

DiffSyn: a generative diffusion approach to materials synthesis planning

Received: 22 April 2025

Accepted: 18 December 2025

Published online: 02 February 2026

Elton Pan  ¹, Soonhyoung Kwon ², Sulin Liu ¹, Mingrou Xie  ², Alexander J. Hoffman  ¹, Yifei Duan ¹, Thorben Prein ³, Killian Sheriff  ¹, Yuriy Roman-Leshkov ², Manuel Moliner  ⁴, Rafael Gomez-Bombarelli  ¹ & Elsa A. Olivetti  ¹ 

 Check for updates

The synthesis of crystalline materials, such as zeolites, remains a notable challenge owing to a high-dimensional synthesis space, intricate structure–synthesis relationships and time-consuming experiments. Here, considering the ‘one-to-many’ relationship between structure and synthesis, we propose DiffSyn, a generative diffusion model trained on over 23,000 synthesis recipes that span 50 years of literature. DiffSyn generates probable synthesis routes conditioned on a desired zeolite structure and an organic template. DiffSyn achieves state-of-the-art performance by capturing the multi-modal nature of structure–synthesis relationships. We apply DiffSyn to differentiate among competing phases and generate optimal synthesis routes. As a proof of concept, we synthesize a UFI material using DiffSyn-generated synthesis routes. These routes, rationalized by density functional theory binding energies, resulted in the successful synthesis of a UFI material with a high Si/Al_{ICP} of 19.0, which is expected to improve thermal stability.

Materials discovery lays the foundation for modern technologies, from catalysis to electronics¹. Recent large-scale computational searches of chemical composition and structures^{2–5} have uncovered millions of potentially stable, synthesizable materials (‘what’ to synthesize)^{2,6–8}. However, finding viable synthesis routes remains a bottleneck in materials discovery (‘how’ to synthesize)^{9–12} because there are many synthesis parameters (composition, conditions and so on) that interact in complex ways. Moreover, the compute required for atomistic simulations scales poorly with system size, precluding accurate modeling of the underlying physical phenomena in complex materials (that is, thermodynamics and kinetics)¹³. Consequently, there is continued interest in machine learning (ML) approaches to directly learn from experimental synthesis data to predict materials synthesis parameters at lower computational cost^{10,11,14}.

Materials synthesis prediction presents a unique challenge for ML for several reasons. First, structure–synthesis relationships are ‘one-to-many’—that is, a single target structure may form through multiple

possible synthesis recipes. Second, the inverse relationship (synthesis–structure) is also one-to-many—that is, a single recipe may result in the formation of a mixture of products (competing phases) owing to the complex interplay of thermodynamic and kinetic pathways¹³. Capturing this phase competition is crucial to selectively synthesize single-phase materials instead of mixtures. Third, complex nonlinear interactions exist between synthesis parameters, such as temperature and time¹¹ (Supplementary Fig. 10), requiring approaches that model joint probabilities across multiple synthesis parameters. Predictions must capture relationships among variables to make trade-offs between parameters leveraging physical information about materials synthesis (for example, crystallization kinetics).

Previous ML approaches to predict synthesis have predominantly used regression approaches^{11,15–17}, which deterministically map a representation (for example, composition^{10,11}, structural features¹⁶, graphs¹⁵) of a material to its synthesis parameters. These approaches are limited because the deterministic mapping is

¹Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Department of Chemistry, Technische Universität München, Munich, Germany. ⁴Instituto de Tecnología Química, Universitat Politècnica de València-Consejo Superior de Investigaciones Científicas, Valencia, Spain.  e-mail: elsao@mit.edu

incompatible with the one-to-many nature of structure–synthesis relationships and assume independence between synthesis parameters^{18,19} (Supplementary Fig. 4).

These factors limit the predictive accuracy of regression approaches and motivate a shift to generative models, which can sample a complex distribution that accounts for the nonlinear interactions between parameters in high-dimensional synthesis space.

To address these challenges, we introduce a diffusion model for materials synthesis. Diffusion models are a powerful class of generative models that have been demonstrated to generate novel, high-quality images conditioned on text^{20,21}. Diffusion models can be guided at each step of the denoising process toward a specific objective (for example, target material)^{22,23}. Unlike generative adversarial networks that suffer from mode collapse²⁴, diffusion models can generate diverse outputs because they are trained to denoise data (for example, a synthesis route) that have been corrupted with noise. Other generative approaches, namely, variational autoencoders and normalizing flows, have limited expressivity due to one-step decoding and affine invertible layers, respectively. In contrast, the iterative denoising process renders diffusion models highly expressive, which enables high sample quality²³. This high expressivity may enable diffusion models to capture boundaries in synthesis space between competing phases.

We propose DiffSyn, a diffusion approach to materials synthesis prediction, and demonstrate it on zeolites, which are crystalline, microporous materials with applications in catalysis, adsorption and ion exchange. Zeolite synthesis is challenging owing to its high dimensionality (Fig. 1a), with numerous variables influencing the synthesis outcome (Supplementary Fig. 1). Moreover, multiple modes of valid synthesis routes exist for a given structure (Supplementary Fig. 5). Progress in zeolite synthesis has focused on trial-and-error experiments guided by domain heuristics²⁵. We leverage guided diffusion models for materials synthesis prediction, which show state-of-the-art performance compared with regression-based and other deep generative approaches. We show that the performance of DiffSyn arises from its ability to capture the one-to-many and multi-modal nature of structure–synthesis relationship in materials. We experimentally validate our approach by synthesizing the UFI zeolite based on DiffSyn-generated synthesis routes. We rationalize these routes using density functional theory (DFT) calculations of inorganic cations that guide UFI synthesis. Together, these results indicate that DiffSyn learns the underlying chemistry that influences synthesis outcomes implicitly from published synthesis recipes.

Results

DiffSyn framework for materials synthesis planning

Chemically guided diffusion model. DiffSyn leverages a chemically guided diffusion model for predicting materials synthesis routes with the target zeolite structure c_{zeo} and the organic structure-directing agent (OSDA) c_{OSDA} as inputs (Fig. 1b). An OSDA is an organic molecule that templates the zeolite's pores (Supplementary Fig. 1e), directing the synthesis toward the formation of a specific structure. Previous work has shown that the optimal OSDA to synthesize a given structure can be identified from atomistic simulations^{26–31}. The goal is to learn $p(x_{\text{comp}}, x_{\text{cond}} | c_{\text{zeo}}, c_{\text{OSDA}})$ to generate an ensemble of synthesis routes consisting of gel compositions $\{x_{\text{comp}}\}$ and synthesis conditions $\{x_{\text{cond}}\}$ given a target structure and OSDA, as shown in Supplementary Fig. 2 (also known as synthesis parameters, defined in Supplementary Table 1). Supplementary Fig. 4 shows an example of the predicted versus ground-truth synthesis parameters (green points). The prediction of synthesis parameters is an under-determined problem with multiple possible valid synthesis routes $\{x_{\text{comp}}, x_{\text{cond}}\}$ for each c_{zeo} and c_{OSDA} .

During training, the forward diffusion process (red arrow in Fig. 1b) adds Gaussian noise to x_{comp} and x_{cond} , progressively mapping them to a Gaussian distribution (noise). During inference, the reverse diffusion process (green arrow in Fig. 1b) starts from Gaussian noise and

iteratively denoises using a U-Net³² conditioned on chemical guidance ('Representation learning of materials' section) via classifier-free guidance (Fig. 1c). After T timesteps of denoising, the model generates synthesis routes for a desired structure. This denoising process can be seen in the improvement of generation metrics (for example, Wasserstein distance and COV-P, defined in 'Metrics' in Methods) throughout the reverse diffusion process (Supplementary Fig. 7). We train DiffSyn on the ZeoSyn dataset¹⁴, which consists of 23,961 synthesis recipes, 233 zeolite topologies and 921 OSDAs (Supplementary Section A).

Representation learning of materials. DiffSyn integrates a dual-encoder approach consisting of separate encoders (Enc_{zeo} and Enc_{OSDA}) for the zeolite structure and OSDA, respectively (Fig. 1b). We use two representations of the zeolite structure: invariant geometric features and an equivariant graph neural network (EGNN). The invariant geometric features are physical descriptors (for example, pore volume) calculated from the zeolite structure using the Zeo++ package. The EGNN encoder directly learns a representation from a graph of the zeolite crystal structure (Supplementary Section D). For the OSDA, we perform molecular geometry relaxation and calculate its physicochemical descriptors (for example, volume and shape) ('Zeolite and OSDA representations' in Methods).

Figure 1d shows that the respective encoders learn smooth and continuous latent spaces with respect to the properties of zeolites and OSDAs. A comprehensive set of properties plotted in embedding space can be found in Supplementary Figs. 8 and 9, indicating chemically meaningful representations of zeolites and OSDAs. These representations are concatenated before a fusion encoder ($\text{Enc}_{\text{fusion}}$) learns a joint representation. We refer to the joint representation as chemical guidance (Fig. 1b). Chemically meaningful representations are pivotal in steering the diffusion model to generate realistic synthesis routes for a desired materials structure. This representation enables DiffSyn to generate synthesis parameters that reflect synthesis routes unseen in training, which have been reported in the literature (Fig. 2e).

Influence of chemical guidance in diffusion model. Classifier-free guidance²² is a critical component of DiffSyn, where the chemical guidance steers the generation process by reweighting the unconditional score function with a conditional score function ('Chemically guided diffusion model' in Methods). We probe the influence of two key hyperparameters related to classifier-free guidance: the probability of the chemical guidance being omitted in score estimation during training (p_{uncond}) and the guidance strength that weighs the conditional score relative to the unconditional score during inference (w ; equation (7)).

Higher values of p_{uncond} and w amplify the conditional score, making the sampling process more dependent on the chemical guidance. This increased dependence on chemical guidance can result in over-constrained outputs, reducing diversity and potentially missing valid synthesis routes. Lower values of p_{uncond} and w reduce the influence of the chemical guidance, potentially decreasing specificity and generating synthesis recipes that are less tailored to the target structure, but improving the diversity of recipes. We find that the balance between diversity and quality of generated synthesis routes occurs at $p_{\text{uncond}} = 0.1$ and $w = 1.0$ (Supplementary Fig. 6).

Modeling structure–synthesis relationships

We evaluate DiffSyn against a suite of previously published approaches in materials synthesis planning. These baseline models fall into three categories: regression-based approaches (average minimum distance (AMD)¹⁶ and Bayesian neural networks (BNN)³³); classical generative models (Gaussian mixture model (GMM)³⁴); and deep generative models (conditional generative adversarial network (GAN)³⁵, normalizing flow (NF)³⁶ and variational autoencoder (VAE)^{11,19}). More information about these baselines is included in 'Baselines' in Methods. We compare approaches using Wasserstein distance (lower is better), which measures the distance between the generated and ground-truth

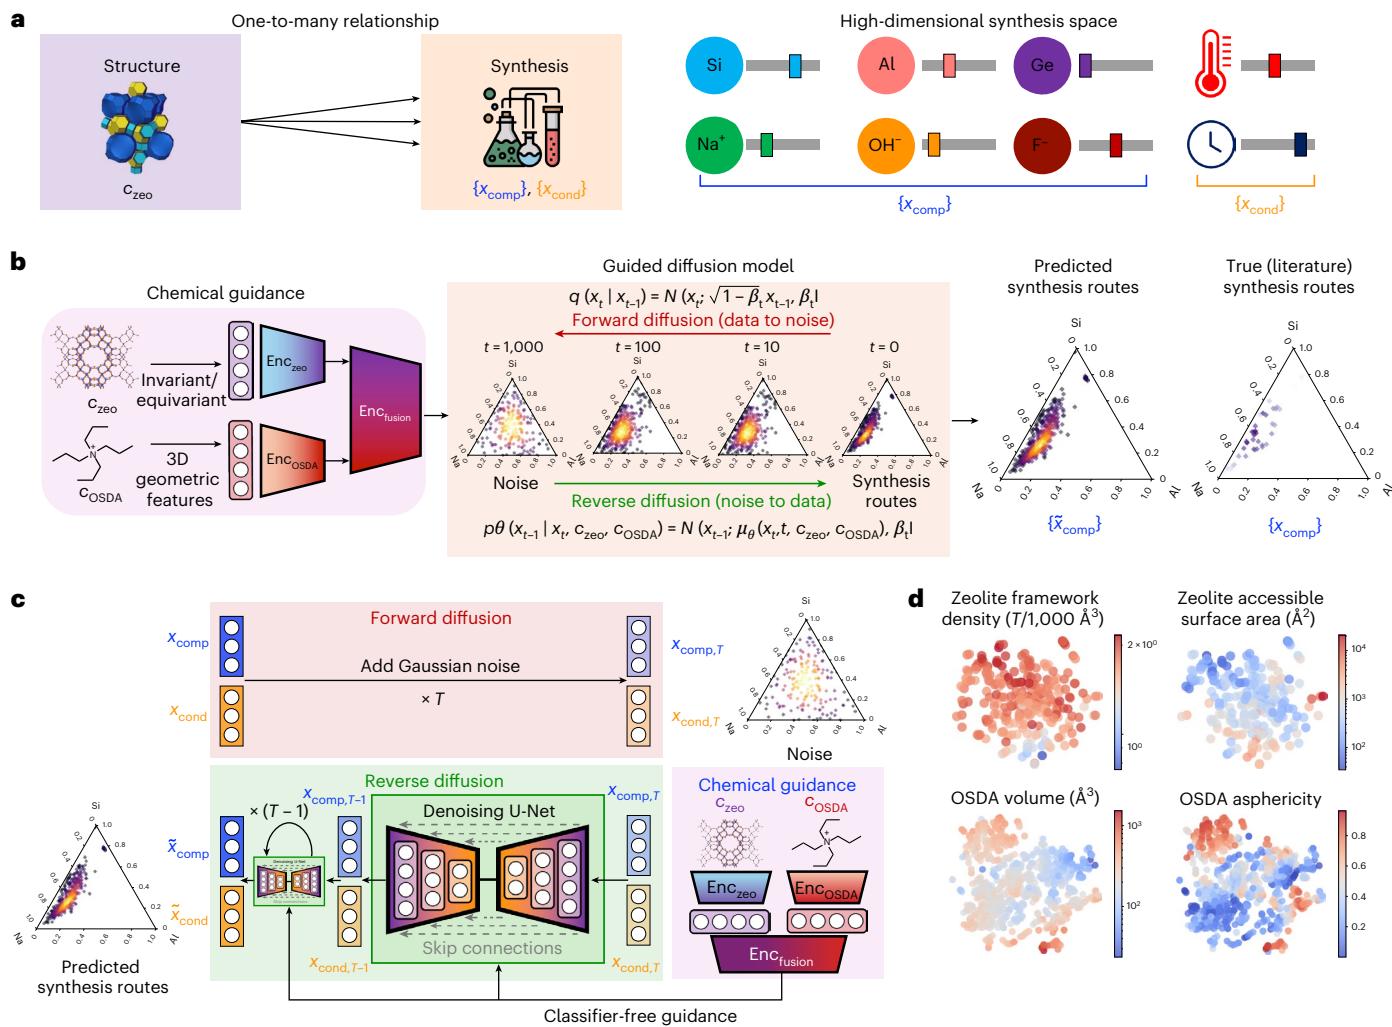


Fig. 1 | A diffusion approach to materials synthesis planning. **a**, Materials, such as zeolites, often have an one-to-many structure–synthesis relationship, where a structure c_{zeo} can be synthesized via multiple possible synthesis routes in a high-dimensional synthesis space (compositions $\{x_{\text{comp}}\}$ and conditions $\{x_{\text{cond}}\}$). **b**, Overview of DiffSyn: given a desired zeolite structure c_{zeo} and an OSDA c_{OSDA} , the model $p_{\theta}(x_{\text{comp}}, x_{\text{cond}} | c_{\text{zeo}}, c_{\text{OSDA}})$ generates an ensemble of synthesis routes $\{\tilde{x}_{\text{comp}}\}$ and $\{\tilde{x}_{\text{cond}}\}$ via reverse diffusion (green arrow) that matches ground-truth synthesis routes $\{x_{\text{comp}}\}$, accurately capturing the aforementioned one-to-many structure–synthesis relationship with which previous methods struggle. Note: only $\{\tilde{x}_{\text{comp}}\}$ is shown. **c**, Model architecture: two separate

encoders embed the zeolite material c_{zeo} and organic template c_{OSDA} before their latent representations are concatenated. The joint representation steers the generation process from noise ($x_{\text{comp}}, x_{\text{cond}}, T \sim \mathcal{N}(0, \mathbf{I})$) to realistic synthesis routes specific to the zeolite and OSDA via classifier-free guidance by learning $\mu_{\theta}(x_t, t, c_{\text{zeo}}, c_{\text{OSDA}}) = \frac{1}{\sqrt{\alpha_t}} (x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_{\theta}(x_t, t, c_{\text{zeo}}, c_{\text{OSDA}}))$ where $\epsilon_{\theta}(\cdot)$ is a conditional U-Net trained to predict the noise at time t ('Denoising diffusion probabilistic models' in Methods). Note that x_{comp} and x_{cond} are jointly noised and denoised. **d**, PCA of learned representations of zeolites and OSDAs. The color bar refers to a specific physical property (as shown in each title).

distributions of literature-reported synthesis parameters for unseen zeolite–OSDA systems. In addition, we propose the coverage metric COV-F1. The model should maximize both COV-P (precision) and COV-R (recall) simultaneously. Therefore, their harmonic mean (COV-F1) measures the degree of generated recipes being both realistic and diverse (ranges from 0 to 1, higher is better). Detailed explanations of these metrics are in 'Metrics' in Methods and Supplementary Section C.

Generative approaches better model structure–synthesis relationships. Wasserstein distances show that deep generative models such as GAN, NF, VAE and DiffSyn outperform the classical approaches, with DiffSyn outperforming the next best baseline (VAE) by over 25%. Classical generative approaches like GMM do not perform much better than a random baseline (Fig. 2a) while a probabilistic regression model (BNN) performs better than GMM.

The COV-F1 of 12 synthesis parameters are shown in Fig. 2b. The models perform better on synthesis parameters related to heteroatoms

(Si/Al, Al/P, Si/Ge, Si/B), cations (Na⁺/T, K⁺/T) and anions (F⁻/T, OH⁻/T), where T refers to the sum of all heteroatoms. However, they struggle to predict crystallization time, which could be attributed to anthropogenic factors; crystallization times are subject to human bias, where experimentalists test and report 'rounded' numbers³⁷. This bias results in the ground-truth time distribution peaking at specific intervals, hence rendering time prediction more challenging.

Deep generative approaches (VAE, NF and DiffSyn) outperform regression-based approaches (AMD and BNN). We hypothesize that generative models have superior performance owing to better recall (higher COV-R; Supplementary Fig. 3). Meanwhile, DiffSyn outperforms other deep generative models owing to higher precision (higher COV-P; Supplementary Fig. 3), where the diffusion model generates higher-quality synthesis routes. Interestingly, DiffSyn achieves the lowest mean absolute error for 10 out of 12 synthesis parameters (Fig. 2c), despite not being explicitly trained on the mean absolute error objective like the regression-based models.

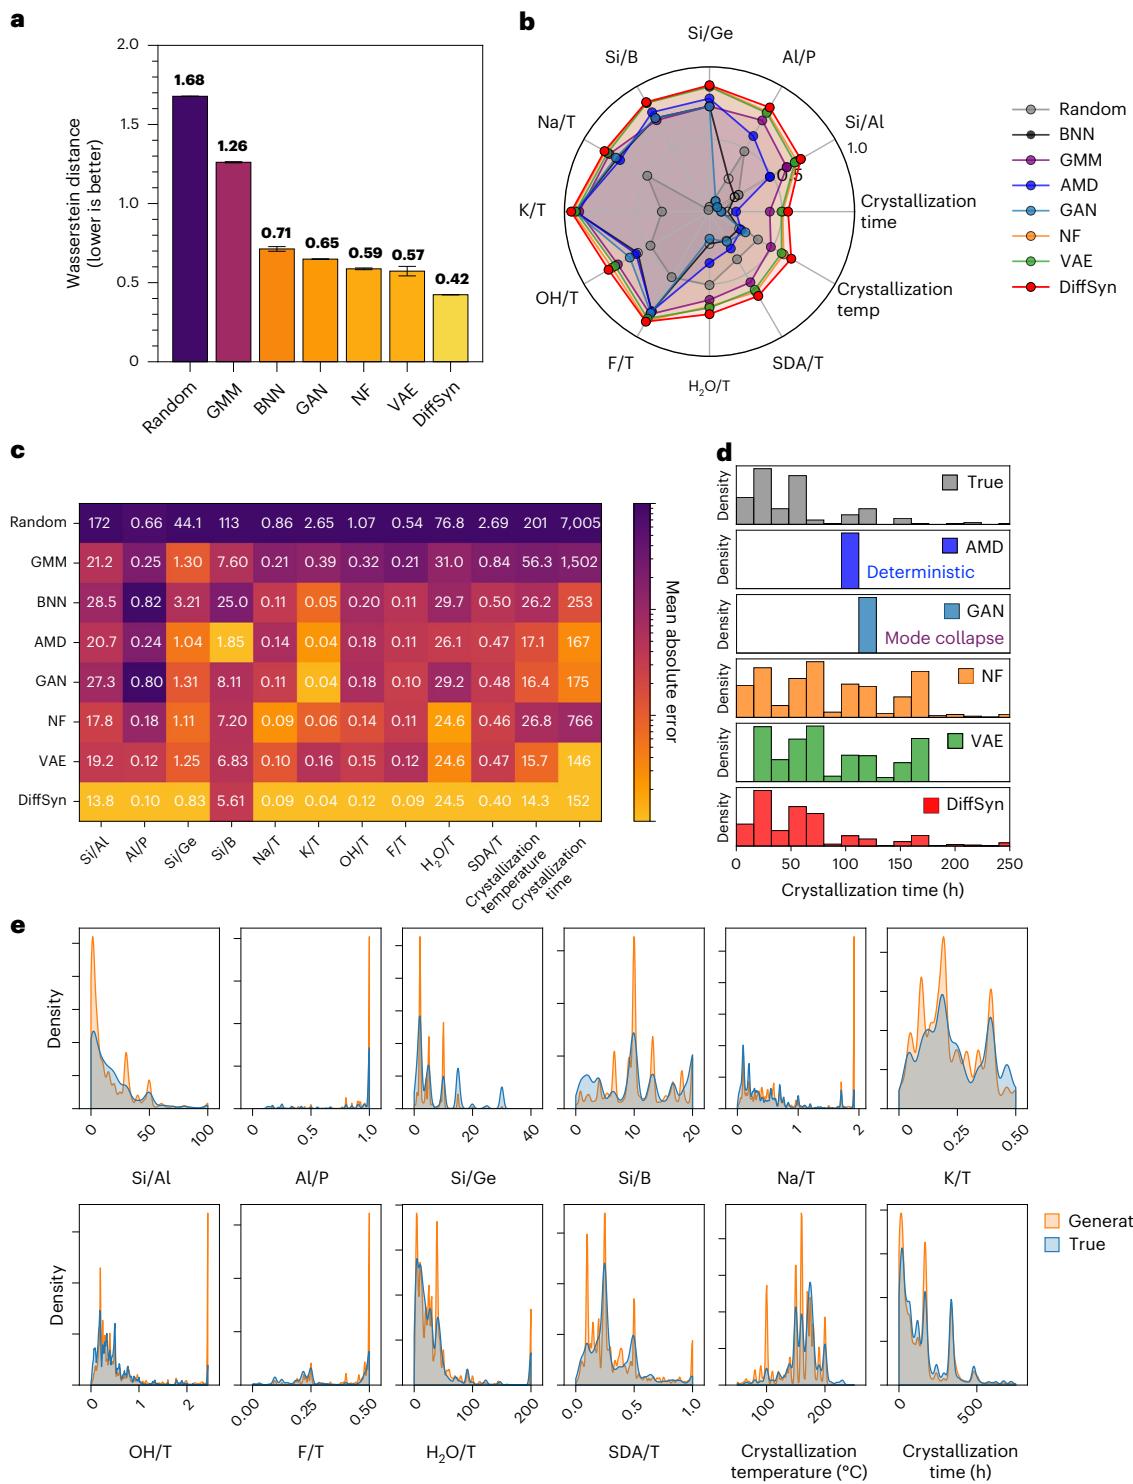


Fig. 2 | Performance in materials synthesis prediction task. **a**, Wasserstein distance (lower is better) between generated and literature synthesis routes. Metric is averaged across all test systems. Error bars correspond to standard deviation across three independent training runs. **b**, COV-F1 (higher is better, ranging from 0 to 1) of individual synthesis parameters (gel composition and conditions). **c**, Mean absolute error between the means of distribution of generated and literature synthesis parameters. **d**, Distributions of predicted and true synthesis parameters

across various different modeling approaches. AMD (dark blue) is regression-based (outputs deterministic, single-point prediction), while GAN (light blue) suffers from mode collapse. NF and VAE output distributions do not match the ground truth (gray). DiffSyn (red) accurately captures the true distribution (gray) of the synthesis parameters. **e**, DiffSyn-generated synthesis routes (orange) and distribution of synthesis routes reported in literature (blue). Synthesis parameters are aggregated across all unseen zeolite–OSDA systems.

Rationalizing the superior performance of generative approaches. For a given target material, there are a range of possible values for each synthesis parameter (for example, temperature) at which a material can be synthesized (that is, synthesis window). We show

the distribution of crystallization times from published syntheses (gray, Fig. 2d).

This distribution of synthesis parameters renders structure–synthesis relationship one-to-many instead of one-to-one.

Regression-based AMD is deterministic, and thus outputs a point prediction that comes from a weighted average of the distribution (Fig. 2d). Generative models like GAN also output a point prediction as they often suffer from mode collapse²⁴. Although generative models such as NF and VAE address mode collapse, they lack expressivity and fail to accurately capture the ground-truth literature distribution. In contrast, DiffSyn captures the ground-truth distribution. Furthermore, we compare the predicted and true joint distributions of multiple synthesis parameters for all of these approaches for the AEL zeolite (Supplementary Fig. 4). Only deep generative approaches (NF, VAE and DiffSyn) capture the ground-truth joint distribution of crystallization temperatures and times for the AEL structure, with DiffSyn most accurately capturing that joint distribution. DiffSyn captures most of the ground-truth points, including some outliers; however, DiffSyn fails to predict points in a minor mode (bottom right of guided diffusion panel in Supplementary Fig. 4), which are extreme outliers (for example, low crystallization temperature and long crystallization time).

The distribution of synthesis parameters is also multi-modal. We plot the principal components of all synthesis parameters for the aforementioned AEL structure in Supplementary Fig. 5, which shows that the true distribution has multiple modes—two in this case. Regression-based models (AMD, BNN) predict only one of the modes. GMM predicts synthesis routes that are far out of distribution. GAN suffers from mode collapse to one of the modes. NF and VAE capture both modes, but also generate a large number of false positives. This behavior arises from the low expressivity of VAEs and NFs, which use one-step decoding and affine invertible layers, respectively. In contrast, DiffSyn accurately predicts the true distribution because it generates high-quality and diverse outputs. Consequently, DiffSyn generates synthesis routes that overlap with unseen literature-reported synthesis parameters (Fig. 2e). For a discussion on diversity of generated samples, refer to Supplementary Section B.

Learning chemically meaningful relationships. We perform an unsupervised, hierarchical clustering of zeolite structures based on their learned representations (Fig. 3a), and observe distinct clusters according to their corresponding structural features (for example, number of channels, largest free sphere diameter). The clustering indicates that the zeolite encoder has learned to separate structurally distinct materials. Consequently, the chemical guidance (Fig. 1b), which requires learning good representations of the zeolite and OSDA, guides the generative process toward the desired target material.

Given that DiffSyn learns the joint distribution of multiple synthesis parameters, we inspect two synthesis parameters (crystallization temperature and time) in Supplementary Fig. 10 for two unseen zeolite–OSDA systems. An inverse relationship is observed between generated temperatures and times. This observation aligns well with the Arrhenius equation

$$k = A e^{-\frac{E_a}{RT}}, \quad (1)$$

where crystallization time (related to rate k) is inversely related to temperature T . A , E_a and R refer to Arrhenius constant, activation energy and gas constant, respectively. In addition, the generated H_2O/T (here, T refers to tetrahedral framework atoms or T-atoms, rather than temperature) and framework density (FD_{Si}) of the zeolite structure correlate

positively (Spearman's coefficient 0.673; Supplementary Fig. 11). This finding agrees with Villaescusa's rule³⁸, which states that denser phases (higher FD_{Si}) are favored at lower concentrations of T-atoms (higher H_2O/T), showing that DiffSyn has learned domain-specific rules in materials synthesis. The model predictions also follow the thermodynamics of zeolite formation, where the generated crystallization temperature and framework density (FD_{Si}) of the zeolite structure positively correlate (Spearman's coefficient 0.931; Supplementary Fig. 12). This finding agrees with the thermodynamic argument from Ostwald's rule of stages, which states that higher temperatures enable the synthesis to overcome the activation barrier to form more stable structures with higher framework densities^{14,39}.

Case studies

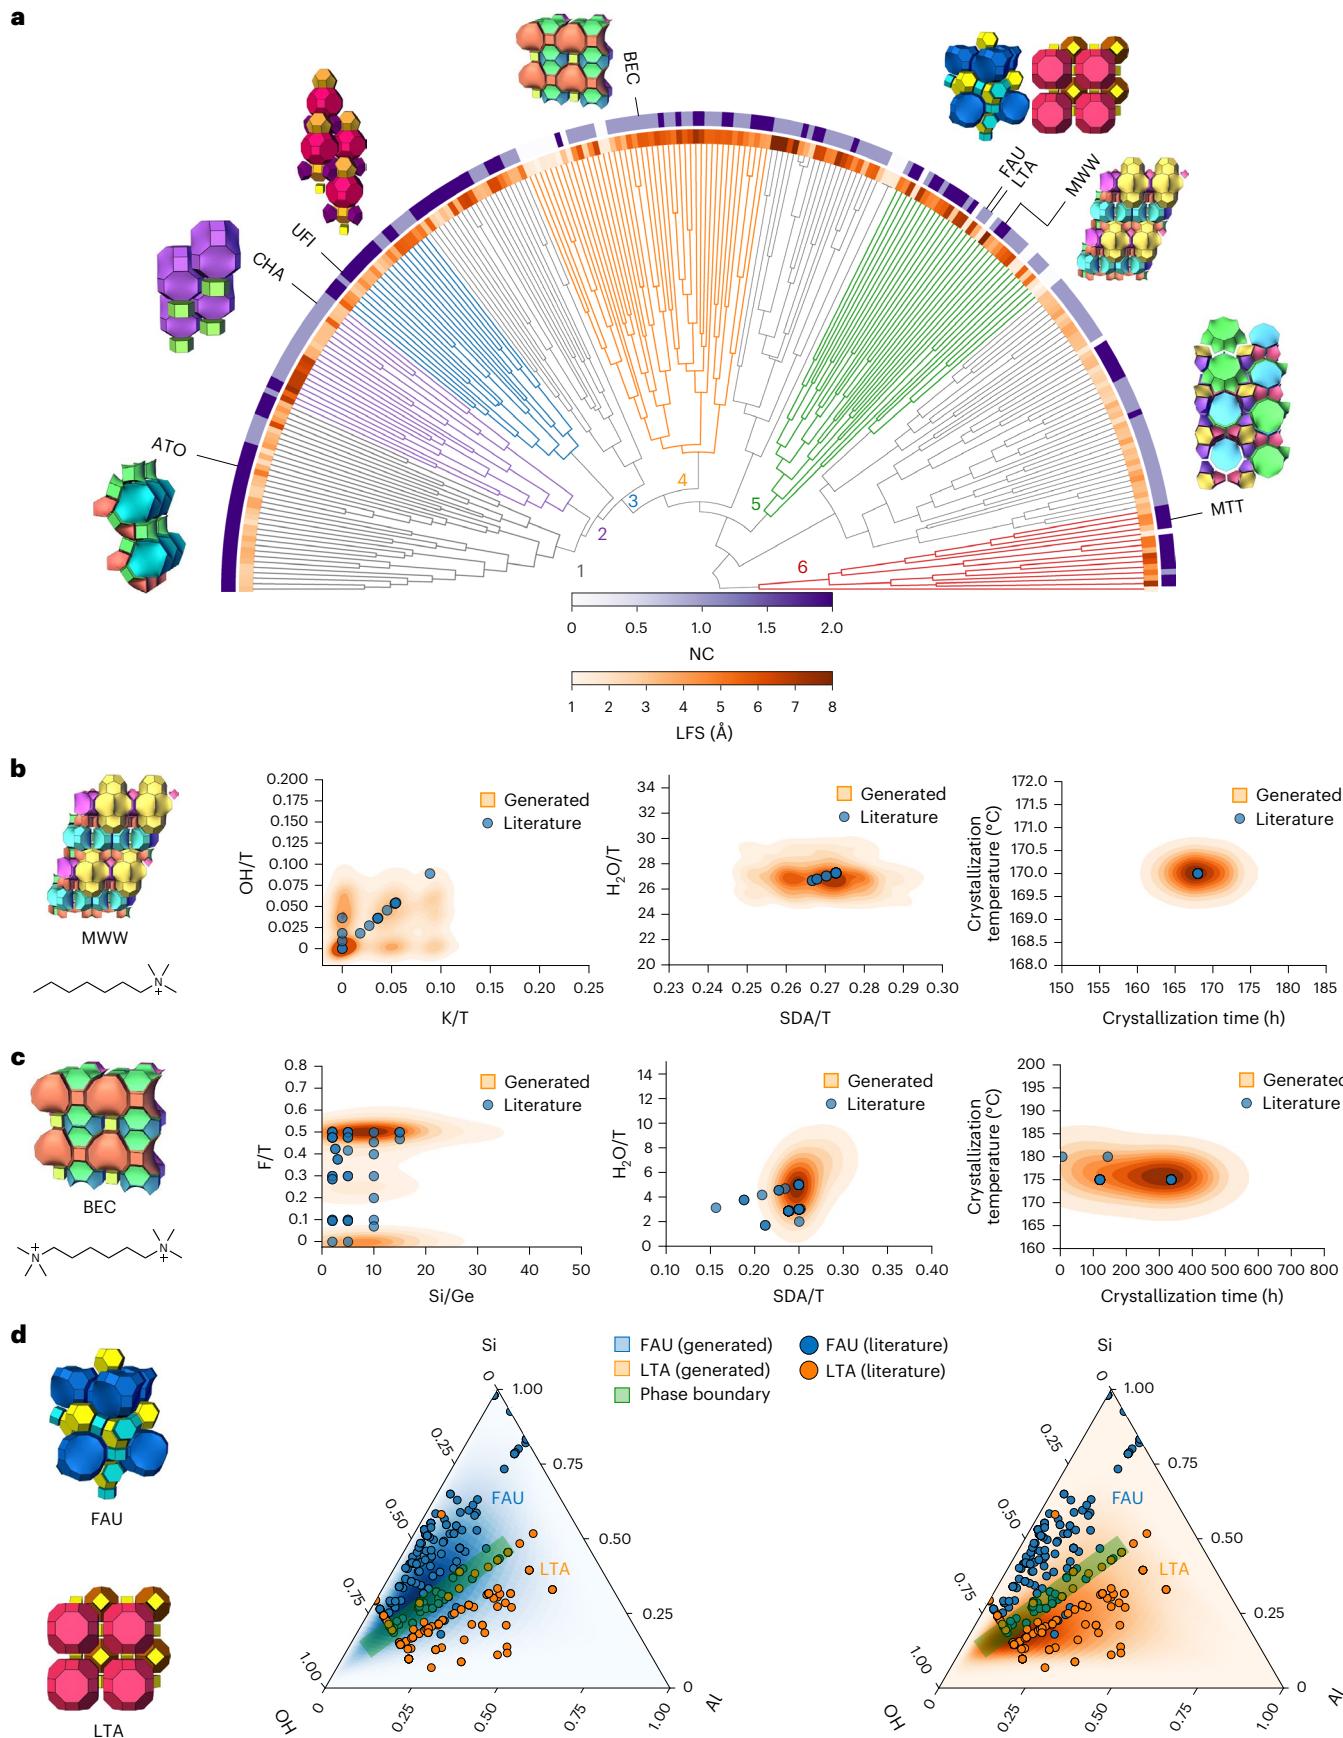
We compare DiffSyn-generated synthesis routes to literature-reported synthesis routes for diverse zeolite–OSDA systems that are synthetically interesting and industrially useful. The generated routes for these unseen systems (MWW, BEC and a pair of competing phases—FAU and LTA) provide evidence of DiffSyn learning meaningful synthesis–structure relationships (Fig. 3). Other systems (MTT and ATO) are analyzed in Supplementary Figs. 13 and 14, respectively.

We first consider the MWW phase, a unique two-dimensional structure with 10-membered rings and large cavities, with applications including isomerization⁴⁰ and aromatization⁴¹. The generated OH^-/T , K^+/T , H_2O/T , SDA/T , and crystallization temperature and time overlap substantially with ground-truth synthesis parameters (Fig. 3b). We also test the model on a significantly different structure: BEC, a large-pore zeolite. BEC has a three-dimensional pore topology with intersecting 12-membered ring channels, with applications in isomerization⁴² and epoxidation⁴³. DiffSyn-generated synthesis parameters closely agree with synthesis parameters reported in the literature, particularly Si/Ge , F^-/T , and crystallization temperature and time (Fig. 3c). This prediction aligns with reports that Ge and F^- stabilize the double four ring ($d4r$) composite building unit of the BEC structure during synthesis⁴⁴. This finding suggests that DiffSyn learns how particular heteroatoms or synthesis conditions favor the formation of specific building units within zeolites. However, the generated synthesis parameters do not always fully recall the ground truth for BEC synthesis. For example, the model fails to predict the full range of possible SDA/T values that have been identified in past recipes ($SDA/T = 0.15–0.25$; Fig. 3c).

Typically, materials synthesis aims to produce a single framework; however, if two phases during synthesis compete, some recipes result in two or more phases. Here we applied DiffSyn to predict OSDA-free synthesis routes for the FAU and LTA zeolites⁴⁵. Synthesis routes generated by DiffSyn align closely with literature-reported recipes (Fig. 3d). Notably, DiffSyn accurately predicts the phase boundary region (green) between FAU and LTA in OSDA-free conditions, delineating the synthesis space under which the competing phases form. This result shows that DiffSyn accurately captures not only the forward relationship (structure–synthesis) but also the decision boundaries of the inverse relationship (synthesis–structure), hence demonstrating its potential to enable phase-selective synthesis. Similarly, we demonstrate this capability for another pair of competing phases (ERI and KFI) in Supplementary Fig. 15. Taken together, these case studies illustrate DiffSyn's ability to generalize to a variety of zeolite frameworks and their corresponding chemistries.

Fig. 3 | Case studies on unseen materials systems. **a**, Hierarchical clustering of zeolite structures. Splits are obtained via agglomerative hierarchical clustering of learned zeolite representations. Each structure (with a three-letter code as its name) is colored by its number of channels (NC) in purple, and its largest free sphere (LFS) in orange. This leads to several distinct clusters: (1) high NC, low LFS; (2) low NC, low + high LFS; (3) high NC, low + high LFS; (4) low NC, high LFS; (5) low + high NC, high LFS; (6) high NC, high LFS. **b,c**, Generated synthesis routes for unseen materials systems. **b**, MWW structure (in cluster 5) with

N,N,N-trimethylhexan-1-aminium as the OSDA. **c**, BEC structure (in cluster 4) with pentane-1,5-diyl-bis(trimethylammonium) as the OSDA. Orange heatmaps refer to synthesis routes generated by DiffSyn, while blue points refer to literature-reported synthesis routes. **d**, Competing phases FAU and LTA (in cluster 5). Heatmaps refer to generated routes, while points refer to literature-reported synthesis routes. Notice that the model accurately predicts the phase boundary (green shaded region) between FAU and LTA.



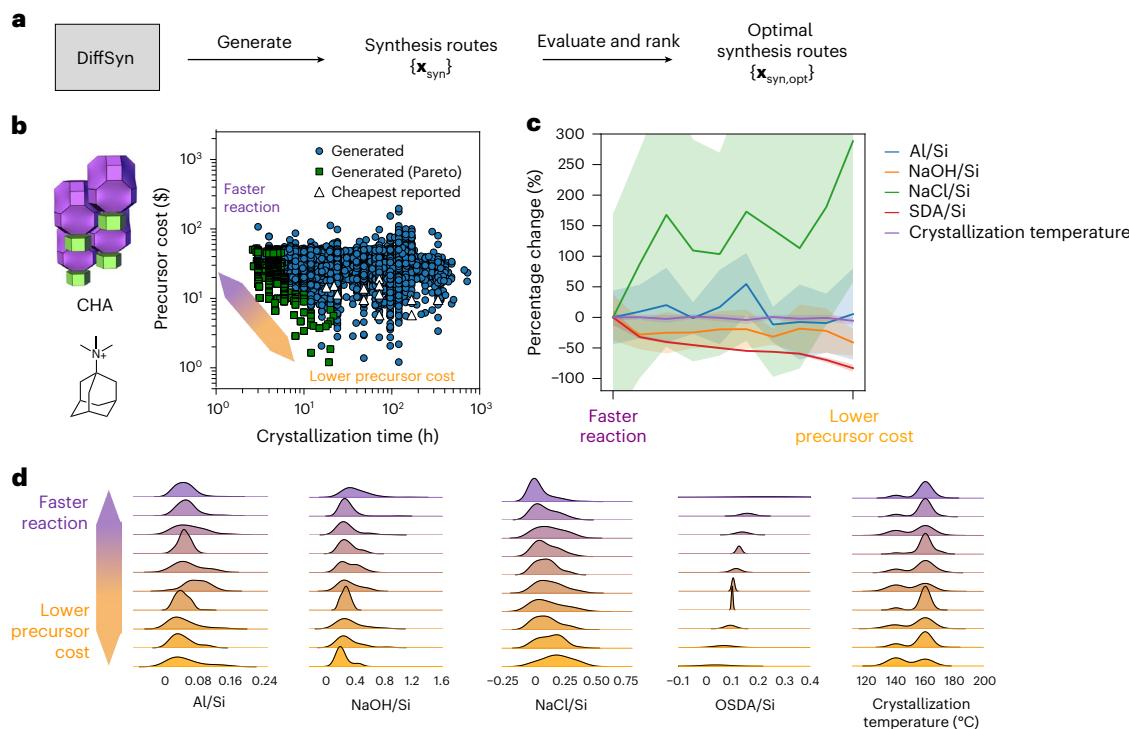


Fig. 4 | Generating optimal synthesis routes. **a**, DiffSyn generates synthesis routes $\{x_{\text{syn}}\}$, which are then evaluated based on some metric of interest (cost, crystallization time) and ranked to find an optimal synthesis route. **b**, Precursor cost and crystallization time of synthesis routes for the CHA structure (in cluster 2 of Fig. 3a) with TMAda as the OSDA. Circles refer to DiffSyn-generated routes and

triangles refer to 20 cheapest routes reported in literature. Notice that the Pareto front (green squares) trades off cost and time. **c**, Percentage change in generated synthesis parameters along the Pareto front. Solid lines and shaded regions refer to mean and standard deviation of the synthesis parameter, respectively. **d**, Distributions of generated synthesis parameters along the Pareto front.

Generating optimal synthesis routes

We use DiffSyn to generate more feasible synthesis routes (Fig. 4a), by evaluating and ranking the generated synthesis routes based on precursor cost and crystallization time (Supplementary Table 2) using methods reported in ref. 46. We apply this approach to the synthesis of CHA, using trimethyladamantyl ammonium (TMAda) as the OSDA⁴⁷, and generate an ensemble of synthesis routes with corresponding precursor costs and crystallization times (Fig. 4b). Among the generated routes, there are Pareto-optimal routes that trade off low precursor cost and fast synthesis time. These two objectives are optimal in different regions in the synthesis space—the OSDAs that produce CHA crystals faster tend to be more expensive (Supplementary Fig. 16). Some of these Pareto-optimal routes generated by DiffSyn have lower crystallization time and lower precursor cost compared with the 20 least expensive synthesis routes reported in the literature.

We inspect the evolution of synthesis parameters on the Pareto front. Notably, Al/Si and NaOH/Si remain relatively constant (Fig. 4c). In contrast, SDA/Si decreases as we prioritize lower cost over faster reaction as the OSDA typically drives precursor cost. The distributions of generated synthesis parameters also change along the Pareto front (Fig. 4d). This analysis reveals that increasing the crystallization temperature from 140 °C to 160 °C while simultaneously increasing NaOH/Si and OSDA/Si would favor faster reaction (Fig. 4d). This assessment shows how varying the joint distributions of synthesis parameters can accelerate synthesis (synthesizing desired structure in shorter time at higher temperature).

Experimental and DFT validation

We validate DiffSyn by experimentally synthesizing a UFI zeolite from recipes generated by our model with Kryptofix 222 (K222) as the OSDA⁴⁸ (Fig. 5). UFI has potential applications in industrially relevant reactions (for example, selective catalytic reduction of NO_x (ref. 49)). The

UFI-K222 system has not been reported in previous literature and is not present in the training dataset, and hence serves as a test of DiffSyn's out-of-distribution generalization.

Figure 5a shows the PCA of 1,000 DiffSyn-generated synthesis routes for UFI (orange), which constitutes a subspace of literature-reported synthesis routes for all zeolites (gray). We retrieve the k -nearest neighbors ($k = 5$) of generated UFI synthesis from the literature-reported synthesis routes (Supplementary Fig. 17). Among the retrieved syntheses, the top-two most similar frameworks (PAU and RHO) share a common *lta* composite building unit (CBU) with UFI (Supplementary Fig. 18). Beyond this, the remaining frameworks do not share any common CBUs with UFI. This observation is supported by a previous work⁵⁰ reporting that the majority of competing phases do not share any CBUs, suggesting that structure–synthesis relationships are complex and cannot be rationalized with building units alone.

We visualize individual synthesis parameters of generated routes in Fig. 5e (orange histograms), including gel compositions (Si/Al, Ge/Si, B/Si, Na⁺/Si, K⁺/Si, H₂O/Si, F⁻/Si) and reaction conditions (crystallization temperature and time). For inorganic cations, DiffSyn recommends high Na⁺/Si and low K⁺/Si for UFI synthesis (Fig. 5e). We rationalize this observation by calculating binding energies (ΔE_b) of inorganic cations (Na⁺ and K⁺) in building units of UFI (*wbc* and *rth*) using DFT as shown in Fig. 5b (details in 'DFT methods' in Methods). The calculations reveal that Na⁺ binds more strongly to *wbc* and *rth* ($\Delta E_b^{\text{sol}} = -70 \text{ kJ mol}^{-1}$ and $+13 \text{ kJ mol}^{-1}$) compared with K⁺ ($\Delta E_b^{\text{sol}} = -50 \text{ kJ mol}^{-1}$ and $+63 \text{ kJ mol}^{-1}$), suggesting that Na⁺ ions stabilize building units of UFI better than K⁺. These binding energies explain the high Na⁺/Si and low K⁺/Si in DiffSyn-generated routes.

We experimentally test the DiffSyn-generated routes for UFI, resulting in four successful experimental syntheses of the UFI material (Fig. 5a, blue). Powder X-ray diffraction patterns for synthesized UFI samples closely match the simulated diffraction pattern (Fig. 5c), confirming

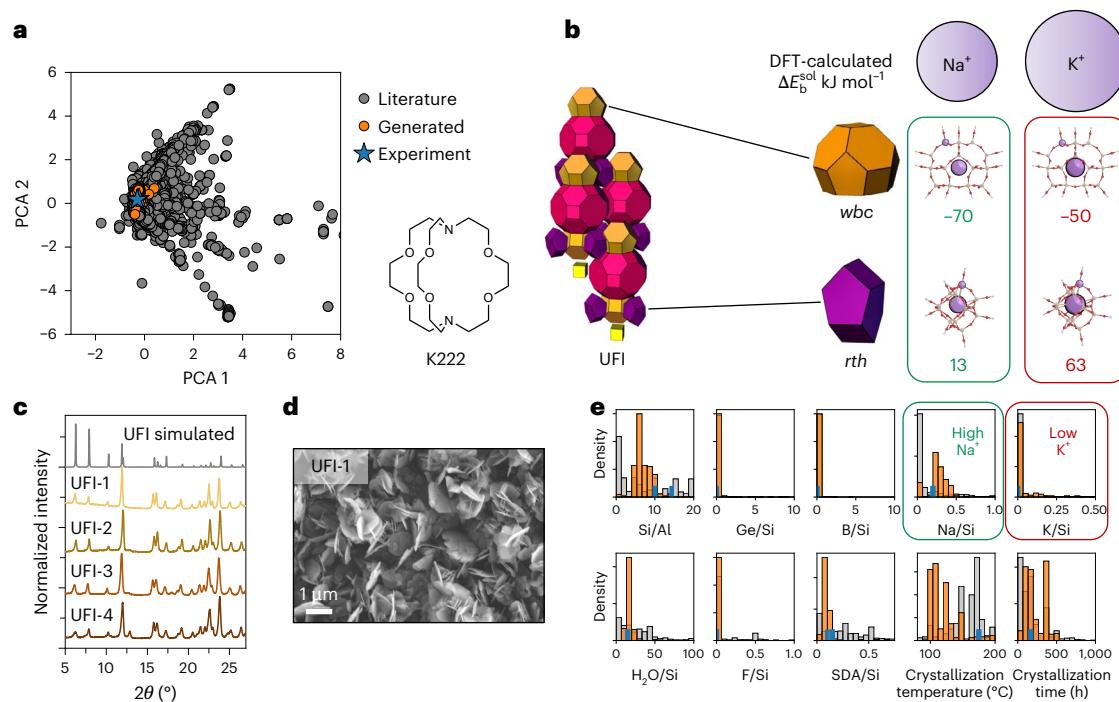


Fig. 5 | Experimental and DFT validation. We leverage DiffSyn to generate synthesis routes for the UFI structure, which resulted in successful crystallization of the desired structure. **a**, PCA of synthesis parameters of all routes reported in the literature (gray), DiffSyn-generated for UFI (orange) and experimentally verified (blue) in this work. **b**, DFT-calculated binding energies (ΔE_b) between cations (Na^+ , K^+) and building units of UFI (*rth*, *wbc*). Notice that Na^+ has stronger (more negative) ΔE_b than K^+ , which suggests that Na^+ (compared

with K^+) favors UFI formation. This rationalizes why DiffSyn suggested synthesis routes with high Na^+/Si and low K^+/Si for UFI (in **e**). **c**, Simulated and measured X-ray diffraction patterns of the synthesized UFI samples. **d**, Scanning electron microscopy image of UFI. **e**, Distributions of individual synthesis parameters. Blue are experimentally verified synthesis routes, orange is DiffSyn-generated and gray is all literature-reported routes.

that the crystallized material has the UFI structure. The resultant crystals had a measured $\text{Si}/\text{Al}_{\text{ICP}}$ of 19.0, one of the highest recorded in UFI synthesis (see 'Elemental analysis' in Methods for details on $\text{Si}/\text{Al}_{\text{ICP}}$). This composition is important because higher $\text{Si}/\text{Al}_{\text{ICP}}$ is associated with higher thermal stability under catalytic conditions⁵¹. In addition, the crystals show a house-of-cards morphology (Fig. 5d), making them highly promising catalysts for a diverse range of applications.

While the experimentally verified synthesis routes fall within the generated distributions, the model generated unusually low crystallization temperatures with a major mode at 100–150 °C (Fig. 5e). The generated minor mode at 175 °C was chosen based on domain expertise because higher temperatures accelerate crystallization kinetics and higher temperatures (>170 °C) avoid the formation of the undesired LTA competing phase. The higher temperature allowed us to use a relatively shorter crystallization time of 7 days (168 hours). These considerations underscore a powerful synergy between model and human expertise to achieve the desired synthesis outcome. Nonetheless, DiffSyn generated nine out of ten synthesis parameters as the major mode. The successful synthesis of UFI, enabled by DiffSyn is a testament to DiffSyn learning the intricacies of zeolite synthesis without explicit chemical encodings and highlights its ability to recommend suitable synthesis parameters for unseen materials systems.

Discussion

Zeolite synthesis is a complex task with a high-dimensional space where synthesis parameters interact. This complexity underscores the need for generative models that capture the intricacies of materials synthesis and circumvent expensive first-principles approaches. By capturing the one-to-many and multi-modal nature of structure–synthesis relationships, DiffSyn outperforms previous methods for zeolite synthesis prediction.

DiffSyn requires the OSDA to be known *a priori*. Including the OSDA as an input ensures that the model captures the influence of the organic template. Recent work on OSDA design has enabled the selection of suitable OSDA inputs into DiffSyn^{29,30,52}. The combination of OSDA design tools and DiffSyn could enable an end-to-end predictive workflow for zeolite synthesis. While our work focuses solely on continuous variables, the synthesis outcome can also be influenced by categorical/discrete variables, such as the precursor choice^{9,53,54} and presence of seed crystals. These discrete variables present a compelling opportunity for future work in ML for synthesis modeling—for example, discrete diffusion⁵⁵, autoregressive models^{56–59} and reinforcement learning^{60,61} for sampling discrete synthesis parameters.

For diffusion models, the inference speed is slower than for other generative approaches owing to sequential denoising. However, as the bottleneck of materials synthesis is experimental synthesis (for example, reaction time of weeks), inference speed is not a significant problem. Even if sampling speed becomes a bottleneck, techniques such as denoising diffusion implicit model sampling⁶² can trade off sampling quality for faster sampling. Flow-matching models⁶³ could be an alternative to balance quality and speed. Notably, training DiffSyn requires optimizing hyperparameters for classifier-free guidance such as p_{uncond} and w as shown in Supplementary Fig. 6, which may be challenging if compute is limited.

The successful synthesis of UFI enabled by DiffSyn exemplifies the potential of our approach to guide experimental zeolite synthesis. However, a comprehensive validation of DiffSyn would ideally also demonstrate that recipes substantially diverging from the generated ones fail to yield the target phase; such an evaluation is beyond the scope of this study and is reserved for future work. We hypothesize that diffusion models could be applied to other materials systems, particularly when data are plentiful. This approach signifies a shift

from regression to generative models for materials synthesis, where the latter is needed to effectively model one-to-many structure–synthesis relationships. This work has enabled the prediction of probable and diverse synthesis pathways of microporous materials, signifying a step toward bridging computational materials design ('what' to synthesize) and synthesis planning ('how' to synthesize).

Methods

Zeolite and OSDA representations

Zeolite. We adopt two different approaches to encoding the zeolite structure. First, invariant features (for example, ring sizes, largest included sphere) are retrieved from the International Zeolite Association (IZA) database⁶⁴. These features serve as inputs into a multi-layer perceptron (MLP) encoder. Second, an EGNN^{65,66} encodes the zeolite as a graph (Supplementary Section D).

OSDA. We optimized each OSDA in the gas phase with the MMFF94 force field from RDKit-generated conformers⁶⁷. Each OSDA is featurized using its physicochemical descriptors (for example, molecular volume and two-dimensional shape descriptors) averaged over all conformers^{28,29}. The OSDA features are defined in Supplementary Table 3. The zeolite embedding is concatenated with the OSDA embedding and further encoded using the fusion encoder before the joint embedding guides the reverse diffusion process to generate synthesis routes (Fig. 1b). Model performance across the denoising diffusion trajectory can be found in Supplementary Fig. 7.

Denoising diffusion probabilistic models

Denoising diffusion probabilistic models⁶⁸ are generative models that use a diffusion process to generate data by reversing a forward process that incrementally adds noise to the data. The forward process gradually corrupts the data x_0 into a noisy sample x_t by adding Gaussian noise in a Markov chain:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I) \quad (2)$$

where q is the forward/noising distribution, β_t is the variance at timestep t , and I is the identity matrix. It is often useful to sample x_t directly from x_0 :

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I) \quad (3)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and s is the diffusion time step. The generative process learns to reverse this corruption step by step (Fig. 1b), sampling from a distribution $p_\theta(x_{t-1}|x_t)$ parameterized by a neural network with parameters θ . This reverse process can also be represented as a Gaussian distribution

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (4)$$

Here, we learn the mean while fixing the variance as

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(x_t, t)), \quad \Sigma_\theta(x_t, t) = \beta_t I \quad (5)$$

where $\epsilon_\theta(x_t, t)$ is a neural network trained to predict the noise ϵ . In this work, we use a U-Net³² (Fig. 1c). The training objective is to minimize a variational bound $\mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2]$, which results in the diffusion model

$$p_\theta(x_0) = \int p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t) dx_{1:T} \quad (6)$$

Chemically guided diffusion model

In standard guided diffusion models, a classifier is used to guide the generation process by adjusting the score to steer the model toward

specific target classes. In contrast, classifier-free guidance²² eliminates the need for a separate classifier by conditioning the diffusion model directly on the desired attributes. During training, the score function s_θ is trained both with and without conditioning c using a null token \emptyset . This training is done by randomly setting c to the unconditional null token \emptyset with some probability p_{uncond} . Sampling is then performed using a linear combination of the conditional and unconditional score estimates:

$$\tilde{s}_\theta(x_t, t, c) = (1 + w)s_\theta(x_t, t, c) - ws_\theta(x_t, t, \emptyset) \quad (7)$$

where c refers to the chemical guidance from zeolite and OSDA embeddings shown in Fig. 1b, and w is the strength of the chemical guidance. Supplementary Section D contains the implementation details.

Model evaluation

Metrics. For each test zeolite–OSDA system, we sample 1,000 synthesis routes using the model and compute the following metrics with reference to unseen synthesis routes reported in literature. Wasserstein distance measures the distance between two probability distributions by finding the minimum cost to move probability mass from one distribution to another⁶⁹. The Wasserstein metric captures differences in both the location and the shape of distributions (including spread and the presence of multiple modes). Although Kullback–Leibler divergence may be used, the Wasserstein distance is chosen as it fulfills all requirements of a metric: non-negativity, identity of indiscernibles, symmetry and triangle inequality. Moreover, a significant drawback of Kullback–Leibler divergence is its behavior with distributions that do not have overlapping support. If there is any point where one distribution has a zero probability and the other has a non-zero probability, the Kullback–Leibler divergence can become infinite or undefined. Wasserstein distance, in contrast, provides a meaningful and finite distance even for distributions with non-overlapping supports. Wasserstein distance has been widely adopted in generative modeling (for example, Wasserstein GANs⁷⁰ use a form of Wasserstein distance) because it provides a smooth, sensitive measure of distributional differences. Small improvements in how well the model captures tails or secondary modes are reflected by a lower Wasserstein distance, guiding us during model development to favor settings that capture the full distribution. In addition, inspired by ref. 71, we use two coverage metrics, COV-R (recall) and COV-P (precision), to measure the similarity between sets of generated and literature-reported synthesis for each zeolite–OSDA system. Intuitively, COV-R measures the fraction of literature synthesis routes being correctly predicted, and COV-P measures the fraction of generated synthesis routes being probable. COV-F1 is computed as the harmonic mean of COV-R and COV-P. Refer to Supplementary Section C for a detailed justification of the metrics.

Baselines. A random dummy baseline corresponds to picking a random point in synthesis space. Reference 16 proposed a deterministic regression-based approach using AMD for zeolite structural featurization for a synthesis prediction task¹⁶. BNNs³³ extend standard neural networks by incorporating Bayesian inference by treating network weights as probability distributions, enabling a distribution of outputs. We implement a classical generative approach, GMM^{34,72}, which models data probabilistically as a sum of Gaussians (each with its mean and covariance). We also implement deep generative baselines: conditional VAE^{11,19,73}, conditional GAN³⁵ and conditional NF. For NF, we use real-valued non-volume preserving (RealNVP) transformations³⁶.

Experimental methods

Synthesis of UFI. We used colloidal silica (Ludox AS-40, 40 wt%, Aldrich), aluminum hydroxide (Al(OH)_3 , SPI Pharmacy), sodium hydroxide (NaOH , 50%, Aldrich), with K222 (98%, Ambeed) and tetramethylammonium hydroxide pentahydrate ($\text{TMAOH}\cdot 5\text{H}_2\text{O}$, 97%, Aldrich)

as OSDAs to synthesize UFI. The OSDAs were selected by the method described in ref. 48. The synthesis mixture had the molar composition 1.0 SiO₂, 0.033 Al₂O₃, 0.1 Na₂O, 15 H₂O, 0.15 K222, 0.025 TMA₂O. We dissolved 0.021 g Al(OH)₃ in a solution containing 0.64 g NaOH, 6.9 g deionized water and 0.36 g TMAOH·5H₂O. After the solution became clear, 6 g Ludox AS-40 and 2.31 g Kryptofix 222 were added. Also, 4 wt% H-UFI seed material was included. The mixture was stirred and aged at room temperature for 24 h, then transferred to Teflon-lined 23-ml autoclaves and heated at 175 °C under dynamic conditions for 7 days. The resulting zeolites were recovered by centrifugation (12,000 rpm, 10 min), washed with deionized water 3 times and dried overnight at 110 °C. Samples were then heated at a ramp rate of 1 °C min⁻¹ to 580 °C in flowing dry air (100 ml g⁻¹ of zeolite) for 6 h to remove OSDA molecules.

X-ray diffraction. Zeolite crystal structures were analyzed using powder X-ray diffraction with a Bruker D8 diffractometer (Cu K α radiation, $\lambda = 1.5418 \text{ \AA}$, 40 kV, 40 mA). Diffraction patterns were recorded over a 2θ range of 5–45°, with an angular step size of 0.02° and a scanning rate of 4° min⁻¹ to confirm phase purity.

Scanning electron microscopy. The morphology of the calcined zeolite crystals was observed using a Zeiss Merlin High-Resolution scanning electron microscope. Zeolite samples were prepared as fine powders and mounted on carbon tape. Images were collected at 2.0 kV, 100 pA and a working distance of 6.7 mm using the HE-SE2 detector in High-Resolution Column mode.

Elemental analysis. The elemental composition of silicon and aluminum (Si/Al_{ICP}) of synthesized UFI zeolites was determined using inductively coupled plasma atomic emission spectroscopy (Agilent 5100). Si/Al_{ICP} is an important property of zeolites as a high Si/Al_{ICP} is correlated with high thermal stability⁵¹. Samples (10 mg) were digested in 100 μl hydrofluoric acid (48 wt%, Sigma-Aldrich) for 24 h, followed by dilution to 10 g with 2 wt% aqueous nitric acid (GFS Chemicals). Calibration curves were constructed with 6-point standards using ICP solutions of 1,000 ppm Si, Al and Na in 2 wt% HNO₃ (Sigma-Aldrich, TraceCERT).

DFT methods

DFT calculations were conducted using the ORCA package (v5.0.4)^{74,75}. For all calculations, the ω B97X-D hybrid functional^{76,77} was used along with the def2 triple- ζ basis set with polarization (def2-TZVP)⁷⁸. During electronic optimization, wavefunctions were converged when energy changes were less than 10^{-8} Ha. Geometry optimizations were performed until the energy varied by less than 5×10^{-6} Ha and the forces on all atoms were below 3×10^{-4} Ha bohr⁻¹, adhering to the default ORCA optimization protocol. Where indicated, the SMD solvation model was applied to estimate solvation effects on the energies of each structure, using water as the solvent with a dielectric constant (ϵ) of 80.4 (ref. 79).

Binding energies were calculated relative to an isolated ion in vacuum and an empty Al-substituted composite building units (CBUs):

$$\Delta E_b = E[\text{M} \cdot \text{CBU}] - E[\text{M}^{n+}] - E[\text{CBU}^{n-}] \quad (8)$$

where $E[\text{M} \cdot \text{CBU}]$ represents the energy of the metal ion positioned within the Al-substituted CBU, $E[\text{M}^{n+}]$ is the energy of the isolated ion, and $E[\text{CBU}^{n-}]$ is the energy of the empty Al-substituted CBU containing n Al atoms. When CBUs offered multiple potential single Al positions or several configurations of two Al atoms, we selected the ΔE_b corresponding to the most stable Al arrangement, as this stability likely reflects the thermodynamically preferred configuration. Given adequate synthesis time, Al atoms tend to occupy stable framework positions, as demonstrated in recent CHA studies⁸⁰. Atoms within CBUs, including terminal hydroxyl groups, were kept fixed during calculations. Energies for empty, Al-substituted CBUs and isolated ions (in vacuum or with implicit solvation) were determined via single-point calculations in ORCA with appropriate charges.

The optimization of the metal-docked ion within the CBU allowed movement of the ion but maintained fixed positions for the CBU atoms, replicating the constraints imposed by a crystal environment where surrounding framework atoms would restrict movement; unconstrained optimization could lead to significant atomic displacements in these molecular forms. This study aims to understand how these ions influence specific CBU structures, given that the OSDA likely directs *Ita* cage formation in these zeolites. Thus, CBU atoms were fixed to preserve a shape close to their zeolite-based structures, and binding energies were calculated with and without solvation effects from water.

Zeolite crystal structures for materials synthesized in this work were obtained from the IZA database⁶⁴. Molecular models of individual CBUs for each synthesized framework were extracted by isolating them from their crystalline counterparts, with frameworks and corresponding CBUs. CBUs were derived from selected zeolite frameworks, as listed in the IZA database. Terminal SiOH groups were added to undercoordinated Si atoms to maintain tetrahedral coordination. These SiOH groups were oriented with Si–O–H bond angles of 180° to minimize hydrogen-bonding interactions between nearby SiOH groups (Supplementary Fig. 19).

Data availability

The ZeoSyn dataset is available on figshare via <https://doi.org/10.6084/m9.figshare.30632942> (ref. 81). Source data is available with this paper.

Code availability

The code for preprocessing the ZeoSyn dataset, training and evaluating the diffusion models is available at https://github.com/eltonpan/zeosyn_gen and on Zenodo⁸².

References

1. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O. & Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
2. Merchant, A. et al. Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023).
3. Sriram, A. et al. The Open DAC 2023 dataset and challenges for sorbent discovery in direct air capture. *ACS Cent. Sci.* **10**, 923–941 (2024).
4. Zhu, Y. et al. Uncovering the network modifier for highly disordered amorphous Li-garnet glass-ceramics. *Adv. Mater.* **36**, 2302438 (2024).
5. Kim, K. J., Balaish, M., Wadaguchi, M., Kong, L. & Rupp, J. L. Solid-state Li–metal batteries: challenges and horizons of oxide and sulfide solid electrolytes and their interfaces. *Adv. Energy Mater.* **11**, 2002689 (2021).
6. Zeni, C. et al. A generative model for inorganic materials design. *Nature* **639**, 624–632 (2025).
7. Barroso-Luque, L. et al. Open Materials 2024 (OMat24) inorganic materials dataset and models. Preprint at <https://arxiv.org/abs/2410.12771> (2024).
8. Saal, J. E., Kirklin, S., Aykolic, M., Meredig, B. & Wolverton, C. Materials design and discovery with high-throughput density functional theory: the Open Quantum Materials Database (OQMD). *JOM* **65**, 1501–1509 (2013).
9. He, T. et al. Precursor recommendation for inorganic synthesis by machine learning materials similarity from scientific literature. *Sci. Adv.* **9**, eadg8180 (2023).
10. Huo, H. et al. Machine-learning rationalization and prediction of solid-state synthesis conditions. *Chem. Mater.* **34**, 7323–7336 (2022).
11. Karpovich, C., Pan, E., Jensen, Z. & Olivetti, E. Interpretable machine learning enabled inorganic reaction classification and synthesis condition prediction. *Chem. Mater.* **35**, 1062–1079 (2023).

12. Mahbub, R. et al. Text mining for processing conditions of solid-state battery electrolytes. *Electrochem. Commun.* **121**, 106860 (2020).

13. Bianchini, M. et al. The interplay between thermodynamics and kinetics in the solid-state synthesis of layered oxides. *Nat. Mater.* **19**, 1088–1095 (2020).

14. Pan, E. et al. ZeoSyn: a comprehensive zeolite synthesis dataset enabling machine-learning rationalization of hydrothermal parameters. *ACS Cent. Sci.* **10**, 729–743 (2024).

15. Prein, T. et al. Reaction graph networks for inorganic synthesis condition prediction of solid state materials. In *38th Conference on Neural Information Processing Systems (NeurIPS 2024)* <https://openreview.net/pdf?id=VGsXQOTs1E> (NeurIPS, 2024).

16. Schwalbe-Koda, D., Widdowson, D. E., Pham, T. A. & Kurlin, V. A. Inorganic synthesis–structure maps in zeolites with machine learning and crystallographic distances. *Digit. Discov.* **2**, 1911–1924 (2023).

17. Luo, Y. et al. MOF synthesis prediction enabled by automatic data mining and machine learning. *Angew. Chem. Int. Ed.* **61**, e202200242 (2022).

18. Pan, E. et al. A chemically-guided generative diffusion model for materials synthesis planning. In *38th Conference on Neural Information Processing Systems (NeurIPS 2024) AI4Mat Workshop* <https://openreview.net/pdf/0d92f126ae589466badd426ddaa31b5aa37de448.pdf> (NeurIPS, 2024).

19. Karpovich, C., Jensen, Z., Venugopal, V. & Olivetti, E. Inorganic synthesis reaction condition prediction with generative machine learning. Preprint at <https://arxiv.org/abs/2112.09612> (2021).

20. Saharia, C. et al. Photorealistic text-to-image diffusion models with deep language understanding. In *36th Conference on Neural Information Processing Systems (NeurIPS 2022)* https://papers.neurips.cc/paper_files/paper/2022/file/ec795aaeadae0b7d230fa35cbaf04c041-Paper-Conference.pdf (NeurIPS, 2022).

21. Ramesh, A. et al. Zero-shot text-to-image generation. In *Proc. 37th International Conference on Machine Learning* <https://proceedings.mlr.press/v139/ramesh21a/ramesh21a.pdf> (PMLR, 2020).

22. Ho, J. & Salimans, T. Classifier-free diffusion guidance. Preprint at <https://arxiv.org/abs/2207.12598> (2022).

23. Dhariwal, P. & Nichol, A. Diffusion models beat GANs on image synthesis. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021)* <https://proceedings.nips.cc/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf> (NeurIPS, 2021).

24. Zhang, Z., Li, M. & Yu, J. On the convergence and mode collapse of GAN. In *SIGGRAPH Asia 2018 Technical Briefs (SA '18)* <https://doi.org/10.1145/3283254.3283282> (Association for Computing Machinery, 2018).

25. Moliner, M., Román-Leshkov, Y. & Corma, A. Machine learning applied to zeolite synthesis: the missing link for realizing high-throughput discovery. *Acc. Chem. Res.* **52**, 2971–2980 (2019).

26. Lewis, D., Freeman, C. & Catlow, C. Predicting the templating ability of organic additives for the synthesis of microporous materials. *J. Phys. Chem.* **99**, 11194–11202 (1995).

27. Zones, S., Nakagawa, Y., Yuen, L.-T. & Harris, T. Guest/host interactions in high silica zeolite synthesis: [5.2.1.02. 6] tricyclodecanes as template molecule. *J. Am. Chem. Soc.* **118**, 7558–7567 (1996).

28. Jensen, Z. et al. Discovering relationships between OSDAs and zeolites through data mining and generative neural networks. *ACS Cent. Sci.* **7**, 858–867 (2021).

29. Schwalbe-Koda, D. et al. A priori control of zeolite phase competition and intergrowth with high-throughput simulations. *Science* **374**, 308–315 (2021).

30. Hoffman, A., Xie, M. & Gómez-Bombarelli, R. Learning descriptors to predict organic structure-directing agent applicability in zeolite synthesis. *Microporous Mesoporous Mater.* **386**, 113467 (2024).

31. Xie, M. et al. A comprehensive mapping of zeolite-template chemical space. *Nat. Comput. Sci.* **5**, 661–674 (2025).

32. Ronneberger, O., Fischer, P. & Brox, T. *U-net: Convolutional Networks for Biomedical Image Segmentation* 234–241 (Springer, 2015).

33. Jospin, L. V., Laga, H., Boussaid, F., Buntine, W. & Bennamoun, M. Hands-on Bayesian neural networks—a tutorial for deep learning users. *IEEE Comput. Intell. Mag.* **17**, 29–48 (2022).

34. Reynolds, D. Gaussian mixture models. in *Encyclopedia of Biometrics* (ed Li, S. Z.) 659–663 (Springer, 2009); https://doi.org/10.1007/978-0-387-73003-5_196

35. Mirza, M. & Osindero, S. Conditional generative adversarial nets. Preprint at <https://arxiv.org/abs/1411.1784> (2014).

36. Dinh, L., Sohl-Dickstein, J. & Bengio, S. Density estimation using real NVP. In *International Conference on Learning Representations 2017* <https://openreview.net/pdf?id=HkpbnH9lx> (ICLR, 2017).

37. Jia, X. et al. Anthropogenic biases in chemical reaction data hinder exploratory inorganic synthesis. *Nature* **573**, 251–255 (2019).

38. Cambor, M. A., Villaescusa, L. A. & Diaz-Cabanas, M. Synthesis of all-silica and high-silica molecular sieves in fluoride media. *Top. Catal.* **9**, 59–76 (1999).

39. Le, T. et al. Process regulation of microwave intensified synthesis of y-type zeolite. *Microporous Mesoporous Mater.* **284**, 476–485 (2019).

40. Corma, A. et al. Proposed pore volume topology of zeolite MCM-22 based on catalytic tests. *Appl. Catal. A* **115**, 121–134 (1994).

41. Ma, D. et al. Mo/HMCM-22 catalysts for methane dehydroaromatization: a multinuclear MAS NMR study. *J. Phys. Chem. B* **105**, 1786–1793 (2001).

42. Zhang, G. et al. Single isomerization selectivity of glucose in methanol over Sn-BEC zeolite of homogenous Sn distribution. *Microporous Mesoporous Mater.* **247**, 158–165 (2017).

43. Moliner, M. et al. Synthesis of the Ti-silicate form of BEC polymorph of β -zeolite assisted by molecular modeling. *J. Phys. Chem. C* **112**, 19547–19554 (2008).

44. Villaescusa, L. et al. Pure silica large pore zeolite ITQ-7: synthetic strategies, structure-directing effects, and control and nature of structural disorder. *Chem. Mater.* **19**, 1601–1612 (2007).

45. Oleksiak, M. D. et al. Organic-free synthesis of a highly siliceous faujasite zeolite with spatially biased Q4 (nAl) Si speciation. *Angew. Chem. Int. Ed.* **129**, 13551–13556 (2017).

46. Muraoka, K., Chaikittisilp, W. & Okubo, T. Multi-objective de novo molecular design of organic structure-directing agents for zeolites using nature-inspired ant colony optimization. *Chem. Sci.* **11**, 8214–8223 (2020).

47. Di Iorio, J. R. & Gounder, R. Controlling the isolation and pairing of aluminum in chabazite zeolites using mixtures of organic and inorganic structure-directing agents. *Chem. Mater.* **28**, 2236–2247 (2016).

48. Lee, H. et al. Data-driven SDA design and high-throughput experimentation for the synthesis of high-silica small-pore zeolites containing lta-cavities. Preprint at *ChemRxiv* <https://doi.org/10.26434/chemrxiv-2025-d7qzq> (2025).

49. Jo, D. et al. Synthesis of high-silica LTA and UFI zeolites and $\text{NH}_3\text{-SCR}$ performance of their copper-exchanged form. *ACS Catal.* **6**, 2443–2447 (2016).

50. Schwalbe-Koda, D., Jensen, Z., Olivetti, E. & Gómez-Bombarelli, R. Graph similarity drives zeolite diffusionless transformations and intergrowth. *Nat. Mater.* **18**, 1177–1181 (2019).

51. Cruciani, G. Zeolites upon heating: factors governing their thermal stability and structural changes. *J. Phys. Chem. Solids* **67**, 1973–1994 (2006).
52. Xie, M. et al. A comprehensive mapping of zeolite-template chemical space. *Nat. Comput. Sci.* **5**, 661–674 (2025).
53. Prein, T. et al. Retro-rank-in: a ranking-based approach for inorganic materials synthesis planning. Preprint at <https://arxiv.org/abs/2502.04289> (2025).
54. Noh, H., Lee, N., Na, G. S. & Park, C. Retrieval-retro: retrieval-based inorganic retrosynthesis with expert knowledge. In *38th Conference on Neural Information Processing Systems (NeurIPS 2024)* https://proceedings.neurips.cc/paper_files/paper/2024/file/2cfa9b0d9be8a5c01cf3eb7f21b4f2b8-Paper-Conference.pdf (2024).
55. Lou, A., Meng, C. & Ermon, S. Discrete diffusion language modeling by estimating the ratios of the data distribution. In *International Conference on Machine Learning 2024* <https://openreview.net/pdf?id=71mqtQdKB9> (ICLR, 2024).
56. Kim, S., Jung, Y. & Schrier, J. Large language models for inorganic synthesis predictions. *J. Am. Chem. Soc.* **146**, 19654–19659 (2024).
57. Okabe, R. et al. Large language model-guided prediction toward quantum materials synthesis. Preprint at <https://arxiv.org/abs/2410.20976> (2024).
58. Prein, T. et al. Language models enable data-augmented synthesis planning for inorganic materials. *ACS Appl. Mater. Interfaces* **17**, 69221–69233 (2025).
59. Song, Z., Lu, S., Ju, M., Zhou, Q. & Wang, J. Accurate prediction of synthesizability and precursors of 3D crystal structures via large language models. *Nat. Commun.* **16**, 6530 (2025).
60. Karpovich, C., Pan, E. & Olivetti, E. A. Deep reinforcement learning for inverse inorganic materials design. *npj Comput. Mater.* **10**, 287 (2024).
61. Pan, E., Karpovich, C. & Olivetti, E. Deep reinforcement learning for inverse inorganic materials design. Preprint at <https://arxiv.org/abs/2210.11931> (2022).
62. Song, J., Meng, C. & Ermon, S. Denoising diffusion implicit models. In *International Conference on Learning Representations* <https://openreview.net/pdf?id=St1giarCHLP> (ICLR, 2021).
63. Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M. & Le, M. Flow matching for generative modeling. In *International Conference on Learning Representations* <https://openreview.net/pdf?id=PqvMRDCJT9t> (ICLR, 2023).
64. Baerlocher, C. & McCusker, L. B. *Database of Zeolite Structures* (International Zeolite Association, 2021); <http://www.iza-structure.org/databases/>
65. Satorras, V. G., Hoogeboom, E. & Welling, M. E(n) equivariant graph neural networks. In *Proc. 38th International Conference on Machine Learning* <https://proceedings.mlr.press/v139/satorras21a/satorras21a.pdf> (PMLR, 2021).
66. Sheriff, K., Cao, Y. & Freitas, R. Chemical-motif characterization of short-range order with E(3)-equivariant graph neural networks. *npj Comput. Mater.* **10**, 215 (2024).
67. Halgren, T. A. MMFF VI. MMFF94s option for energy minimization studies. *J. Comput. Chem.* **20**, 720–729 (1999).
68. Ho, J., Jain, A. & Abbeel, P. Denoising diffusion probabilistic models. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)* <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf> (NeurIPS, 2020).
69. Rubner, Y., Tomasi, C. & Guibas, L. J. A metric for distributions with applications to image databases. In *Proc. 1998 IEEE International Conference on Computer Vision* <https://users.cs.duke.edu/~tomasi/papers/rubner/rubnerlccv98.pdf> (IEEE, 1998).
70. Arjovsky, M., Chintala, S. & Bottou, L. Wasserstein generative adversarial networks. In *Proc. 34th International Conference on Machine Learning* 214–223 (PMLR, 2017).
71. Xie, T., Fu, X., Ganea, O.-E., Barzilay, R. & Jaakkola, T. Crystal diffusion variational autoencoder for periodic material generation. In *Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS 2021)* https://ml4physicalsciences.github.io/2021/files/NeurIPS_ML4PS_2021_48.pdf (NeurIPS, 2021).
72. Sugiyama, M. et al. Conditional density estimation via least-squares density ratio estimation. In *Proc. Thirteenth International Conference on Artificial Intelligence and Statistics* 781–788 (PMLR, 2010).
73. Higgins, I. et al. Beta-VAE: learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations 2017* <https://openreview.net/pdf?id=Sy2fzU9gl> (ICLR, 2017).
74. Neese, F. The ORCA program system. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2**, 73–78 (2012).
75. Neese, F. Software update: the ORCA program system, version 4.0. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **8**, e1327 (2018).
76. Chai, J.-D. & Head-Gordon, M. Long-range corrected hybrid density functionals with damped atom–atom dispersion corrections. *Phys. Chem. Chem. Phys.* **10**, 6615–6620 (2008).
77. Chai, J.-D. & Head-Gordon, M. Systematic optimization of long-range corrected hybrid density functionals. *J. Chem. Phys.* **128**, 084106 (2008).
78. Weigend, F. & Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: design and assessment of accuracy. *Phys. Chem. Chem. Phys.* **7**, 3297–3305 (2005).
79. Marenich, A. V., Cramer, C. J. & Truhlar, D. G. Universal solvation model based on solute electron density and on a continuum model of the solvent defined by the bulk dielectric constant and atomic surface tensions. *J. Phys. Chem. B* **113**, 6378–6396 (2009).
80. Lee, S. et al. Evolution of framework Al arrangements in CHA zeolites during crystallization in the presence of organic and inorganic structure-directing agents. *Cryst. Growth Des.* **22**, 6275–6295 (2022).
81. Pan, E. ZeoSyn: a comprehensive zeolite synthesis dataset enabling machine-learning rationalization of hydrothermal parameters. *figshare* <https://doi.org/10.6084/m9.figshare.30632942.v1> (2025).
82. Pan, E. Code (v1.0.1) for DiffSyn: a generative diffusion approach to materials synthesis planning. *Zenodo* <https://doi.org/10.5281/zenodo.17645371> (2025).

Acknowledgements

We acknowledge funding from the Generalitat Valenciana through the Prometeo Program (CIPROM/2023/34), CSIC through the I-link+ Program (ILINK24035), and MIT International Science and Technology Initiatives (MISTI) Seed Funds between MIT and CSIC. We also acknowledge partial funding from the National Science Foundation DMREF Awards 1922090, 1922311, 1922372, the Office of Naval Research (ONR) under contract N00014-20-1-2280, Kwanjeong Educational Fellowship, ExxonMobil, and the Agency for Science, Technology and Research. We thank D. Du and A. Burton for helpful discussions and suggestions.

Author contributions

E.P. and E.A.O. conceptualized the project. E.P. curated the dataset, developed the generative diffusion framework (and all baselines) for materials synthesis planning and performed downstream model analysis. E.P. created the figures and wrote the original draft of the paper. S.K. synthesized and characterized the zeolite materials. E.P., S.K. and S.L. developed domain-specific metrics on evaluating the

synthesis planning task. M.X. calculated the OSDA descriptors. A.J.H. calculated the DFT binding energies. Y.D. and T.P. provided feedback on the ML modeling. K.S. developed the equivariant graph representations of the zeolite structures. E.A.O., R.G.-B., M.M. and Y.R.-L. supervised the project. All authors reviewed and edited the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s43588-025-00949-9>.

Correspondence and requests for materials should be addressed to Elsa A. Olivetti.

Peer review information *Nature Computational Science* thanks Bo Lei, Nathan Szymanski and the other, anonymous, reviewer(s) for their

contribution to the peer review of this work. Primary Handling Editor: Kaitlin McCardle, in collaboration with the *Nature Computational Science* team.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2026