

Designing athermal disordered solids with automatic differentiation

Mengjie ZU* and Carl GOODRICH†
Institute of Science and Technology Austria
 (Dated: April 24, 2024)

The ability to control forces between sub-micron-scale building blocks offers considerable potential for designing new materials through self-assembly. A typical paradigm is to first identify a particular (crystal) structure that has some desired property, and then design building-block interactions so that this structure assembles spontaneously. While significant theoretical and experimental progress has been made in assembling complicated structures in a variety of systems, this two-step paradigm fundamentally fails for structurally disordered solids, which lack a well-defined structure to use as a target. Here we show that disordered solids can still be treated from an inverse self-assembly perspective by targeting material properties directly. Using the Poisson’s ratio, ν , as a primary example, we show how differentiable programming connects experimentally relevant interaction parameters with emergent behavior, allowing us to iteratively “train” the system until we find the set of interactions that leads to the Poisson’s ratio we desire. Beyond the Poisson’s ratio, we also tune the pressure and a measure of local 8-fold structural order, as well as multiple properties simultaneously, demonstrating the potential for nontrivial design in disordered solids. This approach is highly robust, transferable, and scalable, can handle a wide variety of model systems, properties of interest, and preparation dynamics, and can optimize over 100s or even 1000s of parameters. This result connects the fields of disordered solids and inverse self-assembly, indicating that many of the tools and ideas that have been developed to understand the assembly of crystals can also be used to control the properties of disordered solids.

keywords: self-assembly | disordered solids | auxetic materials | simulation | automatic differentiation

I. INTRODUCTION

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 iso-static 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 $\sim 1\%$ 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_{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 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 $n_{\text{sp}} + n_{\text{sp}}(n_{\text{sp}} + 1)/2$ 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

* Institute of Science and Technology Austria;
 mengjie.zu@ist.ac.at

† carl.goodrich@ist.ac.at




FIG. 1. Tuning the Poisson's ratio of a single configuration. **a.** 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. **b.** 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. **c.** 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. **d.** The evolution of the parameters for 35 successful optimizations. The thick lines correspond to the example shown in **c.** The number of optimization steps is variable because the process terminates after convergence. **e.** 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 SI 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.

$n_{\text{sp}} + n_{\text{sp}}(n_{\text{sp}} + 1)/2$ 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 algorithms [13] for using this information to minimize $\mathcal{L}$, thus tuning the property of interest. 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 $\nu^* \rightarrow -1$. 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.

II. RESULTS

A. 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 1b 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 $n_{\text{sp}} = 2$ 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 $\nu_b \approx 0.453$. We train the system using standard gradient-descent-based algorithms (see Methods) to iteratively 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 1c 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 1d shows how the parameters evolve during training for 35 successfully trained configurations. Clearly, the final parameters depend strongly on the particular configuration; while the systems remain bidisperse and thus disordered, there is no obvious trend in the relative particle diameters nor in the three binding energies.

Interestingly, 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. 1e), with no noticeable spatial correlations in the largest positive and negative values. However, for $\nu < 0.2$, the positive and negative tails of $P(K_{ij})$ widen considerably (blue data in Fig. 1e) but equally so they largely offset each other, and we find that the largest values are spatially localized (inset). 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 SI 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 $\mathcal{L}_{\text{thresh}} = 10^{-6}$ 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.




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 $n_{\text{sp}} < 4$ 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.

These results show that we can accurately and precisely control ν_b over a wide range. Interestingly, Fig. 3a-b show that the parameters learned through this process influence the entire ensemble of configurations even though they were never considered in the training process. For example, the parameters obtained from training a single system 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.

B. 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 $m \rightarrow \infty$ 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 Fig.S1, average gradients of $\langle \nu \rangle$ are predictive in the same way as in Fig. 1b, allowing the use




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 $n_{\text{sp}} = 16$ 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.

of stochastic gradient descent optimization.

Figure 4a shows $\mathcal{L}$ during optimization with a target of $\nu^* = 0.35$ using $n_{\text{sp}} = 2$ 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, $\mathcal{L}_{\text{val}} \approx 10^{-6}$, 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 SI), but the lack of systematic trends over the final 500 optimization

steps suggests that the optimization has converged. The thin lines show the same data for 29 other training runs that also target $\nu^* = 0.35$ and all achieve similarly small $\mathcal{L}_{\text{val}}$. Clearly the final parameters are not unique, meaning that the solution is degenerate, but there are still clear trends in the parameters (e.g. small particles bind more strongly to large particles than to other small particles).

Figure 4c shows the final validated $\langle \nu \rangle_{\text{val}}$ as a function of the target ν^* , showing 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 [14, 15]. 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. 1e. 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. 1d that the final parameters vary dramatically from one configuration to another.

Interestingly, we find that considering more species does not increase the range over which we can tune the system (blue triangles in Fig. 4c). In principle, 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 $0.2 < \langle \nu \rangle < 0.7$.

C. 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 8-fold symmetry in neighbor orientations (see SI). 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 ($p^* = 0.4$), which results in a final validated loss of $\mathcal{L}_{\text{val}} = 6 \times 10^{-4}$, demonstrating success. To be precise, we use a combined objective function $\mathcal{L} = (\langle \nu \rangle - \nu^*)^2 + (\langle p \rangle - p^*)^2$, where the two components are plotted separately.

Not all values of $(\langle p \rangle, \langle \nu \rangle)$ can be obtained, and Fig. 6b shows the final $(\langle p \rangle_{\text{val}}, \langle \nu \rangle_{\text{val}})$ obtained from a series of systematic trials with different (p^*, ν^*) (indicated by light grey dots). This data shows a well-defined two-




FIG. 4. Tuning the ensemble. **a.** The objective (grey 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 states 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.




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.

dimensional region of tunability, within which simultaneous optimization can be obtained.

III. DISCUSSION

We have shown that the properties of 2d athermal disordered solids can be inverse designed using Automatic Differentiation to connect properties of interest to particle-particle interactions. We have focused primarily on the Poisson's ratio, ν , because as a unitless measure of elasticity it cannot be adjusted by trivially changing energy scales. Using a simple model of sticky spheres with tunable species-level interactions that is motivated by DNA-coated colloids [6, 7], we have shown that the Poisson's ratio of a single configuration can be tuned with a high success rate anywhere in the range $-1 \leq \nu \leq 1$ (the theoretical bounds for two-dimensional isotropic systems), and the ensemble average $\langle \nu \rangle$ can be tuned in the range $0.2 < \langle \nu \rangle < 0.7$. Furthermore, we can tune the pressure and local structural order in a similar way, and even design for multiple properties simultaneously.

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 [16–21]. 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.

While other systems exhibit a variable Poisson's ratio

– for example, jammed packings of soft spheres display $\langle \nu \rangle \rightarrow 1$ near the unjamming transition [22] [9, 23], 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.

Unlike for ensemble averages, we are able to precisely tune the Poisson’s ratio of individual configurations with high success even in the auxetic or fully incompressible regimes. 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 1e 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, it is nevertheless telling how designable individual configurations are because it implies that 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.

One of the more surprising results is that increasing the number of species, and thus the number of parameters, does not increase the range over which we can tune $\langle \nu \rangle$. This implies a fundamental constraint that is perhaps related to the general rarity of low and negative Poisson’s ratio materials and leaves open the challenge of using our approach in more sophisticated models, with other types of design parameters, to obtain $\langle \nu \rangle < 0$.

This paper has only scratched the surface of how disordered solids can be designed with Automatic Differentiation. Exciting extensions of this include 1) considering more complicated building blocks, for example with non-trivial shapes or anisotropic potentials, 2) considering different and more complicated preparation protocols, and 3) focusing on a wider range of material properties (e.g. nonlinear stress-strain behavior [24, 25], density of states, and allosteric responses [3]). Recent advances in differentiable programming ecosystems, including the Python packages JAX and JAX-MD used in this work, make it possible to incorporate Automatic Differentiation (AD) into many existing calculations. Important considerations, which are discussed in the context of the present work in the Methods section, include the predictiveness of gradients, strategies for avoiding large memory loads, and the importance of initial parameter guesses.

-
- [1] C. P. Goodrich, A. J. Liu, and S. R. Nagel, The Principle of Independent Bond-Level Response: Tuning by Pruning to Exploit Disorder for Global Behavior, *Physical Review Letters* **114**, 225501 (2015).
 - [2] D. Hexner, A. J. Liu, and S. R. Nagel, Role of local response in manipulating the elastic properties of disordered solids by bond removal, *Soft Matter* **14**, 312 (2017).
 - [3] J. W. Rocks, N. Pashine, I. Bischofberger, C. P. Goodrich, A. J. Liu, and S. R. Nagel, Designing allostery-inspired response in mechanical networks, *Proceedings of the National Academy of Sciences* **114**, 2520 (2017), 1607.08562.
 - [4] D. R. Reid, N. Pashine, J. M. Wozniak, H. M. Jaeger, A. J. Liu, S. R. Nagel, and J. J. d. Pablo, Auxetic metamaterials from disordered networks, *Proceedings of the National Academy of Sciences* **115**, E1384 (2018), 1710.02493.
 - [5] J. W. Rocks, H. Ronellenfitsch, A. J. Liu, S. R. Nagel, and E. Katifori, Limits of multifunctionality in tunable networks, *Proceedings of the National Academy of Sciences* **116**, 2506 (2019), 1805.00504.
 - [6] W. B. Rogers and J. C. Crocker, Direct measurements of DNA-mediated colloidal interactions and their quantitative modeling, *Proceedings of the National Academy of Sciences* **108**, 15687 (2011).
 - [7] F. Cui, S. Marbach, J. A. Zheng, M. Holmes-Cerfon, and D. J. Pine, Comprehensive view of microscopic interactions between DNA-coated colloids, *Nature Communications* **13**, 2304 (2022), 2111.06468.
 - [8] C. S. O’Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, Jamming at zero temperature and zero applied stress: The epitome of disorder, *Physical Review E* **68**, 011306 (2003), cond-mat/0304421.
 - [9] A. J. Liu and S. R. Nagel, The Jamming Transition and the Marginally Jammed Solid, *Annual Review of Condensed Matter Physics* **1**, 347 (2010).
 - [10] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: a survey (2018), arXiv:1502.05767 [cs.SC].
 - [11] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature* **323**, 533 (1986).
 - [12] R. E. Wengert, A simple automatic derivative evaluation program, *Commun. ACM* **7**, 463–464 (1964).
 - [13] S. Ruder, An overview of gradient descent optimization algorithms (2017), arXiv:1609.04747 [cs.LG].
 - [14] P. H. Mott and C. M. Roland, Limits to Poisson’s ratio in isotropic materials, *Physical Review B* **80**, 132104 (2009), 0909.4697.
 - [15] M. C. Rechtsman, F. H. Stillinger, and S. Torquato, Negative poisson’s ratio materials via isotropic interactions, *Phys. Rev. Lett.* **101**, 085501 (2008).
 - [16] M. C. Rechtsman, F. H. Stillinger, and S. Torquato, Optimized Interactions for Targeted Self-Assembly: Application to a Honeycomb Lattice, *Physical Review Letters* **95**, 228301 (2005), cond-mat/0508495.
 - [17] M. Engel, P. F. Damasceno, C. L. Phillips, and S. C. Glotzer, Computational self-assembly of a one-component icosahedral quasicrystal, *Nature Materials* **14**, 109 (2015).
 - [18] M. Z. Miskin, G. Khaira, J. J. d. Pablo, and H. M. Jaeger, Turning statistical physics models into materials design

- engines, *Proceedings of the National Academy of Sciences* **113**, 34 (2016), 1510.05580.
- [19] D. Chen, G. Zhang, and S. Torquato, Inverse Design of Colloidal Crystals via Optimized Patchy Interactions, *The Journal of Physical Chemistry B* **122**, 8462 (2018).
 - [20] R. Kumar, G. M. Coli, M. Dijkstra, and S. Sastry, Inverse design of charged colloidal particle interactions for self assembly into specified crystal structures, *The Journal of Chemical Physics* **151**, 084109 (2019), 1905.11061.
 - [21] F. Romano, J. Russo, L. Kroc, and P. Šulc, Designing Patchy Interactions to Self-Assemble Arbitrary Structures, *Physical Review Letters* **125**, 118003 (2020), 2007.15873.
 - [22] Note that while our model reduces to the soft-sphere model by setting $B_{\alpha\beta} = 0$ and $D_1 = 1.4D_2$, the $\langle\nu\rangle \rightarrow 1$ result is only obtained for finite systems in a fixed-pressure ensemble, and is therefore inaccessible in our fixed-volume ensemble.
 - [23] M. v. Hecke, Jamming of soft particles: geometry, mechanics, scaling and isostaticity, *Journal of Physics: Condensed Matter* **22**, 033101 (2010), 0911.1384.
 - [24] Y. Wang, X. Zhang, Z. Li, H. Gao, and X. Li, Achieving the theoretical limit of strength in shell-based carbon nanolattices, *Proceedings of the National Academy of Sciences* **119**, e2119536119 (2022).
 - [25] H. Cheng, X. Zhu, X. Cheng, P. Cai, J. Liu, H. Yao, L. Zhang, and J. Duan, Mechanical metamaterials made of freestanding quasi-BCC nanolattices of gold and copper with ultra-high energy absorption capacity, *Nature Communications* **14**, 1243 (2023).
 - [26] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, Structural relaxation made simple, *Phys. Rev. Lett.* **97**, 170201 (2006).
 - [27] J. Guérolé, W. G. Nöhring, A. Vaid, F. Houllé, Z. Xie, A. Prakash, and E. Bitzek, Assessment and optimization of the fast inertial relaxation engine (fire) for energy minimization in atomistic simulations and its implementation in lammmps, *Computational Materials Science* **175**, 109584 (2020), 1908.02038.
 - [28] C. P. Goodrich, S. Dagois-Bohy, B. P. Tighe, M. van Hecke, A. J. Liu, and S. R. Nagel, Jamming in finite systems: Stability, anisotropy, fluctuations, and scaling, *Phys. Rev. E* **90**, 022138 (2014).
 - [29] S. S. Schoenholz and E. D. Cubuk, Jax m.d. a framework for differentiable physics, in *Advances in Neural Information Processing Systems*, Vol. 33 (Curran Associates, Inc., 2020).
 - [30] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: composable transformations of Python+NumPy programs (2018).
 - [31] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, Efficient and modular implicit differentiation, *arXiv preprint arXiv:2105.15183* (2021).

IV. MODEL AND 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

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

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 $a = 5$ 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 (provided $B_{\alpha\beta} = B_{\beta\alpha}$), 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 $r_{\text{cut}} = \bar{\sigma} + 9.9/a$. 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 $B_{\alpha\beta} \rightarrow 0$.

For a given set of parameters $\theta = \{\sigma_\alpha\} \cup \{B_{\alpha\beta}\}$, we obtain stable athermal structures by using the FIRE algorithm [26, 27] to minimize the total potential energy $E = \sum_{\langle ij \rangle} V(r_{ij})$ 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 $B_{\alpha\beta} = 0.1$.

Optimization for individual configurations. The original configurations are generated by quenching random packings to local energy minima using FIRE algorithm. Once an initial energy-minimized state with original setting are reached, the Poisson’s ratio ν is calculated explicitly using linear response [28], 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 [29] and the Automatic Differentiation library JAX [30] 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 [31], 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 gradi-

ents. 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 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. 1c 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.

Supplementary information for “Designing athermal disordered solids with automatic differentiation”

Mengjie ZU* and Carl GOODRICH†
Institute of Science and Technology Austria

A. Gradient validation for the ensemble

To demonstrate the predictiveness of the average gradient of the ensemble $\langle \nabla_{\theta} \nu \rangle$, we measure the average change of the Poisson’s ratio over 1000 configurations. The configurations are set up in the same way as in Fig 1b, but with a large number of particles ($N = 10000$) to improve the statistics. In Fig. S1, the dots show the averaged change in the Poisson’s ratio $\langle |\Delta \nu| \rangle = \langle |\nu(\Delta \theta) - \nu(\Delta \theta = 0)| \rangle$ measured at $\Delta \theta$, where $\Delta \theta$ is the change in parameters in an arbitrarily chosen direction, and the red line shows the prediction $|\langle \nabla_{\theta} \nu \rangle \cdot \Delta \theta|$. Due to the finite system size, the distribution of Poisson’s ratio has a width with an order of magnitude of ~ 4 . Therefore, only results above 10^{-3} were displayed. Despite the fluctuations, the averaged calculated gradient was still able to predict the change in the Poisson’s ratio for the ensemble.

B. Number of optimization steps

When training individual configurations, we observed that the probability of success $P_{success}$ improves as the number of optimization steps increases. As illustrated in Fig. S2, $P_{success}$ increases by 10% when employing twice the maximum optimization steps. This means that the $P_{success}$ shown in Fig. 2 in the main text is only a lower bound on the actual probability of success. It is unclear if $P_{success} \rightarrow 1$ if we were to wait long enough.




Figure S1. The averaged change in the Poisson’s ratio $\langle |\Delta \nu| \rangle$ as a function of parameter variations $\Delta \theta$ over 1000 configurations. The solid red line represents the averaged calculated gradient projected onto an arbitrarily chosen direction, with the shaded region indicating the estimated error.

* Institute of Science and Technology Austria; mengjie.zu@ist.ac.at

† carl.goodrich@ist.ac.at




Figure S2. The probability of success, denoted as $P_{success}$, for training individual configurations with varying number of species n_{sp} and a target of $\nu^* = -1$, is depicted for different maximum optimization steps. The coloured regions represent the 95% confidence interval, indicating the estimated probability of success.

C. Bond-level elastic response

When a mechanically stable system is distorted by a symmetric strain tensor $\epsilon_{\alpha\beta}$, the change in energy can be written as

$$\Delta E/V = \sigma_{\alpha\beta}^0 \epsilon_{\alpha\beta} + \frac{1}{2} \epsilon_{\alpha\beta} c_{\alpha\beta\gamma\delta} \epsilon_{\gamma\delta} + O(\epsilon^3) \quad (1)$$

where V is the volume of the system, $\sigma_{\alpha\beta}^0$ is the residual stress tensor and $c_{\alpha\beta\gamma\delta}$ is the elastic modulus tensor.

In general, there are 6 independent components of the elastic modulus tensor for 2d systems, which are calculated from the Hessian of the energy [1].

Because ΔE can be written as a sum over pairs of interactions, so too can the elastic modulus tensor,

$$c_{\alpha\beta\gamma\delta} = \sum_{ij} c_{ij, \alpha\beta\gamma\delta} \quad (2)$$

and the change in energy of a pair of interaction is $\Delta E_{ij} = \frac{1}{2} \epsilon_{\alpha\beta} c_{ij, \alpha\beta\gamma\delta} \epsilon_{\gamma\delta}$. With $c_{ij, \alpha\beta\gamma\delta}$ describing the “bond-level” elastic response for a pair of interaction ij , it allows the calculation of the bond-level bulk modulus K_{ij} and shear modulus G_{ij} .

Figure S3 shows the probability distributions $P(K_{ij})$ and $P(G_{ij})$ at a range of Poisson’s ratios ν . As discussed in the main text, there is a noticeable change in the positive and negative tails of the distributions, especially $P(K_{ij})$, at different ν . The emergence of these broad tails corresponds to the emergence of spatial localization in those largest values, as shown in Fig. S4, meaning that the tails are spatially localized. This also roughly corresponds to when ensemble training of the Poisson’s ratio becomes difficult.

D. Ensemble optimization for desired Poisson’s ratio ν^*

In Fig. S5 and S6, we show the optimization processes for the ensemble with two extreme targets, $\nu^* = -0.1$ and 1.0. The optimizations were performed for systems of $N = 368$ particles evenly divided into 2 species. We use a variable learning rate lr that changes for each step s according to

$$lr(s) = \frac{lr_{\max} - lr_{\min}}{2} \left(\cos \left(\frac{\pi s}{s_{\max}} \right) + 1 \right) + lr_{\min}, \quad (3)$$

where $lr_{\max}$, $lr_{\min}$, and $s_{\max}$ are manually adjusted for different training properties. For the results shown in Fig. 4 in the main text, we use $lr_{\max} = 0.002$, $lr_{\min} = 10^{-5}$, and $s_{\max} = 5000$.




Figure S3. The probability distributions of the contribution of a pair of interactions to the bulk modulus K_{ij} and shear modulus G_{ij} for configurations with various Poisson's ratios. These configurations were obtained by training configurations with two species.




Figure S4. The packings with different Poisson's ratios display the 10 bonds with the largest absolute values of K_{ij} . These packings correspond to the examples presented in Fig S3. Positive K_{ij} are indicated by blue lines, while negative K_{ij} are represented by red lines.




Figure S5. The optimization process for the ensemble with target $\nu^* = -0.1$. From top to bottom and left to right, the plots are the changes of objective function, ν , diameters σ_i , and binding energies B_{ij} .




Figure S6. The optimization process for the ensemble with target $\nu^* = 1.0$. From top to bottom and left to right, the plots are the changes of objective function, ν , diameters σ_i , and binding energies B_{ij} .




Figure S7. The violin plots of validations for the optimization for the ensemble with a target $\nu^* = 0.35$ across various species numbers. The black dots represent validations for 100 independent optimizations.




Figure S8. The probability distribution of validations for the ensemble optimizations with $\nu^* = -0.1, 0.4, 0.8$.

Fig. S7 shows that the training process is reliable and reproducible. For a target of $\nu^* = 0.35$, we performed 100 independent optimizations each for $n_{sp} = 2, 4, 8, 16$. The optimizations were identical except for the seed for the random number generator that was used to choose the configurations to sample at each optimization step. As a result, the optimizations produced different solutions (sets of optimized parameters), but nevertheless resulted in very low validated losses, indicating that they were all successful.

Fig. S8 shows the probability distribution $P(\nu)$ obtained during validation. For $\nu^* = 0.4$, the distribution is narrow. However, for $\nu^* = -0.1$ and 0.8 , the distributions become broader and exhibit long tails at small values, especially for negative Poisson's ratio.

E. Ensemble optimization for desired pressure p^*

Our approach can be readily extended to address other system properties. Here, we have adapted the approach to control the pressure. As depicted in Fig. S9, the magnitude of the objective function for targeting pressure is smaller than that for targeting Poisson's ratio. In addition, the parameter adjustments during optimization for pressure




Figure S9. The evolution of an optimization process with target $p^* = -0.1$ for the ensemble. From top to bottom, left to right, the plots are the changes of objective function, pressure, diameters σ_i , and binding energies B_{ij} .

exhibit smoother transitions compared to those in the optimization for Poisson's ratio.

F. Ensemble optimization for desired structural parameter q_8^*

Although this paper primarily focuses on disordered solids, our approach can also be used to control, either by suppressing or promoting, measures of local structure. As an example, we consider the bond-orientational order parameter q_8 , a measure of local 8-fold symmetry defined by $q_8 = \frac{1}{N} \sum_{k=1}^N \left| \frac{1}{n_b} \sum_{l=1}^{n_b} e^{8i\theta_{kl}} \right|$, where the inner sum is over the neighbours of k and n_b is the number of such neighbours. q_8 is a useful measure for detecting square lattices as it can include next-nearest-neighbor information.

Figure S10 shows how we can learn to assemble a square lattice. This is a good example as it has a known solution of monodisperse repulsive particles at high densities[2]. By setting a target of $q_8^* = 1.0$, we are able to recreate this solution in less than 500 optimization steps (note that the very low values of $B_{\alpha\beta}$ mean that the energy is dominated by the harmonic repulsive interactions). In this case, our objective does not quite vanish because our protocol of instantaneous cooling leads to the formation of polycrystals. As shown in the main text, we are able to precisely tune q_8 between roughly 0.2 and 0.9.

G. Ensemble optimization for desired Poisson's ratio and pressure

Figure S11 illustrates the evolution of the design parameters for the optimization with targets $\nu^* = 0.5$ and $p^* = 0.4$ in the main text.

[1] C. P. Goodrich, S. Dagois-Bohy, B. P. Tighe, M. van Hecke, A. J. Liu, and S. R. Nagel, Jamming in finite systems: Stability, anisotropy, fluctuations, and scaling, Phys. Rev. E **90**, 022138 (2014).




Figure S10. The evolution of an optimization process with target $q_8^* = 1.0$ for the ensemble. From top to bottom, left to right, the plots are the changes of objective function, q_8 , diameters σ_i , and binding energies B_{ij} . The insert shows an example configuration with parameters at the final optimization step, noting that the diameter of particles is shrunk by 0.5 times for clearer visualization.




Figure S11. The evolution of parameters for the optimization with targets $\nu^* = 0.5, p^* = 0.4$ for the ensemble in the main text.

- [2] M. Zu, P. Tan, and N. Xu, Forming quasicrystals by monodisperse soft core particles, *Nature Communications* **8**, 2089 (2017).