A Generative Diffusion Model for Amorphous Materials

Kai Yang

University of California, Los Angeles kyang42@ucla.edu

Daniel Schwalbe-Koda

University of California, Los Angeles dskoda@ucla.edu

Abstract

Generative models show great promise for the inverse design of molecules and inorganic crystals, but remain largely ineffective within more complex structures such as amorphous materials. Here, we present a diffusion model that reliably generates amorphous structures up to 1000 times faster than conventional simulations across processing conditions, compositions, and data sources. Generated structures recovered the short- and medium-range order, sampling diversity, and macroscopic properties of silica glass, as validated by simulations and an information-theoretical strategy. Conditional generation allowed sampling large structures at low cooling rates of 10^{-2} K/ps to uncover a ductile-to-brittle transition and mesoporous silica structures. Extension to metallic glassy systems accurately reproduced local structures and properties from computational datasets, demonstrating how synthetic data can be generated from characterization results. Our methods provide a roadmap for the design and simulation of amorphous materials previously inaccessible to computational methods.

1 Introduction

Solving the inverse design of materials is one of the most coveted goals of computational materials science,[1] but sampling the space of materials exhibiting outstanding structures and properties remains challenging. In recent years, deep generative models provided an effective framework for sampling data points conditioned on labels.[2] Diffusion models, in particular, excel at learning the joint probability distributions of complex data and labels, denoising random inputs into high-fidelity samples such as images, text, and more.[3, 4] Within materials science, diffusion models have been shown to produce novel inorganic crystal structures conditioned on properties, symmetries, compositions, and other constraints, in an encouraging step for the inverse design of crystalline materials.[5–8]

Many technologically relevant solids, however, are not crystalline, but amorphous. Polymers, metallic glasses, battery electrolytes, and phase-change memory alloys are examples of materials lacking the long-range order that makes inorganic crystalline systems comparatively simpler to represent and, in many cases, generate.[9, 10] Furthermore, the rugged, high-dimensional energy landscapes of amorphous structures demand very long molecular dynamics (MD) or Monte Carlo simulations to achieve adequate convergence.[11] Thus, modeling synthesis-structure-property relationships that can enable the inverse design of amorphous materials remains challenging. Several works have attempted to use generative models, including variational autoencoders, generative adversarial networks, or diffusion models, to produce amorphous configurations.[12–17] However, comparisons between generated and simulated structures show that most models fail to either produce physically correct samples[12] or sample structures well outside the training set.[14, 18] Even when structures are reasonable, there is limited evidence that generated samples exhibit the correct macroscopic properties,[17] severely limiting the applicability of generative models in computational studies (see Appendix A.1).

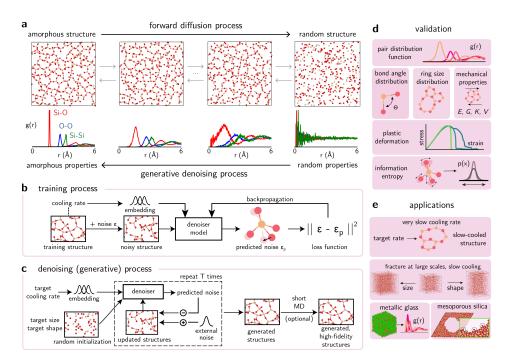


Figure 1: Training and validating a generative diffusion model for amorphous materials. a, Forward diffusion and reverse denoising for amorphous materials, illustrated with SiO_2 structure and partial pair distribution functions (PDF). b, The denoiser is trained to predict Gaussian-sampled displacements ε added to training structures, with structure-level labels (e.g., cooling rates) embedded via a Gaussian basis. c, New structures are generated from random inputs and target cooling rates, denoised over multiple steps using a DDPM-like noise schedule.[3] d, Validation metrics include short/medium-range order, network connectivity, mechanical properties, and information entropy. e, Applications include generating glassy structures under slow cooling, large-scale systems, porous structures, and reproducing metallic glass phase space.

Here, we show that a graph neural network (GNN)-based diffusion model can faithfully generate amorphous structures across diverse compositions, densities, processing conditions, and data sources (Fig. 1). Following a denoising diffusion probabilistic model (DDPM) framework[3] for atomistic systems,[19] the GNNs are trained to denoise atomic environments added to known amorphous structures (Fig. 1a,b, Sec. 2). Later, trained models are used to sample new amorphous configurations from random distributions of atoms (Fig. 1a,c). Generated amorphous structures recover short- and medium-range order, create distributions of outputs nearly indistinguishable from simulated structures, exhibit elastic and plastic behaviors compatible with their simulated counterparts, and can be conditioned on processing parameters such as cooling rates (Fig. 1d). The methods are demonstrated with applications on amorphous silica (a-SiO₂) and simulated metallic glass. The capabilities are used to generate large simulation cells, create processing-property relationships, and produce porous structures, enabling computational studies previously expensive or impractical (Fig. 1e). These results provide a step toward the inverse design of amorphous materials and an efficient tool for sampling their vast configurational space.

2 Methods

2.1 Generative diffusion model

Intuition of the generative diffusion model for amorphous structures: To create a generative model for amorphous structures, we adapted the denoising diffusion probabilistic models (DDPM) approach to amorphous materials by learning a mapping between a random distribution of atoms and a physically meaningful amorphous structure using a graph neural network (GNN).[3, 19] Figure 1a exemplifies the training and generative denoising processes for amorphous SiO₂ (a-SiO₂), with their

partial pair distribution functions (PDFs) illustrating how short- and medium-range order varies across the system. The forward diffusion process (left to right in Fig. 1a) gradually introduces Gaussian noise to atomic positions from simulated amorphous structures until the final structure resembles a random distribution of particles. After training, the model is used to generate amorphous structures from randomly sampled particles with the denoising process (Fig. 1a, right to left). Given a target cooling rate and an initial configuration with fixed density and shape, the model sequentially denoises the configuration in a process that resembles "Langevin dynamics" with a learned score. This process adds an extra noise at each denoising step, which acts as an effective "temperature" for the system and escaping local minima in the learned score function (see Appendix A.2). Within the DDPM, this extra noise follows a schedule and becomes progressively smaller throughout the denoising process.[3]

Score-based denoiser for amorphous structures: Mathematically, a generative model approximates the distribution $p(\mathbf{x})$ over atomic structures \mathbf{x} of the material system of interest. Under this DDPM strategy for atomistic structures, we train a model to approximate the distribution of atomic environments of amorphous configurations $p(\mathbf{x})$ inferred from a dataset of known configurations. This approximation of $p(\mathbf{x})$ is first performed by considering an input structure \mathbf{x}_0 and defining a Markov chain process such that

$$\mathbf{x}_0 \to \mathbf{x}_1 \to \ldots \to \mathbf{x}_T,$$
 (1)

where T is the number of diffusion steps. This forward process is defined by a distribution $q(\mathbf{x}_t|\mathbf{x}_{t-1})$, $1 \leq t \leq T$. For simplicity, the distribution q is often chosen as a Gaussian distribution, where

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}\left(\boldsymbol{\mu}_t(\mathbf{x}_{t-1}), \sigma_t^2 \mathbb{I}\right), \tag{2}$$

with (μ_t, σ_t) functions of the denoising timestep t and \mathbb{I} the identity matrix. Within the diffusion model framework, given that the ground-truth distribution p is often intractable, a model p_{θ} with parameters θ is trained to approximate, at each step t, the score function of the noise distribution $q(\mathbf{x}_t|\mathbf{x}_{t-1}),[20]$

$$\min_{\theta} \mathbb{E}_t \left[\sigma_t^2 \mathbb{E}_{\mathbf{x}_0 \sim \mathcal{D}} \mathbb{E}_{q(\mathbf{x}_t | \mathbf{x}_0)} \left[\| p_{\theta}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t | \mathbf{x}_0) \|^2 \right] \right]. \tag{3}$$

Within this notation, $p_{\theta}(\mathbf{x}_t, t)$ is called the score model, and the expectation $\mathbb{E}_{\mathbf{x}_0 \sim \mathcal{D}}$ is performed over the dataset \mathcal{D} of initial (training) structures \mathbf{x}_0 .

When the distribution q is chosen to be Gaussian with $\mu_t = 0$ and we define $\mathbf{x}_t = \mathbf{x}_{t-1} + \sigma_t \varepsilon$, $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$, its score can be written directly as

$$\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t | \mathbf{x}_{t-1}) = -\frac{\varepsilon}{\sigma_t},\tag{4}$$

which defines a per-atom displacement ε_t that maps the noisy atomic positions \mathbf{x}_t to its previous, denoised state \mathbf{x}_{t-1} .[3] By redefining the model $p_{\theta}(\mathbf{x}_t, t)$ as $\varepsilon_{\theta}(\mathbf{x}_t, t) = -\sigma_t p_{\theta}(\mathbf{x}_t, t)$ and adopting a denoising score-matching loss,[21] which simplifies the objective function of Eq. (3) to

$$\min_{\theta} \mathbb{E}_{t} \left[\mathbb{E}_{\mathbf{x}_{0} \sim \mathcal{D}} \mathbb{E}_{q(\mathbf{x}_{t} | \mathbf{x}_{t-1})} \left[\left\| \boldsymbol{\varepsilon}_{\theta}(\mathbf{x}_{t}, t) - \boldsymbol{\varepsilon} \right\|^{2} \right] \right]. \tag{5}$$

Equation (5) indicates that a model ε_{θ} is given an input structure \mathbf{x}_{t} and is trained to predict the noise ε that maps from the denoised state \mathbf{x}_{t-1} into \mathbf{x}_{t} given a noise schedule σ_{t} .

Model architecture: Following previous approaches,[19] the denoiser model uses a customized version of the E(3)-equivariant NequIP model[22] to implement the training from in Eq. (5). The model utilizes an initial embedding layer that maps atomic species information onto an 8-dimensional scalar feature space (8x0e) for both node attributes and atomic numbers, with interactions truncated at the specified cutoff radius (5 Å). The network architecture consists of three convolution layers (num_convs = 3) within the message-passing framework. Hidden layer representations are

configured as 64 scalar features and 32 vector (rotation-equivariant) features. Edge features are represented with 4x0e + 4x1o + 2x2e features encoding scalar, vector, and rank-2 tensor components. The radial basis functions are parameterized through a two-layer neural network with 16 and 64 neurons, respectively, to model distance-dependent interactions. The model considers up to 12 nearest neighbors per atom (num_neighbors = 12) and produces vector outputs, ensuring that the predicted noises are equivariant under rotations and translations while maintaining the fundamental symmetries of the physical system.

Cooling rate embedding: To incorporate cooling rate information into the model architecture, we implemented a Gaussian basis embedding module, which maps a scalar cooling rate value into a high-dimensional feature vector. Specifically, we project each cooling rate onto a set of 9 Gaussian basis functions with means distributed uniformly across the normalized cooling rate range in a \log_{10} scale, thus mapping the interval of cooling rates from $[10^{-2}, 10^3]$ K/ps to [-2, 3]. Each basis function's activation represents the similarity between the input cooling rate and that basis function's mean value, effectively creating a smoothly varying representation. These activations are then processed through a two-layer neural network with a SoftPlus nonlinearity, producing a dense embedding vector that captures the cooling rate's influence on structural properties. This embedding is summed to the hidden representation at each convolution layer of the GNN.

Training procedure: the model was trained on a single NVIDIA RTX A6000 GPU. The percoordinate, per-atom noise perturbation ε was added at train time and sampled from a normal distribution $\mathcal{N}(0,\sigma)$, where σ ranged from 0.001 to 0.75 Å based on the noise schedule. At train time, the schedule t is sampled from a uniform distribution, where $1 \le t \le T$. Maximum displacement parameters were calibrated to approximately 50% of typical bond lengths, ensuring that the diffusion process leads to fully random or highly noised configurations. Cooling rates associated with individual training configurations were embedded using a Gaussian basis set, as details described in the above section. The conditional denoiser model was trained for 40,000 iterations with a batch size of 16 and learning rate of 2×10^{-4} , with validation loss monitored to assess convergence. For unconditional denoiser models (CuZr and experimental nanoparticle systems), the number of training iterations was 20,000.

Generative denoising process: The number of atoms and lattice parameters were selected as inputs based on desired density, lattice shape, and atom types. All atoms were initially sampled from a uniform distribution within the simulation cell, meaning each of their fractional coordinates (f_x, f_y, f_z) was sampled from $\mathcal{U}(0,1)$. The denoiser model was applied to generate atomic structures from these initialized configurations with extra noise (Fig. 1c). The per-atom, per-coordinate extra noise is sampled from a normal distribution $\mathcal{N}(0,\sigma_t)$ with mean 0 and standard deviation σ_t represents the noise magnitude following a noise schedule dependent on the denoising step t. The noise magnitude was decreased linearly to zero over the extent of 2,900 generation steps. Thus, each initial configuration underwent 2,900 denoising steps with extra noise (with $\sigma_{t=0}=1.0$) followed by 100 denoising steps without extra noise. Despite being denoised from the same initial structure, adding "large" initial extra noise and using different random seeds ensures that generated structures are different from each other.

Post-generation MD refinement: Generated configurations were subsequently refined using molecular dynamics simulations consisting of 25 ps in the *NVT* ensemble followed by 25 ps in the *NPT* ensemble using the potential and settings for MD described below. This last-step MD refinement was performed for a-SiO₂ and $Cu_{50}Zr_{50}$ systems but was omitted for structures used to reproduce experimental data. Comparison studies between refined and unrefined structures are presented in Figs. 15 and 16 (see Appendix).

Density dependence on the cooling rate: For conditional denoising based on different cooling rates, simulation boxes are kept constant during the denoising process, but the cooling rate is known to influence the density. To account for this fact, we extracted densities from simulated structures across different cooling rates and identified a reasonably linear correlation between a-SiO₂ density and the logarithm of the cooling rate (see Fig. 14 and Appendix A.10) within the ranges of interest. [23, 24] Thus, when selecting the targeted cooling rate, we fix the density of the initial, randomly sampled configuration using a linear regression model valid within the range of cooling rates of interest. Future work can focus on sampling amorphous structures with variable densities throughout the denoising process.

Denoising a porous structure: To create an amorphous a-SiO₂ structure, we performed the denoising process described above with a different initial condition, where atoms were sampled according to a non-uniform distribution throughout the lattice. Atoms were sampled from a random distribution over a hexagonal lattice with predefined wall density, number of atoms, and lattice length. However, an excluded volume in the shape of a cylinder prevented atoms from being sampled within the "pore", thus mimicking a uniform atomic distribution of atoms for all regions outside the pore (Fig. 4c). Then, the porous structure was denoised following the exact same procedure described above, with a target cooling rate of 1 K/ps.

3 Results

3.1 Generating valid and novel amorphous materials

As a first case study, we trained a GNN-based diffusion model to generate a-SiO₂ structures sampled with MD simulations at a constant cooling rate of 1 K/ps (see Appendix A.3). Then, we verified whether generated structures displayed the correct short- and medium-range order for these materials. Figure 2 shows that generated structures exhibited excellent structural agreement with their simulated counterparts adopted as ground-truth. Generated structures captured the short-range order characteristics of a-SiO₂, with partial PDFs and bond angle distributions (BADs) reproducing peak positions and intensities of the reference data (Fig. 2a). At the medium range, Fig. 2c shows that generated structures reproduced ring distributions (see Appendix A.5) with excellent accuracy, and only slightly overestimated the number of 6-membered rings. Beyond structural distributions, which can average out outliers, macroscopic quantities such as mechanical properties are highly sensitive to unphysical environments, as they often lead to large variations in energy (see Appendix A.7). Thus, reproducing properties is a rigorous test to validate generative models, but have been largely underutilized within amorphous materials. To enable this test, we calculated the stiffness tensor of simulated and generated structures according to the ground-truth potential used to produce the training data (see Appendix A.3 for details). Figure 2d shows that the elastic properties computed for the generated silica structures, including bulk modulus (K), shear modulus (G), Young's modulus (E), and Poisson's ratio (ν), are statistically indistinguishable from their simulated counterparts. In contrast with methods such as reverse Monte Carlo, where structural features are well-reproduced but properties can be severely wrong, [25, 26] our models produced structures in excellent agreement with reference simulation data across all relevant figures of merit.

As generative models learn a distribution of input data points, they are not only assessed according to the validity of samples, but also according to their novelty, thus verifying whether the model exhibits "mode collapse." [2] However, in contrast to crystalline structures or molecular graphs, novelty is challenging to define in the amorphous space (see Appendix A.8 for a complete discussion). We computed a probability distribution over environments for amorphous structures following an information-theoretical strategy, [27] and quantified the overlap between distributions. This comparison between distributions resembles the Fréchet Inception Distance, [28] which accounts for novelty in terms of distribution statistics rather than point-wise metrics. Figure 2e shows that the distribution of overlap scores between different structure pairs is nearly identical. Generated structures exhibited as much overlap with simulated structures as the overlap between two simulated structures, ruling out the existence of mode collapse across the system when a large external noise is used to generate structures. On the other hand, Fig. 2f shows that the information entropy of generated samples is unreasonably large if no external noise is used during the denoising process (see Appendix A.2). These results demonstrate that our diffusion model generates valid structures according to structural features[12, 17, 18] and property distributions, as well as novel structures, as quantified by their sampling distributions.

3.2 Generating amorphous structures conditioned on processing parameters

Controlling processing conditions such as cooling rates in glassy materials is essential to model and design amorphous structures and their resulting properties.[23, 24, 29] Given the cost of performing long MD simulations, cooling rates in computational models typically revolve around 10^{12} K/s (or 1 K/ps), exceeding typical experimental cooling rates (1–100 K/s) by several orders of magnitude. Thus, learning to sample structures at very low cooling rates can unlock the ability to produce increasingly realistic models for amorphous systems. We trained our generative model on a-SiO₂

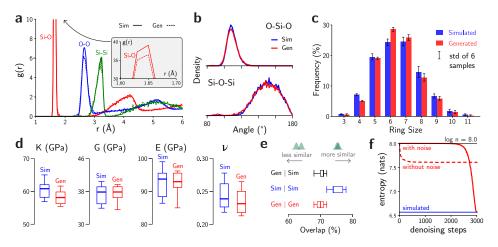


Figure 2: **a**, Partial pair distribution functions (PDFs) of simulated (solid) and generated (dashed) structures match closely, with similar trends in other metrics: **b**, Si–O–Si and O–Si–O bond angles; **c**, ring size distributions. **d**, Bulk modulus (K), shear modulus (G), Young's modulus (E), and Poisson's ratio (ν) are indistinguishable between generated and simulated structures. **e**, Generated structures follow the expected distribution of amorphous environments,[27] with Gen–Sim overlaps nearly equal to Sim–Sim overlaps. In **d** and **e**, thicker lines, boxes, and whiskers show median, interquartile range, and range of distributions, respectively. **f**, Without a large external noise schedule, generated structures have excessive outlier environments (high entropy). Adding extra noise yields more realistic results.

structures simulated with cooling rates between 10^{-1} to 10^2 K/ps, then used an input to the GNN that conditioned the generation of amorphous structures to a target cooling rate (Fig. 1c). Figure 3a demonstrates that structures generated with higher cooling rate targets exhibited higher information entropy even at constant density, reproducing known relationships between configurational entropy and cooling rates in glasses (see Appendix A.9).[30] Moreover, Fig. 3b shows that the dependence of average Si-O-Si angles with respect to the thermal history are reasonably captured in the generated a-SiO₂ structures, in agreement with simulated counterparts and the literature.[23, 24, 29] Beyond structural features, Fig. 3c shows the Young's modulus (E) for simulated and generated a-SiO₂ structures across different cooling rates. The values of Young's moduli remain consistent across cooling rates for both simulated and generated a-SiO₂ structures, ranging between 85–100 GPa. This agreement is also true for all other elastic properties, including bulk modulus, shear modulus, and Poisson's ratio (Figs. 6–9 in Appendix), confirming our ability to conditionally generate structures across cooling rates while reproducing their property distributions (see Appendix A.10 for an extended analysis).

Remarkably, the trends in Figs. 3b,c generalize to both faster (10^3 K/ps) and slower (10^{-2} K/ps) cooling rates beyond the training regime (see Fig. 10 for more examples), which can be used to generate structures at low cooling rates that are expensive for traditional simulations. To illustrate the impact of this observation, Fig. 3d compares the estimated computational cost (in CPU-h) for simulating amorphous structures with 3000 atoms across cooling rates compared to the time required to evaluate the generative model (see Appendix A.4 for hardware details). As slow cooling rates can only be simulated by longer MD simulations, the computational cost of simulating realistic glasses follows a power law with the cooling rate, while the generative model produces structures at constant time. At slow cooling rates such as 0.01 K/ps, the diffusion model produces a-SiO₂ structures at least 1000 times faster (in CPU-h) than MD simulations (Fig. 3d) for a system with 3000 atoms. Performing melt-quench simulations for the largest system in this work (112,848 atoms, see Appendix 3.3) at 0.01 K/ps would require about 89,702 CPU-h using MD in our hardware (see Appendix A.5), a substantial investment of computational resources. In practice, we use GPUs to generate structures with the GNN, further reducing the wall times. Even when compared to GPU-accelerated MD, the generative process drastically improves the effectiveness of sampling amorphous structures at low cooling rates and large sizes, demonstrating that the correct processingstructure-property relationships can be captured with the model.

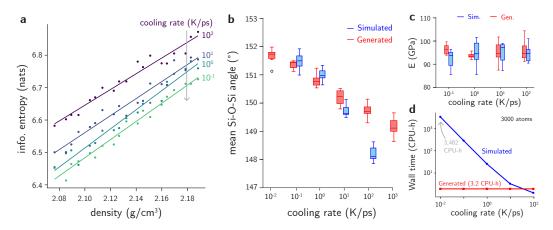


Figure 3: Conditional generation of a-SiO₂ using diffusion models recover structural features and properties across cooling rates. a, Diffusion models reproduce the correct information entropy trends for amorphous structures across cooling rates and densities; faster cooling yields more outlier environments and higher entropy. b, Bond angle distribution trends are well captured across cooling rates, even beyond the training domain $(10^{-2}-10^3 \text{ K/ps})$ c, Young's modulus (E) of generated a-SiO₂ matches simulations despite sensitivity to outliers. In b and c, boxplots show median, interquartile range, and range of the distribution, respectively. d, Simulation cost in melt-quench rises sharply at low cooling rates, while generative model inference remains constant; at 10^{-2} K/ps the time difference reaches three orders of magnitude.

3.3 Fracture tests with large and slow-cooled amorphous silica

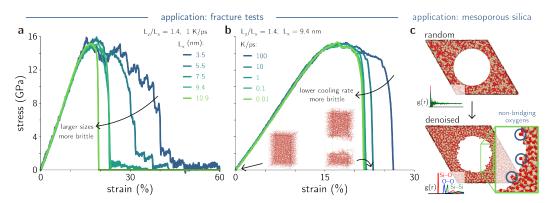


Figure 4: Shape- and size-agnostic amorphous structure generation enables simulations on fracture and structural porosity. a, Stress-strain curves show that generated a-SiO₂ structures recover the right behavior of elastic and plastic deformation behavior across length scales, with more ductile fracture at smaller length scales while maintaining a constant ultimate tensile strength at fixed aspect ratio ($L_z/L_x=1.4$). b, Stress-strain curves from generated a-SiO₂ structures of the same shape and size ($L_z/L_x=1.4$, $L_x=9.4$ nm) under conditional cooling rates show that slower cooling rate leads to a more brittle behavior for amorphous silica. c, The model generates mesoporous a-SiO₂ structures without training on surfaces or pores. Partial pair distribution functions (PDFs) confirm the structures remain amorphous, and pore surfaces reproduce experimental non-bridging oxygen concentrations, enabling modeling of silanol defects and other interfacial behaviors in amorphous mesoporous silica.

As our generative model samples structures for any input size or shape and for a range of cooling rates, we investigated the fracture behavior of generated a-SiO₂ structures for different length scales and processing conditions. Experimental observations have shown that a-SiO₂ exhibited a brittle-to-ductile transition at the nanoscopic regime,[31, 32] while computational studies were typically limited by small sizes and fast cooling rates.[23, 24, 33] We used our diffusion model to generate

a-SiO₂ with different sizes and cooling rates, and performed fracture test simulations for each of them (see Appendix A.3). Figure 4a shows the stress-strain curves obtained for a-SiO₂ structures generated with a fixed aspect ratio of $L_z/L_x = 1.4$, as z being the direction of the tensile strain, and L_x varying from 3.5 nm (3,735 atoms) to 10.9 nm (112,848 atoms). Generated structures reproduced the increasingly brittle behavior of amorphous silica at large length scales, while smaller simulation lengths ($L_x < 6$ nm) were more ductile (see Appendix A.11). The influence of low cooling rate effects on fracture, however, is often not studied at large length scales due to its associated computational cost. [24, 34] On the other hand, our generative model allows sampling these large, slow-cooled structures at minimal computational cost. Figure 4b compares the behavior of slowand fast-cooled structures for a-SiO₂ structures generated at large sizes ($L_z/L_x = 1.4$, $L_x = 9.4$ nm) and different target cooling rates. The results show that lower cooling rates decreased the ductility of the amorphous material, [35] but up to a certain limit. The fracture strain for samples generated with cooling rates of 0.1 and 0.01 K/ps are nearly identical even though the structures are slightly different (Fig. 3b), showing how the fracture behavior converges at very low cooling rates. This contrasts with the continuous trend observed in simulated glasses at much smaller system sizes of 3000 atoms, for which the ductility was drastically overestimated and continued to decrease at low cooling rates (Figs. 11,12 in Appendix).

3.4 Generating mesoporous amorphous structures

Mesoporous silica particles exhibiting pores with diameters between 2–50 nanometers are ubiquitous in many applications ranging from catalysis and separation to drug delivery. [36, 37] Despite their wide uses in experiments, computational sampling of these structures is laborious and can lead to incorrect pore structures or defect concentrations.[38, 39] We used our generative approach to produce a porous a-SiO₂ structure by denoising a structure from atoms sampled in a hexagonal lattice with a cylindrical pore (see Appendix A.5, Fig. 4c). Despite never been trained on surfaces or porous structures, our diffusion model generated excellent amorphous surfaces and a wall that exhibited the right short-range ordering. Interestingly, generated porous structures correctly captured the atomic structure at the pore surface, comprised mostly of pristine SiO₄ tetrahedra and non-bridging oxygens (NBO). The concentrations of NBOs were also in agreement with typical ranges of 2-4 NBOs/nm² from experimental data, [40, 41] and were a function of the initial wall density used as starting configuration for the denoiser. At higher densities, the pore tended to shrink and form more densed NBOs at the pore wall, while low-density walls tended to form less NBOs at the surface (Fig. 13 in Appendix). As NBOs are responsible for a variety of properties of mesoporous silica in the presence of hydrogen, such as catalytic or hydrophobicity, obtaining these substructures directly from the generative approach can drastically accelerate computational modeling in these fields.

3.5 Generating metallic glasses

To extend our generative approach beyond a-SiO₂, we tested our methodology against Cu-Zr metallic glass structures and properties, a non-covalent amorphous system with a challenging potential energy surface. Amorphous Cu-Zr structures with a fixed composition of Cu₅₀Zr₅₀ and 5000 atoms were obtained from Ref. [42] to train a diffusion model without conditional targets. Then, given a box with randomly sampled atoms at a chosen composition and density, we used the diffusion model to generate the structure of the metallic glass, which can be followed by a short, post-denoising MD equilibration (Sec. 2). Figure 5a shows that the partial PDFs between simulated and generated Cu₅₀Zr₅₀ glasses are in excellent agreement with each other, with generated structures showing near-perfect peak positions and only slightly lower intensities compared to the reference structures. This is true even for long-range peaks around 7 and 9 Å, which were outside of the cutoff of the GNN (5 Å), but were treated indirectly by the message-passing layers. The atomic volume in Fig. 5b and fractions of Voronoi indices in Fig. 5c for simulated and generated Cu₅₀Zr₅₀ also show good agreement with each other and only underestimated (0,0,12,0) icosahedral motifs in generated samples. Given the sensitivity of Voronoi polyhedra to the structural fidelity, the ability of the model to correctly recover these trends is remarkable. Furthermore, Fig. 5d shows that the elastic and plastic regimes of the generated amorphous Cu-Zr structures (6 independent samples) are also compatible with the reference curves (2 samples). Though the yield stress is slightly underestimated in the generated structures, mirroring the results in Fig. 12 in Appendix, the strength curves of both metallic glasses are within good agreement with each other, especially when considering the sensitivity of stress-strain curves from these metallic glasses depending on the sample preparation[43]

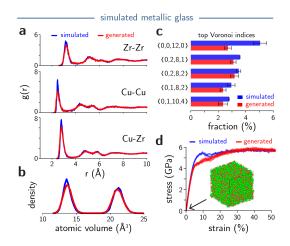


Figure 5: Generating metallic glass structures. a, Partial pair distribution functions (PDFs) g(r) of generated $Cu_{50}Zr_{50}$ structures (red lines) exhibit good agreement with PDFs of their simulated counterparts (blue lines), including capturing the shorter Cu-Cu pair interactions compared to Cu-Zr and Zr-Zr pairs. b, Distributions of atomic volume exhibit great agreement between the coordinations and local environments of generated $Cu_{50}Zr_{50}$ structures (red) and their simulated, ground-truth counterparts (blue). c, Generated $Cu_{50}Zr_{50}$ structures reproduce the fractions of major Voronoi polyhedra compared to their simulated counterparts, only under-representing $\langle 0,0,12,0\rangle$. Error bars for indices indicate the range of the distribution. d, Stress-strain behavior of generated $Cu_{50}Zr_{50}$ structures (red, 5 samples) approach the ones from their simulated $Cu_{50}Zr_{50}$ counterparts (blue, 2 samples), with equivalent Young's modulus and strength, despite slightly underestimated yield strength.

and the limited training data used for the denoiser model. This proves that the generative model is not limited to amorphous silica networks, but can also produce valid amorphous structures and properties for metallic glassy systems.

4 Discussion

Our generative diffusion model substantially advances computational modeling of amorphous materials, accurately producing physically realistic structures and properties across various compositions, data sources, and processing conditions. Unlike previous generative methods for amorphous materials, whose unphysical artifacts or constrained structural diversity limited their broad use in simulations, our model consistently replicates key short- and medium-range structural characteristics and mechanical properties, highlighting its robustness and reliability. Importantly, this framework offers efficient computational performance, drastically reducing the computational cost typically associated with molecular dynamics simulations. Our work provides a roadmap not only for the training process of diffusion models for amorphous materials, but also how to quantitatively assess model quality in terms of structural novelty, properties, and applications.

As with any model, our approach exhibits limitations. Occasional outlier environments, although rare, require short MD simulations to refine the generated structure. Furthermore, though several properties are predicted with high accuracy, yield strengths are still slightly underestimated, showing this metric may be a good indicator of model quality within generative models for amorphous structures as the field progresses. Future research can extend this approach toward broader classes of complex materials, including metallic, semiconducting alloys, polymers or multi-component systems. Additionally, integrating generative methods with experimental characterization can automate structural elucidation, particularly in fields where experimental data is scarce or costly to obtain. The adaptability and diversity of applications demonstrated here suggest that generative diffusion models can already be used in a range of computational materials science problems, accelerating their modeling and discovery.

Data Availability

The datasets generated in this work will be released before AI4Mat-NeurIPS-2025 workshop starts and can be accessed through project blog post https://kaiyang1010.github.io/DM2/. The datasets of metallic glasses were obtained from the original source[42] at: https://figshare.com/articles/dataset/Heterogeneous_thermal_activation_energy_in_Cu-Zr_metallic_glasses/12485795.

Code Availability

The code from this work was based on the code from Hsu et al.,[19] available at https://github.com/LLNL/graphite. The information theoretical analysis was performed with the QUESTS package from Schwalbe-Koda et al.,[27] available at https://github.com/dskoda/quests. The codebase to reproduce the results/plots from this work will be released before AI4Mat-NeurIPS-2025 workshop starts and can be accessed through project blog post https://kaiyang1010.github.io/DM2/.

Acknowledgements

This work was supported by Toyota Research Institute under the Synthesis Advanced Research Challenge. This work used computational and storage services associated with the Hoffman2 Shared Cluster provided by UCLA Office of Advanced Research Computing's Research Technology Group. The authors thank Linda Hung, Amalie Trewartha, Steven Torrisi, Fei Zhou, Jiawei Guo, and Long Qi for discussions and suggestions regarding this work. The authors also thank Jianwei (John) Miao and his group for making the experimental dataset available.

Conflicts of Interest

The authors have no conflicts to disclose.

Author Contributions

Kai Yang: Formal Analysis; Data Curation; Investigation; Software; Validation; Visualization; Writing - Original Draft; Writing - Review & Editing. **Daniel Schwalbe-Koda:** Conceptualization; Formal Analysis; Investigation; Methodology; Project Administration; Visualization; Writing - Original Draft; Writing - Review & Editing; Funding Acquisition; Supervision.

References

- [1] Alex Zunger. Inverse design in search of materials with target functionalities. *Nature Reviews Chemistry*, 2(4):0121, 2018.
- [2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 27, 2014.
- [3] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [4] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [5] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S. Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=03RLpj-tc_.

- [6] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, 639:624–632, 2025.
- [7] Gabe Guo, Tristan Luca Saidi, Maxwell W Terban, Michele Valsecchi, Simon JL Billinge, and Hod Lipson. Ab initio structure solutions from nanocrystalline powder diffraction data via diffusion models. *Nature Materials*, 2025. doi: 10.1038/s41563-025-02220-y. URL https://doi.org/10.1038/s41563-025-02220-y.
- [8] Peichen Zhong, Xinzhe Dai, Bowen Deng, Gerbrand Ceder, and Kristin A Persson. Practical approaches for crystal structure predictions with inpainting generation and universal interatomic potentials. *arXiv:2504.16893*, 2025.
- [9] Daniel Schwalbe-Koda and Rafael Gómez-Bombarelli. Generative models for automatic chemical design. In *Machine Learning Meets Quantum Physics*, pages 445–467. Springer, 2020.
- [10] Yuanbin Liu, Ata Madanchi, Andy S Anker, Lena Simine, and Volker L Deringer. The amorphous state as a frontier in computational materials design. *Nature Reviews Materials*, 10: 228–241, 2024.
- [11] Shingo Urata, Marco Bertani, and Alfonso Pedone. Applications of machine-learning interatomic potentials for modeling ceramics, glass, and electrolytes: A review. *Journal of the American Ceramic Society*, 107(12):7665–7691, 2024.
- [12] Massimiliano Comin and Laurent J Lewis. Deep-learning approach to the structure of amorphous silicon. *Physical Review B*, 100(9):094107, 2019.
- [13] Michael Kilgour, Nicolas Gastellu, David YT Hui, Yoshua Bengio, and Lena Simine. Generating multiscale amorphous molecular structures using deep learning: a study in 2d. *The Journal of Physical Chemistry Letters*, 11(20):8532–8537, 2020.
- [14] Qiyuan Chen, Ajay Annamareddy, Ying-Fei Li, Dane Morgan, and Bu Wang. Physical regularized hierarchical generative model for metallic glass structural generation and energy prediction. arXiv:2505.09977, 2025.
- [15] Honglin Li, Chuhao Liu, Yongfeng Guo, Xiaoshan Luo, Yijie Chen, Guangsheng Liu, Yu Li, Ruoyu Wang, Zhenyu Wang, Jianzhuo Wu, et al. Conditional generative modeling for amorphous multi-element materials. *arXiv*:2503.07043, 2025.
- [16] Xiang Xu and Jingyi Hu. A generative adversarial networks (GAN) based efficient sampling method for inverse design of metallic glasses. *Journal of Non-Crystalline Solids*, 613:122378, 2023.
- [17] Hyuna Kwon, Tim Hsu, Wenyu Sun, Wonseok Jeong, Fikret Aydin, James Chapman, Xiao Chen, Vincenzo Lordi, Matthew R Carbone, Deyu Lu, Fei Zhou, and Tuan Anh Pham. Spectroscopy-guided discovery of three-dimensional structures of disordered materials with diffusion models. *Machine Learning: Science and Technology*, 5(4):045037, 2024.
- [18] Adrian Xiao Bin Yong, Tianyu Su, and Elif Ertekin. Dismai-Bench: benchmarking and designing generative models using disordered materials and interfaces. *Digital Discovery*, 3(9): 1889–1909, 2024.
- [19] Tim Hsu, Babak Sadigh, Nicolas Bertin, Cheol Woo Park, James Chapman, Vasily Bulatov, and Fei Zhou. Score-based denoising for atomic structure identification. *npj Computational Materials*, 10(1):155, 2024.
- [20] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005. URL http://jmlr.org/papers/v6/hyvarinen05a.html.
- [21] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.

- [22] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13(1): 2453, 2022.
- [23] Katharina Vollmayr, Walter Kob, and Kurt Binder. Cooling-rate effects in amorphous silica: A computer-simulation study. *Physical Review B*, 54(22):15808, 1996.
- [24] J Matthew D Lane. Cooling rate and stress relaxation in silica melts and glasses via microsecond molecular dynamics. *Physical Review E*, 92(1):012320, 2015.
- [25] Robert L McGreevy. Reverse monte carlo modelling. *Journal of Physics: Condensed Matter*, 13(46):R877, 2001.
- [26] Phillip M Maffettone, William JK Fletcher, Thomas C Nicholas, Volker L Deringer, Jane R Allison, Lorna J Smith, and Andrew L Goodwin. When can we trust structural models derived from pair distribution function measurements? *Faraday Discussions*, 255:311–324, 2025.
- [27] Daniel Schwalbe-Koda, Sebastien Hamel, Babak Sadigh, Fei Zhou, and Vincenzo Lordi. Model-free estimation of completeness, uncertainties, and outliers in atomistic machine learning using information theory. *Nature Communications*, 16:4014, 2025. doi: 10.1038/s41467-025-59232-0.
- [28] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30, 2017.
- [29] Xin Li, Weiying Song, Kai Yang, NM Krishnan, Bu Wang, Morten M Smedskjaer, John C Mauro, Gaurav Sant, Magdalena Balonis, and Mathieu Bauchy. Cooling rate effects in sodium silicate glasses: Bridging the gap between molecular dynamics simulations and experiments. *The Journal of Chemical Physics*, 147(7), 2017.
- [30] Ludovic Berthier, Misaki Ozawa, and Camille Scalliet. Configurational entropy of glass-forming liquids. *The Journal of Chemical Physics*, 150(16):160902, 2019.
- [31] F Célarié, S Prades, Daniel Bonamy, L Ferrero, E Bouchaud, C Guillot, and C Marliere. Glass breaks like metal, but at the nanometer scale. *Physical Review Letters*, 90(7):075504, 2003.
- [32] Junhang Luo, Jiangwei Wang, Erik Bitzek, Jian Yu Huang, He Zheng, Limin Tong, Qing Yang, Ju Li, and Scott X Mao. Size-dependent brittle-to-ductile transition in silica glass nanofibers. *Nano Letters*, 16(1):105–113, 2016.
- [33] Fenglin Yuan and Liping Huang. Molecular dynamics simulation of amorphous silica under uniaxial tension: From bulk to nanowire. *Journal of Non-Crystalline Solids*, 358(24):3481–3487, 2012.
- [34] Zhen Zhang, Simona Ispas, and Walter Kob. The critical role of the interaction potential and simulation protocol for the structural and mechanical properties of sodosilicate glasses. *Journal of Non-Crystalline Solids*, 532:119895, 2020.
- [35] Meng Fan, Minglei Wang, Kai Zhang, Yanhui Liu, Jan Schroers, Mark D Shattuck, and Corey S O'Hern. Effects of cooling rate on particle rearrangement statistics: Rapidly cooled glasses are more ductile and less reversible. *Physical Review E*, 95(2):022611, 2017.
- [36] a CT Kresge, M E Leonowicz, W Jelal Roth, JC Vartuli, and JS Beck. Ordered mesoporous molecular sieves synthesized by a liquid-crystal template mechanism. *Nature*, 359(6397):710–712, 1992.
- [37] Fangqiong Tang, Linlin Li, and Dong Chen. Mesoporous silica nanoparticles: synthesis, biocompatibility and drug delivery. *Advanced Materials*, 24(12):1504–1534, 2012.
- [38] Sebastian Calderon V, Tânia Ribeiro, José Paulo S Farinha, Carlos Baleizão, and Paulo J Ferreira. On the Structure of Amorphous Mesoporous Silica Nanoparticles by Aberration-Corrected STEM. Small, 14(40):1802180, 2018.

- [39] Ellie L Fought, Yong Han, Theresa L Windus, Igor I Slowing, Takeshi Kobayashi, and James W Evans. Modeling of linear nanopores in a-SiO2 tuning pore surface structure. *Microporous and Mesoporous Materials*, 341:112077, 2022.
- [40] Xiu S Zhao, GQ Lu, AK Whittaker, GJ Millar, and HY Zhu. Comprehensive study of surface chemistry of MCM-41 using 29Si CP/MAS NMR, FTIR, pyridine-TPD, and TGA. *The Journal of Physical Chemistry B*, 101(33):6525–6531, 1997.
- [41] Svetlana A Kozlova and Sergey D Kirik. Post-synthetic activation of silanol covering in the mesostructured silicate materials MCM-41 and SBA-15. *Microporous and Mesoporous Materials*, 133(1-3):124–133, 2010.
- [42] Qi Wang, Jun Ding, Longfei Zhang, Evgeny Podryabinkin, Alexander Shapeev, and Evan Ma. Predicting the propensity for thermally activated β events in metallic glasses via interpretable machine learning. *npj Computational Materials*, 6(1):194, 2020.
- [43] Aya Nawano, Weiwei Jin, Jan Schroers, Mark D Shattuck, and Corey S O'Hern. Characterizing the mechanical response of metallic glasses to uniaxial tension using a spring network model. *Physical Review Materials*, 7(7):073605, 2023.
- [44] Zhenze Yang, Yifan Wu, Xu Han, Ziqing Zhang, Haoen Lai, Zhenliang Mu, Tianze Zheng, Siyuan Liu, Zhichen Pu, Zhi Wang, et al. A unified predictive and generative solution for liquid electrolyte formulation. *arXiv*:2504.18728, 2025.
- [45] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton. LAMMPS a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comp. Phys. Comm.*, 271:108171, 2022. doi: 10.1016/j.cpc.2021.108171.
- [46] M Bouhadja, N Jakse, and A Pasturel. Structural and dynamic properties of calcium aluminosilicate melts: a molecular dynamics study. *The Journal of Chemical Physics*, 138(22): 224510, 2013.
- [47] Leandro Martínez, Ricardo Andrade, Ernesto G Birgin, and José Mario Martínez. PACKMOL: A package for building initial configurations for molecular dynamics simulations. *Journal of Computational Chemistry*, 30(13):2157–2164, 2009.
- [48] YQ Cheng, E Ma, and HW Sheng. Atomic level structure in multicomponent bulk metallic glass. *Physical Review Letters*, 102(24):245501, 2009.
- [49] Alexander Stukowski. Visualization and analysis of atomistic simulation data with OVITO-the Open Visualization Tool. *Modelling and Simulation in Materials Science and Engineering*, 18 (1):015012, JAN 2010. ISSN 0965-0393. doi: {10.1088/0965-0393/18/1/015012}.
- [50] Sébastien Le Roux and Philippe Jund. Ring statistics analysis of topological networks: New approach and application to amorphous GeS2 and SiO2 systems. *Computational Materials Science*, 49(1):70–83, 2010.
- [51] Kai Yang, Xinyi Xu, Benjamin Yang, Brian Cook, Herbert Ramos, NM Anoop Krishnan, Morten M Smedskjaer, Christian Hoover, and Mathieu Bauchy. Predicting the young's modulus of silicate glasses using high-throughput molecular dynamics simulations and machine learning. *Scientific Reports*, 9(1):8739, 2019.
- [52] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018.
- [53] Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. GuacaMol: benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, 2019.

- [54] Josh Leeman, Yuhan Liu, Joseph Stiles, Scott B Lee, Prajna Bhatt, Leslie M Schoop, and Robert G Palgrave. Challenges in high-throughput inorganic materials prediction and autonomous synthesis. *PRX Energy*, 3(1):011002, 2024.
- [55] JH Wu and Q Jia. Thermodynamic universality of configurational entropy in glass-forming liquids. *Journal of Alloys and Compounds*, 692:698–704, 2017.
- [56] Ondrej Gedeon. Configurational entropy relaxation of silica glass—molecular dynamics simulations. *Entropy*, 23(7):885, 2021.
- [57] Raúl Barciela, Thiruvilla S Mahadevan, Félix Quintero, Juan Pou, and Jincheng Du. Size effects on the fracture behavior of amorphous silica from molecular dynamics simulations. *Journal of Non-Crystalline Solids*, 631:122935, 2024.

A Supplementary Text

A.1 Evaluating generative models for amorphous materials

Good generative models for amorphous materials should be able to reproduce quantitative distributions of structural features and properties for these systems. Traditional validation strategies in amorphous materials rely on coordination environments, order parameters, energy distributions, or similar structural features that can be easily computed from the structure. Deep learning methods often fall short of producing physically reasonable amorphous structures compared to simulation results — exhibiting unphysical coordination environments, large deviations from expected pair correlation functions, or unrealistic defects — let alone reproducing their properties or accounting for processing parameters such as cooling rates.[18] As examples of recent work related to amorphous materials, Yang et al. presented a forward predictive model and an inverse generative approach for liquid electrolyte compositions,[44] though not all tasks are performed at the atomistic scale. Recently, Chen et al. used a graph variational autoencoder, with energy and radial distribution function (RDF) regularizations, to generate glass structures but only validate their generations with energy distributions and RDFs.[14] Furthermore, qualitative visualizations or distance/descriptor distributions are insufficient to prove whether generated structures are simultaneously accurate, novel, and valid. The work from Yong et al.[18] thoroughly benchmarks recent generative models' ability to produce realistic samples of amorphous and disordered systems, highlighting that most models cannot produce amorphous structures even at the qualitative level. In that work, structural characteristics and properties relevant to the analysis of amorphous materials have not been explored, however. Kwon et al. showed that a generative diffusion model can be used to generate amorphous carbon matching conditioned on characterization data.[17] Nevertheless, the resulting structures still showed outlier environments, and it remained unclear whether generative models can be used to generate and reproduce the properties of more complex amorphous materials such as multi-component ones, especially when conditioned to processing parameters.

In this work, we proposed a suite of tests to validate the quality of generated amorphous structures. As a first figure of merit, structures are evaluated according to their short- and medium-range order, which includes pair distribution functions, bond angles, and network topology. These distributions are commonly used to understand and analyze the structure of amorphous materials, and often correlated with experimental observations. As a second tier of tests, we computed the macroscopic properties of simulated and generated structures such as their elastic tensor and plastic behavior. These mechanical properties are highly sensitive to outlier environments, as minor unphysical features can produce large variations in energy (see Appendix A.7 below), thus providing an excellent test case to assess the robustness of the generative models. Finally, we introduced an information-theoretical metric to evaluate whether structures are sampled from similar probability distribution of atomistic environments, which is useful to quantify novelty (see Appendix A.8 below). These strategies are summarized in Fig. 1d of the main text.

A.2 Importance of an extra noise in denoising process

Adding a high level of extra noise during the denoising process is critical to generate valid and novel amorphous structures. The noise magnitude σ is analogous to a temperature in conventional meltquench molecular dynamics simulations, a connection that has been explored before in the context of Langevin dynamics within diffusion models.[3] High σ values in the beginning of the denoising process are analogous to elevated temperatures, where atoms possess sufficient kinetic energy to escape local minima and break structural correlations imposed by the initial state. Furthermore, the progressive reduction of σ to zero during denoising is similar to the controlled cooling process in MD quenching. The only difference in this analogy is that the generative model can produce structures with different cooling rates at a fixed number of steps, whereas an MD-based melt-quench process requires long simulations to properly cool the structures. Denoising without extra noise resembles instantaneous quenching from a liquid to a glassy state, preventing adequate structural relaxation.

Partial pair distribution function (PDF) analysis in Fig. 15 shows that amorphous structures with superior structural quality are obtained when extra noise is employed during the denoising process. Importantly, when approximating the distribution of atomistic environments in the training data, the generative model makes no consideration about the energy of the system. Sometimes, this leads to a few outlier environments that, while rare, can be nonphysical for the simulated system, such as

Si-O bonds shorter than 1.5 Å in a-SiO₂ (Fig. 15a,b). While both approaches generate amorphous structures with reasonable short-range order, structures generated with extra noise exhibit denser bond pair distributions and more accurate peak intensities compared to reference simulations (Fig. 15c,d). Directly applying MD fine tuning process on structures generated without extra noise can eliminate unphysical short Si-O bonds and enhance partial PDF performance; however, the second coordination shell remains poorly reproduced (Fig. 15b). On the other hand, the small amounts of outlier environments in structures generated with extra noise can be readily corrected through picosecond-long MD simulations after the denoising process and show excellent agreement with the simulated reference (Fig. 15d).

As discussed in the main text and Methods, structural diversity of amorphous structures was quantified using overlap scores between six independently generated a-SiO₂ structures. Without extra noise, overlap scores were distributed either excessively high (approaching 100%, indicating mode collapse) or unusually low (below 45%, suggesting structural inaccuracy) as shown in Fig. 18a. Conversely, structures generated with extra noise exhibited consistent overlap scores around 70%, matching nearly perfectly the balance between diversity and validity observed in simulated structures (Fig. 18b). When compared to MD simulated reference structures, configurations generated without extra noise showed poor correlation (overlap scores around 29%, Fig. 19a), whereas those with extra noise maintained consistent accuracy with overlap scores around 70% (Fig. 19b). Altogether, these figures of merit (PDF analysis, overlap score distributions, and structural diversity metrics) provide evidence on the importance of using a large extra noise when generating amorphous materials using denoiser models.

A.3 Molecular dynamics simulations

All molecular dynamics (MD) simulations were performed using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) software[45] (v. 17 Apr 2024). A fixed integration timestep of 1.0 fs was used throughout all simulations.

 ${
m SiO_2}$ simulations: To generate training data for the ${
m SiO_2}$ systems, we conducted MD simulations of amorphous silica comprised of 3000 atoms across multiple cooling rates. We executed six independent simulations at each cooling rate, maintaining the same settings while varying initial configurations. We adopt the Jakse interatomic potential,[46] parametrized from *ab initio* calculations, to model a-SiO₂. Short-range interactions are calculated with an 8.0 Å cutoff. Coulomb interactions are calculated by adopting the Fennell damped shifted force model with a damping parameter of 0.25 Å⁻¹ and a global cutoff of 8.0 Å. Periodic boundary conditions were applied in all three directions during MD simulations.

All a-SiO₂ samples were prepared via conventional melt-quench method. Initial structures were generated by randomly placing SiO₂ molecules in cubic boxes using PACKMOL (v. 20.13.0),[47] enforcing a 2.0 Å minimum separation between molecules to prevent unphysical overlaps. Following energy minimization, configurations underwent sequential 50 ps relaxations in the canonical (*NVT*) and isothermal-isobaric (*NPT*) ensembles at 300 K. Complete melting was achieved with a 50 ps-long *NPT* equilibration at 5000 K and zero pressure, eliminating the "memory" of the initial configuration. Subsequently, liquid systems were cooled from 5000 to 300 K under zero pressure *NPT* conditions at rates of 10^{-1} , 10^{0} , 10^{1} , and 10^{2} K/ps. The obtained glass structures are further relaxed one last time at 300 K for 100 ps in the *NPT* ensemble. This process follows the standards from previous work[29] shown to reproduce multiple experimental features from a-SiO₂ glasses.

CuZr simulations: We adopted the published dataset from Wang *et al.*[42], which used a set of optimized embedded-atom method (EAM) potentials to simulate CuZr metallic glass models.[48] Specifically, we used two of their $Cu_{50}Zr_{50}$ configurations, each containing 5,000 atoms, to train our denoiser model. The potential from Cheng *et al.*[48] was later used to perform the short MD refinement and the fracture test (described below) for generated CuZr structures. The post-denoising MD simulation was 25 ps-long at the *NVT* ensemble, followed by a 25 ps-long simulation at the *NPT* ensemble, both at 300 K.

Computation of elastic properties: the stiffness tensor C_{ij} of the equilibrated glasses is computed by performing a series of 6 deformations (i.e., 3 axial and 3 shear deformations along the 3 axes) and calculating the curvature of the potential energy U:

$$C_{ij} = \frac{1}{V} \frac{\partial^2 U}{\partial e_i \partial e_j},\tag{6}$$

where V is the glass volume, e is the strain, and i,j are the indices corresponding to each Cartesian direction. All simulated configurations exhibited near-complete isotropy. Bulk modulus (K), shear modulus (G), Young's modulus (E), and Poisson's ratio (ν) are then calculated from the stiffness tensors.

Fracture tests: fracture simulations are performed on different starting bulk configurations by deforming the structures along a single direction during an MD simulation. Uniaxial deformation is achieved by imposing a constant strain rate of 10^9 s⁻¹ along the z direction while allowing lateral dimensions (x and y) to relax freely under zero lateral pressure, approximating realistic uniaxial tension conditions. The deformation process is controlled using an *NPT* ensemble maintained at 300 K and 1 bar. True stress is calculated from the negative z-component of the stress tensor.

A.4 Computational cost estimates

We used a single CPU core of Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz from UCLA's Hoffman2 Cluster to estimate the computational cost of MD simulation and denoiser generation for 3000 atoms a-SiO₂. The entire MD simulation process includes initialization, melting, cooling and relaxation parts as detailed above, where only the cooling part is affected by the choice of cooling rate. The denoising process (inference) consists of a total of 3000 denoising steps (2900 steps with extra noise and 100 steps without extra noise). The trained denoiser model is also saved, loaded, and evaluated on the same single CPU for this comparison. In practice, we used GPUs for the generative inference process, which perform faster than using CPUs for inference of the GNNs.

To estimate the computational cost of traditional melt-quench MD simulations for large systems at slow cooling rates, we benchmarked structures at target sizes for 25 ps. All simulations include fixed computational overhead include 50 ps *NVT* initialization, 50 ps *NPT* initialization, 50 ps *NPT* melting, and 100 ps *NPT* relaxation (250 ps total). The quenching process duration varies with cooling rate, requiring 47, 470, 4,700, 47,000 and 470,000 ps for 10^2 , 10^1 , 10^0 , 10^{-1} , and 10^{-2} K/ps, respectively. Total computational cost is calculated by multiplying the average time per step obtained for a 25 ps MD simulation by the total number of steps to perform the simulation. For the large system estimate, the cost of performing a 25 ps MD simulation of 112,848 atoms within a-SiO₂ and the simulation settings described above is 4.77 CPU-h. This results on the total estimated cost of 89,702 CPU-h for the same size under 0.01 K/ps cooling rate described in the main text.

A.5 Structural analysis

All structural features of simulated a-SiO $_2$ structures were performed with six independent samples of 3000 atoms each. Samples are obtained either through the traditional melt-quench process, or using the generative models. All structural properties reported in the main text (e.g., Fig. 2,3) are reported as either averages over these six samples (e.g., structural features in Fig. 2a–c) or distributions (e.g., mechanical properties in Fig. 2d). This allows us to perform statistically relevant comparisons between generated and simulated structures by accounting for the distribution of structures and properties in amorphous materials.

Pair distribution functions: short-range structural ordering was characterized through pair distribution function (PDF) analysis. All PDFs were calculated using OVITO[49] with 100 bins between 0 and 8 Å.

Ring size distribution: the medium-range order structure of amorphous silica was characterized through ring size distribution analysis. These calculations were performed using the RINGS package,[50] with Si–O bonds defined using a 2.0 Å cutoff distance for ring identification within the network structure, which is in agreement with the partial PDF for Si-O pairs in Fig. 2a.

Concentration of non-bridging oxygens: we quantified the concentration of non-bridging oxygens (NBOs) on generated mesoporous silica surfaces by calculating the ratio of NBO atoms to the estimated surface area of the pore. As described in the main text, the pore morphology exhibits density-dependent shape and distortions from ideal cylindrical structures, as illustrated in Fig. 13. To address this non-ideal geometry, we analyzed the radial distribution of Si and O atoms relative

to the central simulation axis. The effective pore radius was estimated as the distance at which the density of Si and O atoms plateaus, indicating the boundary between pore space and silica framework. Whereas our generation has been performed only for the SiO_2 composition, NBOs are of high relevance to applications of mesoporous silica, where they can become silanol groups in the presence of hydrogen.

Atomic volume: We employed Voronoi analysis in OVITO[49] to compute the atomic volumes of CuZr structures. Particle radii of 1.35 Å for Cu and 1.55 Å for Zr atoms are applied for analysis, with relative face area threshold set to 1%.

Voronoi indices: Voronoi indices for CuZr structures were computed using OVITO[49] for polydisperse Voronoi tessellation with particle radii of 1.35 Å for Cu and 1.55 Å for Zr atoms. Surface atoms occupying less than 1% of the total surface area were excluded to minimize effect from experimental measurement and structural reconstruction errors. Voronoi indices are represented using reduced Schlaefli notation, $\langle n_3, n_4, n_5, n_6 \rangle$, where n_i denotes the number of polyhedral faces containing i edges. The polyhedral face distribution quantifies the frequency of faces with varying edge numbers.

A.6 Information entropy and QUESTS method

Representation: the representation of atomic environments was computed as described in our previous work.[27] In summary, a number of k=32 neighbors was used to represent the atomic environment, with a cutoff of 5 Å. No distinction was made between element types, which is implicitly recovered based on the coordination environments, and thus captured in the information entropy of the system.

Information entropy: the information entropy of descriptor distributions was computed as described before.[27] Given a set of feature vectors $\{X\}$, their information entropy is computed as follows:

$$\mathcal{H}(\{\mathbf{X}\}) = -\frac{1}{n} \sum_{i=1}^{n} \log \left[\frac{1}{n} \sum_{i=1}^{n} K_h(\mathbf{X}_i, \mathbf{X}_j) \right],\tag{7}$$

where our choice of the kernel K_h is the Gaussian kernel,

$$K_h(\mathbf{X}_i, \mathbf{X}_j) = \exp\left(\frac{-||\mathbf{X}_i - \mathbf{X}_j||^2}{2h^2}\right),\tag{8}$$

with a constant bandwidth $h = 0.015 \text{ Å}^{-1}$, as studied before.[27]

We define the differential entropy $\delta \mathcal{H}$ of a data point Y with respect to a reference dataset $\{X\}$ as

$$\delta \mathcal{H}(\mathbf{Y}|\{\mathbf{X}\}) = -\log \left[\sum_{i=1}^{n} K_h(\mathbf{Y}, \mathbf{X}_i) \right]. \tag{9}$$

Overlap: From the definition of $\delta\mathcal{H}$, the overlap between two discrete distributions of atomic environments $\{\mathbf{Y}\}$ and $\{\mathbf{X}\}$ is the fraction of environments $\mathbf{Y}_i \in \{\mathbf{Y}\}$ with $\delta\mathcal{H}(\mathbf{Y}_i|\{\mathbf{X}\}) \leq 0$. Therefore, a zero overlap between two distributions indicates that the distributions have disjoint supports, whereas a 100% overlap indicates that the two distributions have the same support, even when the point-wise probabilities are different. Within the context of generative models described in the main text, a high overlap indicates that the generated structures follow the same distribution of atomic environments, while a low overlap suggests that the structures do not share the same environments.

Units: Throughout this work, the natural logarithm was used for the entropy in Eq. (7), which scales the information to natural units (nats).

A.7 Sensitivity of macroscopic properties to outliers

While structures generated with the denoising process without extra noise can capture peak positions and demonstrate reasonable agreement with simulated partial PDFs (Fig. 15a), they fail to reproduce the macroscopic properties of amorphous silica. Macroscopic properties such as mechanical properties and stress-strain behavior are quite sensitive to structural fidelity and are thus excellent tests to the quality of generated structures. Nonphysical atomic environments resulting from imperfect denoising can produce unrealistic bond lengths and angles that remain nearly undetected at short- and medium-range order (PDF, bond angle distributions, and ring size distributions such as Fig. 2 of the main text). However, even a single structurally invalid atom can result in major deviations in macroscopic properties, severely limiting the utility of generated structures for subsequent investigations.

Figure 16 compares stress-strain curves for structures generated under identical conditions using four methods: without extra noise, with extra noise, and each condition followed by a short MD equilibration (refinement). Despite very similar partial PDF performance across all samples (Fig. 15), their mechanical properties differ dramatically. Structures generated without extra noise fail to exhibit realistic fracture behavior and yield strength compared to simulated references even after a post-generation MD equilibration to reduce the occurence of outliers. Adding substantial extra noise during the denoising stage produces much more reasonable stress-strain curves, with brittle fracture characteristics. Both non-refined approaches exhibit residual stress due to fixed density constraints during denoising. In contrast, structures generated with extra noise and MD refinement result in similar stress-strain curves given the expected variability of ductility across different samples, reproducing both yield strength and elastic modulus.

Quantitative analysis of the elastic properties further demonstrates this sensitivity. For all a-SiO₂ structures, we computed the average bulk modulus (K), shear modulus (G), Young's modulus (E), and Poisson's ratio (ν) across simulated and generated results. Table 1 summarizes these results. For the reference simulated structures, we obtain K=61.2 GPa, G=37.1 GPa, E=92.4GPa, and $\nu = 0.247$, compatible with the mechanical properties for this interatomic potential.[51] Generated structures without extra noise, despite exhibiting somewhat reasonable PDFs (Fig. 15), produce severely aberrant values: K=166.7 GPa, G=2.9 GPa, E=8.7 GPa, and $\nu=0.491$. The post-generation MD refinement corrects the elastic properties of generated structures, resulting in values of K=63.2 GPa, G=37.5 GPa, E=93.9 GPa, and $\nu=0.256$. Interestingly, although these values are close to simulated reference, the failure of exhibiting fracture behavior (Fig. 16) demonstrates that multiple validation approaches are essential for assessing generated structure fidelity, and underscores the critical importance of adding extra noise during the generation process. Adding extra noise during denoising significantly improves the properties of generated structures, but still exist small discrepancies, leading to K=67.0 GPa, G=32.6 GPa, E=83.6GPa, $\nu = 0.287$. Finally, after a short MD equilibration, structures generated with extra noise achieve excellent agreement with simulated results: K = 60.6 GPa, G = 37.8 GPa, E = 93.8 GPa, and $\nu = 0.241$. These findings underscore that mechanical properties provide a stringent test of structural fidelity, revealing deficiencies that remain hidden in conventional structural analyses.

A.8 Overlap as novelty metric of amorphous structures

In the main text, we introduced the use of information theoretical quantities to measure the "novelty" of newly sampled amorphous configurations. In generative models for molecules, for instance, novel data points are molecular graphs distinct from the training set,[52, 53] and thus can be trivially computed by comparing the molecular graphs. In generative models for inorganic crystalline materials, novelty is defined by comparing the space groups, atomic basis, compositions, and other discrete quantities that uniquely define a crystal.[5, 6] Whereas this is not as trivial as in the case of molecules and can lead to deviations in the assignment of "novelty" of a crystal,[54] it still allows for quantitative comparisons of novelty in crystal structure generation.

However, in contrast to crystalline structures or molecular graphs, novelty is challenging to define in the amorphous space. Even amorphous materials obtained from identical simulation conditions will not be identical unless they use exactly the same initial configuration and random seeds, leading to an ill-defined concept of "novelty" in the amorphous materials space. Thus, analyzing whether a generative model can approximate the data distribution and produce new samples requires measuring probabilities in the data domain. To do that, we quantified the similarity between amorphous

configurations using an information-theoretical method. Specifically, we computed a probability distribution over environments for a-SiO₂ structures using information theory, [27] and evaluated the similarity between distributions using an information-theoretical overlap score. This comparison between distributions resembles the Fréchet Inception Distance in generative adversarial networks, [28] which accounts for novelty in terms of distribution statistics rather than point-wise metrics. As discussed in the main text, we propose that the overlap of distributions of atomic environments can provide insights on the "novelty" of amorphous structures. These overlaps are computed per pair of structures, and thus can be depicted as a non-symmetrical matrix, as $P(A|B) \neq P(B|A)$. Figure 17 shows how these matrices look like when six a-SiO₂ structures are compared against other six structures. The diagonal of Figs. 17a,b indicate that the overlap of a structure with itself is 100%, as expected. Off-diagonal elements show that simulated a-SiO₂ samples ("Sim | Sim") exhibited, on average, 75% overlap with each other despite deriving from the same processing condition, with narrow variability between the overlaps (Fig. 17b). This result shows that, for 3000-atom systems sampled in our simulations, about 75% of the environments are consistently sampled, and 25% environments are usually considered "new" even among deterministic simulations. Similarly, generated samples showed 70% overlap with each other ("Gen | Gen" in Fig. 2e). This overlap indicates that the generative model is capturing both the information relevant to describe an amorphous state and the approximate expected novelty in the amorphous states.

To test for mode collapse, we also computed the overlap between the distribution of generated environments and simulated ones ("Gen | Sim"), shown in Fig. 17c. If the model displayed a mode collapse and only learned to reproduce the training data, this overlap would be close to 100%. On the other hand, a model that consistently produces unphysical data points would exhibit very low overlap with the simulated structures even when the "Gen | Gen" overlap is high. Finally, a model that perfectly reproduces the distribution of simulated data should exhibit an overlap close to 75%. In our case, the overlap between the generated and simulated structures was approximately 70%, thus demonstrating that our generated a-SiO₂ samples maintain excellent structural fidelity while preserving appropriate configurational diversity. Figure 2e of the main text summarizes these findings by comparing the distribution of overlaps.

A.9 Information theoretical analysis of conditionally generated silica glasses

Given the known relationships between free volume, configurational entropy, and cooling rates in glasses,[30, 55] we used our information entropy approach as a surrogate of configurational entropy for glassy materials.[27] As explained in Section A.8, this method computes an information entropy of distributions of atomic environments and thus translates atomistic systems into probability distributions. Within glassy materials, even at constant density, input cooling rates should be able to steer the final generated structure to different states containing different distributions of atomistic environments, and thus different values of configurational entropy. [56] In particular, rapid cooling produces high-energy states with incomplete structural relaxation, leading to glasses with more diverse local atomic motifs compared to slowly cooled counterparts, which is quantified by a higher information entropy. The relationship between information entropy averaged across six generated structures, cooling rate, and density of generated structures is shown in Figure 3a of the main text. The results confirm that higher information entropy is achieved for higher cooling rates, even at constant densities, validating the model's ability to produce conditional generation of cooling rates. For generated a-SiO₂ samples at the right densities (Fig. 20), average information entropy values were 6.87, 6.70, 6.53, and 6.43 nats for cooling rates of 10^2 , 10^1 , 10^0 , and 10^{-1} K/ps, respectively. These results are in excellent agreement with the average information entropy values of 7.01, 6.76, 6.58, and 6.43 nats for the reference simulated structures.

A.10 Effect of cooling rate on simulated and generated properties

MD simulations are widely employed to study glass structure and properties, yet their inherent limitation to short timescales imposes unrealistically fast cooling rates on simulated structures.[29] This constraint makes it challenging to provide a one-to-one comparison with experimental glasses processed under typical laboratory conditions. In simulations of a-SiO₂, high cooling rates typically introduce spurious structural defects, though bulk structural characteristics remain relatively stable despite thermal history. Conversely, properties such as density and thermal expansion coefficients exhibit strong dependence on cooling rate. Since our denoising approach generates structures at fixed density (i.e., number of particles and volume), we developed a regression model that enables

us to specify the density as a function of cooling rate (Fig. 14). Within the range of cooling rates of interest, simulated amorphous silica densities are reasonably linear with the logarithm of the cooling rate, simplifying the determination of initial densities for denoising. Although the density of silica does not follow a linear trend across a wider range of cooling rates, [24] they are valid within our study and simplify the analysis of model generalization. The post-denoising MD equilibration employs both NVT and NPT ensembles, but the latter typically only increases the density by around 0.01 g/cm³.

Figure 20 illustrates how the information entropy changes at the final steps of conditional denoising across different cooling rates. Following the methodology described in the Methods section, entropy values remain approximately 8.0 nats before step 2200, reflecting the maximum possible entropy with a system with 3000 atoms ($\ln 3000 = 8.0$).[27] As the noise is progressively lowered due to the noise schedule, the information entropy of the system decreases and structures develop local ordering, with slower cooling rates leading to lower final entropy values, all consistent with simulated trends.

Validation of structures as a function of cooling rate employs mean Si–O–Si bond angles and mechanical properties, including Young's modulus (Fig. 6), bulk modulus (Fig. 7), shear modulus (Fig. 8) and Poisson's ratio (Fig. 9) across different cooling rates. Generated structures accurately capture Si–O–Si bond angle trends with respect to cooling rate, demonstrating excellent agreement in both interpolation regions (0.3, 0.5, 3, 5, 30, and 50 K/ps) and "generalization" domains, i.e., targeted cooling rates outside of the interval of $[10^{-1}, 10^2]$ K/ps. Mechanical properties, which exhibit high sensitivity to structural accuracy (see Appendix A.7 above), remain physically consistent across all cooling rates for generated samples. On the other hand, outliers, such as structures generated without extra noise (Fig. 15a), exhibit aberrant elastic properties (Figs 6–9), despite having similar partial PDFs compare to simulated samples. This confirms that our conditioned structure generation produces structurally valid amorphous configurations across a range of cooling rates.

A.11 Size and shape effects in fracture of amorphous silica

Despite known variability in simulated plastic properties such as strength, generated structures consistently lead to stress-strain curves with good approximation of the elastic regime and ultimate tensile strength of the simulated material. This is true even at large length scales, when cracks are strongly influenced by structural outliers. Moreover, larger generated structures exhibit an increasingly brittle behavior, with substantial reduction in ductility from structures with 3,000 and 30,000 atoms, consistent with past experimental results.[57] Figure 12 shows the stress-strain curves, averaged from six independent tests, for the fracture of generated cubes ($L_z/L_x = 1.0$) with system sizes ranging from 3,000 to 30,000 atoms. Given the larger aspect ratio of the generated samples compared to cubic structures in Fig. 12, brittle fracture is observed at lower strains in Fig. 4a. This agrees with previous studies stating that the length along the deformation axis is the most important factor in minimize boundary effects and provide adequate volume for realistic crack propagation.[57]

Table 1: Average mechanical properties (bulk modulus K, shear modulus G, Young's modulus E, and Poisson's ratio ν) of a-SiO $_2$ structures for different structure generation methods. All values are in GPa except for ν (dimensionless). Errors correspond to the standard deviation of properties across six generated and simulated samples. Samples generated with extra noise and MD refinement exhibit the exact distribution of elastic properties of the simulated reference configurations.

Method	K (GPa)	G (GPa)	E (GPa)	ν
Generated (no extra noise)	166.7 ± 8.2	2.9 ± 0.3	8.7 ± 1.1	0.491 ± 0.001
Generated (no extra noise) + MD refinement	63.2 ± 2.1	37.5 ± 2.3	93.9 ± 4.5	0.256 ± 0.016
Generated (with extra noise)	67.0 ± 8.8	32.6 ± 4.3	83.6 ± 8.7	0.287 ± 0.042
Generated (with extra noise) + MD refinement	60.6 ± 3.1	37.8 ± 0.8	93.8 ± 1.4	0.241 ± 0.014
Simulated reference	61.2 ± 2.9	37.1 ± 3.0	92.4 ± 5.6	0.247 ± 0.025

B Figures

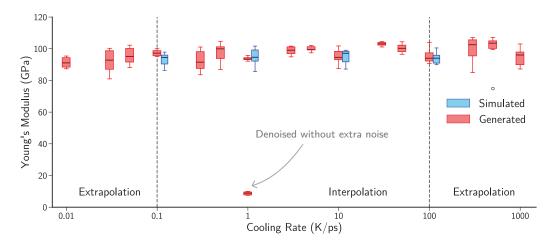


Figure 6: Young's modulus of simulated and generated a-SiO₂ samples across different cooling rates. The moduli are highly sensitive to outliers, yet the denoiser is able to generate structures with accurate Young's moduli compared to simulations, even outside the training domain $(10^{-2} \text{ and } 10^3 \text{ K/ps})$. The results of structures generated without adding extra noise demonstrate that high fidelity structures are necessary to reproduce reasonable elastic property.

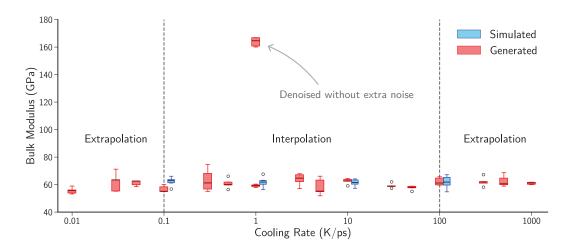


Figure 7: Bulk modulus of simulated and generated a-SiO₂ samples across different cooling rates. The moduli are highly sensitive to outliers, yet the denoiser is able to generate structures with accurate bulk moduli compared to simulations, even outside the training domain (10^{-2} and 10^{3} K/ps). The results of structures generated without adding extra noise demonstrate that high fidelity structures are necessary to reproduce elastic properties.

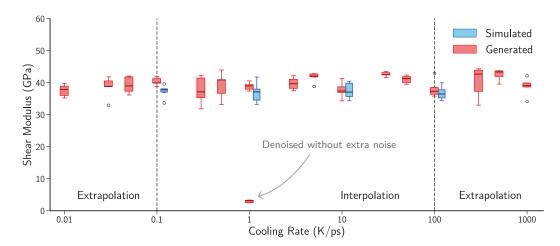


Figure 8: Shear modulus of simulated and generated a-SiO₂ samples across different cooling rates. The moduli are highly sensitive to outliers, yet the denoiser is able to generate structures with accurate shear moduli compared to simulations, even outside the training domain (10^{-2} and 10^{3} K/ps). The results of structures generated without adding extra noise demonstrate that high fidelity structures are necessary to reproduce elastic properties.

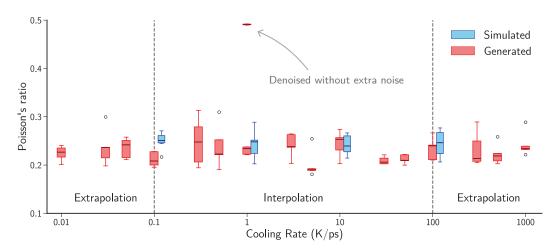


Figure 9: Poisson's ratio of simulated and generated a-SiO₂ samples across different cooling rates. The elastic properties are highly sensitive to outliers, yet the denoiser is able to generate structures with accurate Poisson's ratios compared to simulations, even outside the training domain $(10^{-2}$ and 10^3 K/ps). The results of structures generated without adding extra noise demonstrate that high fidelity structures are necessary to reproduce elastic properties.

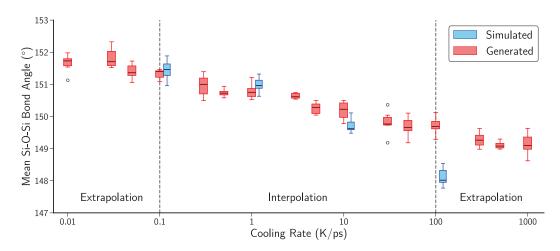


Figure 10: Mean Si–O–Si bond angle of a-SiO $_2$ structures as a function of cooling rate for simulated and generated samples. Each data point represents the average of six independent structures. The generated samples accurately reproduce the cooling rate dependence observed in simulated structures, with correct trends maintained even for cooling rates outside the training domain $(10^{-2}$ and 10^3 K/ps).

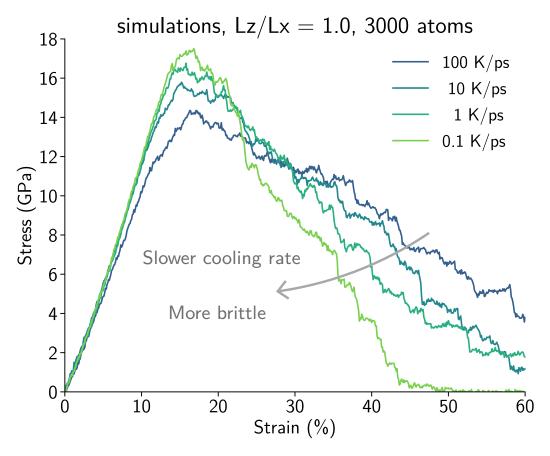


Figure 11: Stress-strain curves from simulated a-SiO₂ structures under different cooling rates show that slower cooling rate leads to a more brittle behavior. The initial structure was a cube with a total of 3000 atoms, and the fracture simulation was performed as described in the Methods. These results show that the ductility of the glassy system is high at this small scales, in contrast with results at large scales.

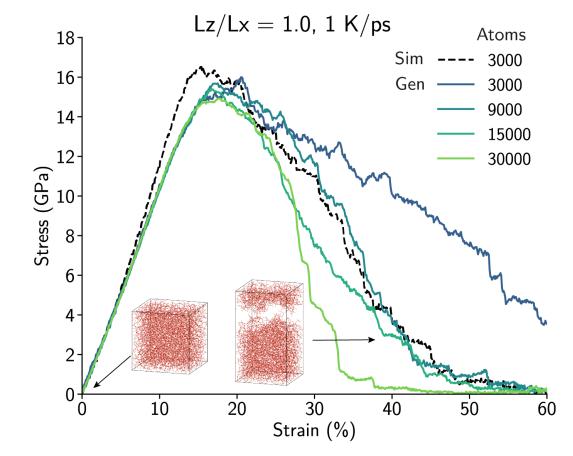


Figure 12: Stress-strain curves show that generated a-SiO $_2$ structures (solid lines) recover the expected trends of elastic and plastic deformation behavior across length scales, with more ductile fracture at smaller length scales. The stress-strain curve of a simulated a-SiO $_2$ structure with 3000 atoms is shown with a dashed line. The initial simulation box was a cube ($L_z/L_x=1$).

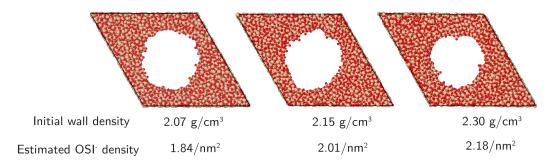


Figure 13: Mesoporous a-SiO2 structures generated by the denoiser model at varying initial wall densities. All structures are generated from the same initial geometry and processing conditions. Higher wall densities result in correspondingly denser non-bonding oxygen groups on pore surface.

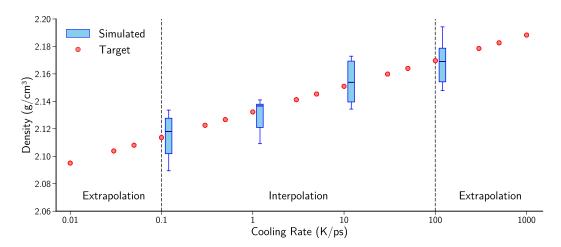


Figure 14: Densities of simulated a- SiO_2 structures and target densities across different cooling rates. Target densities are determined based on a linear regression model valid with simulated densities and extended to 0.01 and 1000 K/ps.

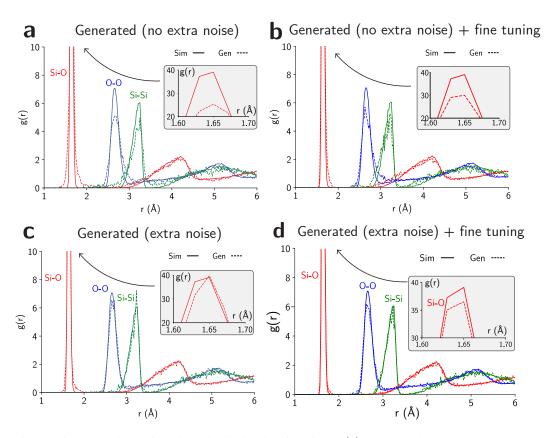


Figure 15: Comparison of partial pair distribution functions g(r) between simulated and generated a-SiO₂ for cases: **a**, generated without adding extra noise, **b**, generated without adding extra noise, followed by MD refinement, **c**, generated with extra noise, and **d**, generated with extra noise and MD refinement). All cases demonstrate that generated configurations reproduce good structures according to structural quality metrics. Adding extra noise and fine tuning with MD removes most nonphysical behaviors (such as short Si–O pairs) and give accurate peaks for PDFs without intensity decreases.

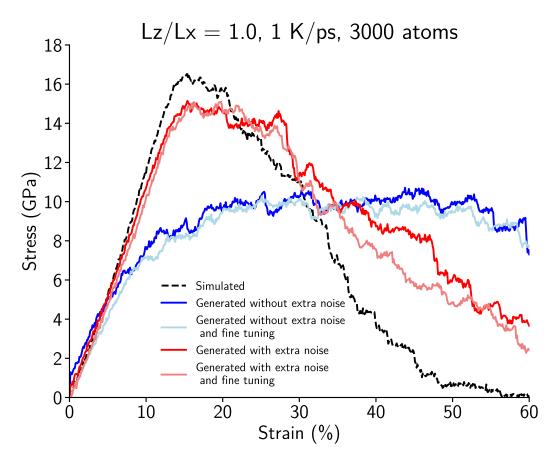


Figure 16: Stress-strain curves of simulated a-SiO $_2$ structure (dashed line) and generated structures (solid lines). Generated structures with extra noise and fine tuning show the best agreement compared to simulated results in terms of strength and ductility. Generated structures with extra noise show reasonable agreement compared to simulated result, but overestimate the stress at low strain and show more ductility. Generated structures without extra noise, even with fine tuning process, fail to reproduce the stress-strain behavior, showing a complete ductile behavior and underestimate the ultimate strength. All atomic structures were cube (L_z/L_x = 1) with 3000 atoms, and generated with a target cooling rate of 1 K/ps.

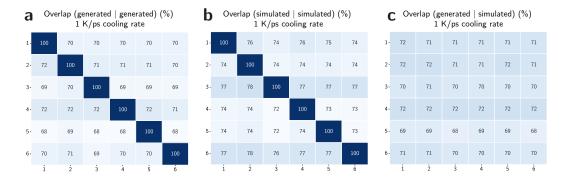


Figure 17: Overlap score matrices for $\bf a$, generated | generated, $\bf b$, simulated | simulated, and $\bf c$, generated | simulated a-SiO₂ structures. Each matrix computes the overlap scores between two series of six independent simulations or generations. High overlap values indicate strong structural similarities. The results show two main results: (1) the distribution of generated atomic environments is very similar to the ones obtained from MD simulations, showing that the model produces valid amorphous structures; and (2) the generative model does not simply "copy-and-paste" the known amorphous structures, and instead produces novel configurations, with a novelty rate nearly identical to the ones obtained from simulations. All structures are prepared with 3000 atoms and under 1 K/ps cooling rate.

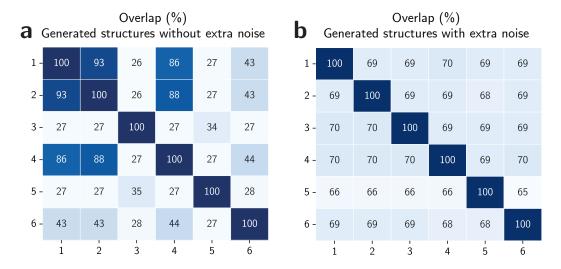


Figure 18: Overlap score matrices for generated a-SiO $_2$ structures $\bf a$, without and $\bf b$, with extra noise during the diffusion model. High overlap values indicate structural similarity, while low values indicate diversity. Without extra noise, structures show variable similarity. For instance, samples 1, 2, and 4 generated without noise are overly similar to each other, indicating a mode collapse into a single configuration. On the other hand, others differ substantially, indicating that the right distribution of environments is not captured. When extra noise is added, structures exhibit more consistent similarity across all samples.

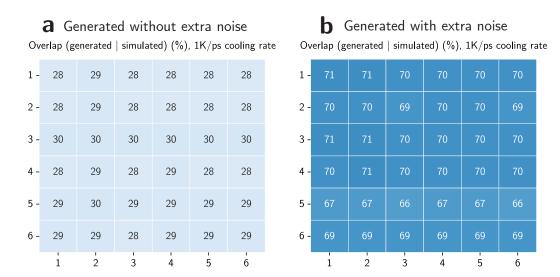


Figure 19: Overlap score matrices for generated against simulated a- SiO_2 structures \mathbf{a} , without and \mathbf{b} , with extra noise during generation. High overlap values indicate that the distribution of environments is similar. Generated samples without extra noise are mostly dissimilar to the simulated samples. On the other hand, the samples generated with extra noise are nearly as similar to simulated ones as simulated samples are to themselves.

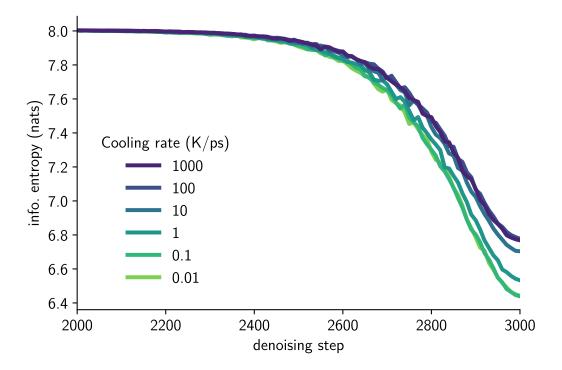


Figure 20: Information entropy of generated a-SiO $_2$ structures sampled along the denoising process with different target cooling rates. The denoising process was the same across the different conditions. The maximum information entropy that can be obtained for a system of this size with 3000 atoms is $\ln 3000 = 8.0$ nats.