parameters through the gradient $\nabla_\theta \mathcal{L}$ enables the targeted design of multiple properties, which is otherwise challenging. Together, this presents a scalable approach to inverse design and reveals a direct strategy for the targeted manipulation of disordered solids.

We have also shown that the Poisson's ratio of a single configuration can be tuned with a high success rate anywhere in the range $-1 \le \nu \le 1$ (the theoretical bounds for two-dimensional isotropic systems). This is analogous to the bond-level tuning in refs. 1,4 except it is at the level of species-species interactions and maintains the constraint of force-balance on every particle. Figure 1f also hints at an unexpected localized mechanism for creating auxetic materials. While these individual configurations are likely difficult to obtain experimentally due to the overwhelming number of competing structures, this result is important because it demonstrates extreme tunability at the level of individual configurations. Thus, the lower bound of $\langle\nu\rangle$ is inherently a collective effect: at $\langle\nu\rangle \approx 0.2$, any particular ν_b can still be lowered, but doing so necessarily increase the Poisson's ratio of other configurations so that $\langle\nu\rangle$ is unchanged.

This paper has only scratched the surface of how disordered solids can be designed with Automatic Differentiation, and there are many exciting ways that this work could be extended. First, our results are obtained with spherically symmetric particles with very simple, Lennard-Jones-like interactions. However, there are numerous experimental systems, from nonspherical colloids[26] to colloidal-scale DNA-origami[27] to de novo proteins[28,29], that combine controllable interactions with nontrivial shape, and that could be used to increase designability. Next, AD could be used to optimized over nontrivial preparation protocols, including temperature ramps or oscillatory shear. Finally, there are numerous other material properties that could be (simultaneously) targeted, including thermal properties, nonlinear stress-strain behavior[30,31], density of states, and allosteric responses[3]. Critical to the future success of such extensions will be the continued integration of differentiable programming ecosystems with classical molecular dynamics through packages like JAX-MD[24].

Methods

Model

We consider a system composed of N particles divided evenly into n_{sp} species. Particles interact via a "harmonic-Morse" pairwise potential given by

$$V(r) = \begin{cases} \frac{k}{2}(r-\bar{\sigma})^2 - B, & r < \bar{\sigma} \\ B\left(e^{-2a(r-\bar{\sigma})} - 2e^{-a(r-\bar{\sigma})}\right), & r \ge \bar{\sigma} \end{cases} \quad (1)$$

where r is the center-center distance between two particles, $\bar{\sigma}$ is the mean of their diameters, k characterizes the short-ranged repulsions, and B determines the strength of the medium-range attractions, whose extent is proportional to $1/a$. Unless otherwise stated, we use $k = 5.0$ and $a = 5.0$ for all pairs of particles. However, B and $\bar{\sigma}$ depend on the species type of the particles in question. Specifically, we independently vary the attractive strength $B_{\alpha\beta}$ for every pair of species α and β (with $B_{\alpha\beta} = B_{\beta\alpha}$), as well as the particle diameter σ_α (so that $\bar{\sigma}_{\alpha\beta} = (\sigma_\alpha + \sigma_\beta)/2$). We then use the XPLOR smoothing function to truncate $V(r)$ at a distance $r_{cut} = \bar{\sigma} + 9.9/a$. This model is commonly used to describe, for example, DNA-coated colloids where diameters and binding affinities can be manipulated at the species level[6,7]. Note that we use a harmonic repulsive force to ensure numerical stability at a random set of initial position, and that the strength k of this repulsion is decoupled from the attractive forces to allow $B_{\alpha\beta} \to 0$.

For a given set of parameters $\theta = \{\sigma_\alpha\} \cup \{B_{\alpha\beta}\}$, we obtain stable athermal structures by using the FIRE algorithm[32,33] to minimize the total potential energy $E = \sum_{<ij>} V(r_{ij})$ starting from a random set of initial positions. We always consider an even number of species, and set the initial diameters for half the species to be 0.8 times that of the other half, with an overall number density of 1.6. The initial binding strengths are all set to $B_{\alpha\beta} = 0.1$.

Calculating the Poisson's ratio

The Poisson's ratio is expressed in terms of the elements of the elastic modulus tensor C_{ijkl}, which describe the second order terms of the Taylor expansion of the energy of a system with respect to boundary deformations:

$$\frac{\Delta U}{V^0} = \sigma^0_{ij}\epsilon_{ji} + \frac{1}{2}C_{ijkl}\epsilon_{ij}\epsilon_{kl} + O(\epsilon^3), \quad (2)$$

where σ^0 is the stress tensor describing residual stresses in the initial state, ϵ is the strain tensor, describing boundary deformations, V^0 is the volume of the initial state, and ΔU is the change in energy of the system.

If we choose a particular strain of magnitude γ, i.e. $\epsilon = \tilde{\epsilon}\gamma$ for some $\tilde{\epsilon}$, Eq. (2) becomes

$$\frac{\Delta U}{V^0} = \tilde{\sigma}^0\gamma + \frac{1}{2V^0}\left(\frac{\partial^2 U}{\partial\gamma^2} - \Xi^T(H^0)^{-1}\Xi\right)\gamma^2 + O(\gamma^3), \quad (3)$$

where $\tilde{\sigma}^0 = \sigma_{ij}\tilde{\epsilon}_{ji}$, $\Xi = -\frac{\partial \nabla U}{\partial \gamma}$ is the derivative of the force on every particle with respect to γ, and $H^0 = \nabla^2 U$ is the Hessian matrix of the energy with respect to particle positions. Equation (3) is derived elsewhere[34,35], and can be used to back out all elements of C_{ijkl} through appropriate choices of $\tilde{\epsilon}_{ij}$. Our implementation is part of the core JAXMD library, and more detailed documentation and explanations can be found there.

Optimization for individual configurations

The original configurations are generated by quenching random packings to local energy minima using the FIRE algorithm. Once an initial energy-minimized state is reached, the Poisson's ratio ν is calculated explicitly using linear response[36], and combined with the target Poisson's ratio ν^* to construct the objective function $\mathcal{L} = (\nu - \nu^*)^2$.

Our model, the energy minimization, and the calculation of ν and $\mathcal{L}$ are performed entirely using the library JAXMD[24] and the Automatic Differentiation library JAX[37] on which it is built. As a result, the entire calculation is end-to-end differentiable, allowing us to accurately and efficiently calculate $\nabla_\theta \mathcal{L}$. Implicit differentiation, implemented using the JAXOPT library[38], is used to propagate derivatives through the energy minimization as well as a key step in the linear response calculation, which ensures gradient accuracy and mitigates any extensive memory overhead. We then employ the RMSProp algorithm (implemented as part of JAX) to iteratively update the parameters θ based on the gradients. To improve the convergence speed, we dynamically adjust the learning rate using a meta-learning program that adaptively generates per-step hyperparameters. At each iteration, we use the final configuration of the previous step for the initial positions of the next energy minimization process, ensuring that we consistently track a given energy minimum. Occasionally, the energy minimum transforms into a saddle point and the system undergoes a rearrangement. This results in a small spike in $\mathcal{L}$, e.g. see Fig. 1d at around 1300 steps.

Optimization for ensemble-averaged quantities

The primary difference in optimization between individual and ensemble systems lies in the input configurations. In the optimization process for ensemble-averaged quantities, we generate 8 new configurations (by minimizing from random initial positions) at every step. Given our relatively small system size, this is well below the number of configurations needed to obtain accurate ensemble averages, meaning our estimates of $\mathcal{L}$ and $\nabla_\theta \mathcal{L}$ are highly noisy. However, stochastic optimization algorithms like RMSProp are well-suited to very noisy gradients and this does not prevent convergence.

Data availability

The authors declare that the data supporting the findings of this study are available within the paper and its Supplementary Information files.

Code availability

Computer code is available at https://github.com/MJ-Zu/design-disordered.

Received: 12 April 2024; Accepted: 18 July 2024;
Published online: 01 August 2024

References

1. Goodrich, C. P., Liu, A. J. & Nagel, S. R. The principle of independent bond-level response: Tuning by pruning to exploit disorder for global behavior. *Phys. Rev. Lett.* **114**, 225501 (2015).
2. Hexner, D., Liu, A. J. & Nagel, S. R. Role of local response in manipulating the elastic properties of disordered solids by bond removal. *Soft Matter* **14**, 312–318 (2017).
3. Rocks, J. W. et al. Designing allostery-inspired response in mechanical networks. *Proc. Natl. Acad. Sci. USA* **114**, 2520–2525 (2017).
4. Reid, D. R. et al. Auxetic metamaterials from disordered networks. *Proc. Natl. Acad. Sci. USA* **115**, E1384–E1390 (2018).
5. Rocks, J. W., Ronellenfitsch, H., Liu, A. J., Nagel, S. R. & Katifori, E. Limits of multifunctionality in tunable networks. *Proc. Natl. Acad. Sci. USA* **116**, 2506–25115 (2019).
6. Rogers, W. B. & Crocker, J. C. Direct measurements of DNA-mediated colloidal interactions and their quantitative modeling. *Proc. Natl. Acad. Sci. USA* **108**, 15687–15692 (2011).
7. Cui, F., Marbach, S., Zheng, J. A., Holmes-Cerfon, M. & Pine, D. J. Comprehensive view of microscopic interactions between DNA-coated colloids. *Nat. Commun.* **13**, 2304 (2022).
8. O'Hern, C. S., Silbert, L. E., Liu, A. J. & Nagel, S. R. Jamming at zero temperature and zero applied stress: The epitome of disorder. *Phys. Rev. E* **68**, 011306 (2003).
9. Liu, A. J. & Nagel, S. R. The Jamming transition and the marginally Jammed solid. *Annu. Rev. Condens.* **1**, 347–369 (2010).
10. Baydin, A.G., Pearlmutter, B. A., Radul, A. A. & Siskind, J. M. Automatic differentiation in machine learning: a survey. *J. Mach. Learn. Res.* **18**, 1–43 (2018).
11. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
12. Wengert, R. E. A simple automatic derivative evaluation program. *Commun. ACM* **7**, 463–464 (1964).
13. Ruder, S. An overview of gradient descent optimization algorithms. *arXiv* 1609.04747 (2017).
14. Zaccone, A. & Terentjev, E. M. Short-range correlations control the G/K and Poisson ratios of amorphous solids and metallic glasses. *J. Appl. Phys.* **3**, 115 (2014).
15. Schlegel, M., Brujic, J., Terentjev, E. M. & Zaccone, A. Local structure controls the nonaffine shear and bulk moduli of disordered solids. *Sci. Rep.* **6**, 18724 (2016).
16. Mott, P. H. & Roland, C. M. Limits to Poisson's ratio in isotropic materials. *Phys. Rev. B* **80**, 132104 (2009).
17. Rechtsman, M. C., Stillinger, F. H. & Torquato, S. Negative Poisson's Ratio Materials via Isotropic Interactions. *Phys. Rev. Lett.* **101**, 085501 (2008).
18. Rechtsman, M. C., Stillinger, F. H. & Torquato, S. Optimized Interactions for Targeted Self-Assembly: Application to a Honeycomb Lattice. *Phys. Rev. Lett.* **95**, 228301 (2005).
19. Engel, M., Damasceno, P. F., Phillips, C. L. & Glotzer, S. C. Computational self-assembly of a one-component icosahedral quasicrystal. *Nat. Mater.* **14**, 109–116 (2015).
20. Miskin, M. Z., Khaira, G., de Pablo, J. J. & Jaeger, H. M. Turning statistical physics models into materials design engines. *Proc. Natl. Acad. Sci. USA* **113**, 34–39 (2016).
21. Chen, D., Zhang, G. & Torquato, S. Inverse design of colloidal crystals via optimized patchy interactions. *J. Phys. Chem. B* **122**, 8462–8468 (2018).
22. Kumar, R., Coli, G. M., Dijkstra, M. & Sastry, S. Inverse design of charged colloidal particle interactions for self assembly into specified crystal structures. *J. Chem. Phys.* **151**, 084109 (2019).
23. Romano, F., Russo, J., Kroc, L. & Šulc, P. Designing patchy interactions to self-assemble arbitrary structures. *Phys. Rev. Lett.* **125**, 118003 (2020).
24. Schoenholz, S. S. & Cubuk, E. D. JAX M.D. A framework for differentiable physics. *Advances in Neural Information Processing Systems*, vol. 33 (Curran Associates, Inc., 2020).
25. Hecke, M. V. Jamming of soft particles: geometry, mechanics, scaling and isostaticity. *J. Phys. Condens. Matter* **22**, 033101 (2010).
26. He, M. et al. Colloidal diamond. *Nature* **585**, 524–529 (2020).
27. Sigl, C. et al. Programmable icosahedral shell system for virus trapping. *Nat. Mater.* **20**, 1281–1289 (2021).

28. Huang, P. S., Boyken, S. E. & Baker, D. The coming of age of de novo protein design. *Nature* **537**, 320–327 (2016).
29. Boyken, S. E. et al. De novo design of protein homo-oligomers with modular hydrogen-bond network-mediated specificity. *Science* **352**, 680–687 (2016).
30. Wang, Y., Zhang, X., Li, Z., Gao, H. & Li, X. Achieving the theoretical limit of strength in shell-based carbon nanolattices. *Proc. Natl. Acad. Sci. USA* **119**, e2119536119 (2022).
31. Cheng, H. et al. Mechanical metamaterials made of freestanding quasi-BCC nanolattices of gold and copper with ultra-high energy absorption capacity. *Nat. Commun.* **14**, 1243 (2023).
32. Bitzek, E., Koskinen, P., Gähler, F., Moseler, M. & Gumbsch, P. Structural relaxation made simple. *Phys. Rev. Lett.* **97**, 170201 (2006).
33. Guénolé, J. et al. Assessment and optimization of the fast inertial relaxation engine (fire) for energy minimization in atomistic simulations and its implementation in lammps. *Comput. Mater. Sci.* **175**, 109584 (2020).
34. Lemaˆıtre, A. & Maloney,C. Sum rules for the quasi-static and visco-elastic response of disordered solids at zero temperature. *J. Stat. Phys.* **123**, 415–453 (2006).
35. Maloney, C. E., & Lemaˆıtre, A. Amorphous systems in athermal, quasistatic shear. *Phys. Rev. E* **74**, 016118 (2006).
36. Goodrich, C. P. et al. Jamming in finite systems: Stability, anisotropy, fluctuations, and scaling. *Phys. Rev. E* **90**, 022138 (2014).
37. Bradbury, J. et al. JAX: Composable transformations of Python +NumPy programs. http://github.com/google/jax (2018).
38. Blondel, M. et al. Efficient and modular implicit differentiation. *Advances in Neural Information Processing Systems* (eds Koyejo, S. et al.) vol. 35, 5230–5242 (Curran Associates, Inc., 2022).

Acknowledgements

We thank M. Lechner for helpful discussions.

Author contributions

M.Z. performed the numerical simulations. M.Z. and C.P.G. designed and developed the study, analyzed the data, and wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s43246-024-00583-4.

Correspondence and requests for materials should be addressed to Carl P. Goodrich.

Peer review information *Communications Materials* thanks Natsuhiko Yoshinaga and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editor: John Plummer.