

A recommendation system to predict missing adsorption properties of nanoporous materials

Arni Sturluson¹, Ali Raza², Grant D. McConachie¹, Daniel W. Siderius³, Xiaoli Fern^{†2},
and Cory M. Simon^{†1}

¹School of Chemical, Biological, and Environmental Engineering. Oregon State University. Corvallis, OR.

²School of Electrical Engineering and Computer Science. Oregon State University. Corvallis, OR, USA

³National Institute of Standards and Technology. Chemical Sciences Division. Gaithersburg, MD, USA.

[†]{xiaoli.fern, cory.simon}@oregonstate.edu

March 15, 2021

Abstract

Nanoporous materials (NPMs) selectively adsorb and concentrate gases into their pores, and thus could be used to store, capture, and sense many different gases. Modularly synthesized classes of NPMs, such as covalent organic frameworks (COFs), offer a large number of candidate structures for each adsorption task. A complete NPM-property table, containing measurements of the relevant adsorption properties in the candidate NPMs, would enable the matching of NPMs with adsorption tasks. However, in practice the NPM-property matrix is only partially observed (incomplete); (i) many properties of any given NPM have not been measured and (ii) any given property has not been measured for all NPMs.

The idea in this work is to leverage the observed (NPM, property) values to impute the missing ones. Similarly, commercial recommendation systems impute missing entries in an incomplete item-customer ratings matrix to recommend items to customers. We demonstrate a COF recommendation system to match COFs with adsorption tasks by training a low rank model of an incomplete COF-adsorption-property matrix. A low rank model, trained on the observed (COF, adsorption property) values, provides (i) predictions of the missing (COF, adsorption property) values and (ii) a “map” of COFs, wherein COFs with similar (dissimilar) adsorption properties congregate (separate). We find the performance of the COF recommendation system varies for different adsorption tasks and diminishes precipitously when the fraction of missing entries exceeds 60 %. The concepts in our COF recommendation system can be applied broadly to many different materials and properties.

1 Introduction

Nanoporous materials (NPMs) [1] often exhibit permanent porosity and possess large-area internal surfaces [2] decorated with functional groups. This enables them to (selectively) adsorb and concentrate gas in their pores [3–5]. As a result, NPMs have applications in storing [5], separating [4,6], and sensing [7] gases, as well as in catalysis [8].

Advanced families of NPMs, such as metal-organic frameworks (MOFs) [9], covalent organic frameworks (COFs) [10], porous polymer networks (PPNs) [11], porous organic cages (POCs) [12], and metal-organic polyhedra (MOPs) [13,14], are constructed modularly from molecular building blocks. The copiousness of building blocks, together with post-synthetic modifiability [15], make the number of possible NPM structures extremely large [16].

The vastness of NPM “space” [17] and the many adsorption-based applications of NPMs give rise to two important optimization problems. The first is application-led NPM search [1], where the goal is to search for an NPM structure with an optimal adsorption property for a desired application. The second is material-led application search [18], where the goal is to search for the most suitable application of a given NPM.

Given a list of candidate NPM structures and a list of adsorption properties relevant to various applications, the complete *NPM-property matrix* containing measurements for every (NPM, property) pair reduces both application-led NPM search and NPM-led application search to look-up problems. However, in practice, the NPM-property matrix, constructed from experimental data collected from the literature [19] and/or databases of simulated gas adsorption in libraries of NPMs [18], is likely incomplete, because many (NPM, property) values have not been observed. I.e., (i) for any given NPM, only a proportion of its adsorption properties have been measured, and (ii) for any given adsorption property, it has been measured in only a proportion of the NPMs. See Fig. 1.

The idea in this work is to leverage the observed (NPM, property) values to predict the missing ones—i.e., to impute the missing values of, or complete, the NPM-property matrix. This predic-




Figure 1: A recommendation system for nanoporous materials (NPMs). A toy NPM-adsorption-property matrix is illustrated. Entry (i, j) of the matrix represents the value of adsorption property j of NPM i . Many entries are unobserved (white) because measurements are missing. The goal of our NPM recommendation system is to use the observed entries to impute the unobserved entries, allowing recommendation of NPMs for various adsorption tasks (requiring a certain adsorption property). This is analogous to commercial recommendation systems that aim to recommend customer-specific items to customers, with NPM : item :: adsorption property :: customer.

tion task, of interest primarily to recommend NPMs for specific adsorption tasks, is analogous to a commercial recommendation system for items to customers (see Box and Fig. 1), materials : items :: properties : customers. A machine learning strategy to complete the NPM-property matrix is much less expensive and time-consuming than experimentally measuring or computationally simulating these missing properties. The machine-completed NPM-property matrix is valuable because it can be used to direct higher-fidelity but more expensive experimental measurements of properties towards the most promising materials—allowing a more efficient use of resources in both application-led material search and material-led application search.

Box: analogy with commercial recommendation systems

In commercial recommendation systems, observed (item, customer) ratings are used to predict missing ratings for the recommendation of items to customers [20]. For example, movie ratings by Netflix users can be stored in a movie-user ratings matrix (rows=movies, columns=users, entries=ratings) [21]. Most entries in the matrix are missing, as (i) each user has rated only a small proportion of the movies and (ii) each movie is rated by only a small proportion of the users. A movie recommendation system imputes the missing (movie, user) ratings using the observed ones (perhaps, using features of the movies and users as well) in order to make user-specific recommendations of movies. Thus, our (material, property) values are analogous to (item, customer) ratings in commercial recommendation systems. Though, two distinctions are (i) typically, there are many more customers than items in commercial recommendation systems, compared to more materials than properties in material recommendation systems, and (ii) the entries in an item-customer ratings matrix have the same units and scale, whereas the units and scale could vary across different properties in the material-property matrix. Distinction (i) is inconsequential for low rank models because the model is transpose-equivariant.

Our hypothesis, which would permit accurate matrix completion, is that the NPM-property matrix exhibits a low rank structure [22,23], owing to underlying structural and chemical similarities among both NPMs and gas species that dictate their interactions. A low rank structure implies both NPMs and adsorption properties can be represented by low-dimensional vectors that together express the affinity between a (NPM, property) pair; these latent representations can be machine-learned, jointly, using the observed (NPM, adsorption property) values and then used to impute the missing values [21,24].

Herein, we demonstrate the imputation of missing data in an incomplete material-property matrix through learning low rank models [24]. Particularly, we train low rank models of COF-gas-adsorption-property matrices, pertaining to 560 experimentally-reported COFs [25] and the simulated uptake of CH₄, H₂O, H₂S, Xe, Kr, CO₂, N₂, O₂ and H₂ at various conditions [18] that apply to different gas storage and separation applications of COFs. Advantageously, this COF-gas-adsorption-property matrix is in reality complete, allowing us to artificially introduce different fractions of missing values and investigate the effect of sparsity on the performance of the recommendation system. From the observed (COF, gas adsorption) values, the low rank model machine-learns low-dimensional latent vector representations of both the COFs and the adsorption properties, al-

lowing the (i) imputation of the missing values of the adsorption properties and (ii) drawing of a COF “map” that clusters together COFs with similar adsorption properties.

2 Review of previous work

2.1 Computational methods for materials discovery

Virtual screenings of NPMs for adsorption-based applications use molecular models and simulations [26, 27] to [cheaply, relative to conducting an experiment in the lab] predict the adsorption property of each candidate material [16, 28, 29]. As opposed to an exhaustive virtual screening, genetic algorithms [30–32] and Monte Carlo tree search [33] have been used to more efficiently search for the NPM(s) with the optimal adsorption property. Supervised machine learning models have been widely used to predict the adsorption properties of the NPMs at a lower cost than but similar fidelity to the molecular simulations [34–41]. In this approach, (i) molecular simulations are used to label a small subset of the candidate NPMs with the gas adsorption property, (ii) these examples are used to train a machine learning model to predict the gas adsorption of an NPM from cheaply computed structural and chemical features, and then (iii) the supervised machine learning model is used, as a surrogate model for the molecular simulations, to (cheaply) predict the properties of the remaining materials. See reviews in Refs. [42–48]. Limitations of such a task-specific supervised learning approach are (a) knowledge is not transferred between tasks, (b) the correlation between different targets is not leveraged, and (c) a task-specific, information-rich feature vector representation of the NPM is required. Though, transfer learning [49] and graph neural networks [50–52] alleviate limitations (a) and (c), respectively.

Unsupervised machine learning holds promise for materials discovery by clustering together materials with similar structures and thus properties and learning a low-dimensional (e.g., 2D) embedding of their structures into a “map” of materials [53–55]. Such a map facilitates lead-optimization approaches to materials discovery and the selection of structurally diverse sets of materials to adequately explore materials space.

Finally, autoencoders enable inverse design [56–58], where one specifies a desired adsorption property, and the machine learning model generates a NPM structure with that property.

2.2 Open data to enable machine-learning accelerated materials discovery

To spur and enable developments of machine learning approaches to materials discovery, open, structured databases [59–61] of (i) crystal structure models of NPMs [25, 62–66], (ii) simulated [18, 65, 67] or experimentally measured [68] adsorption properties of NPMs, and (iii) electronic properties of NPMs [69, 70], have been curated.

2.3 Placing our work in context

Our material recommendation system deviates from previous data-driven approaches to predict adsorption properties of NPMs by combining observations of many different properties to impute missing ones without using explicitly hand-crafted [71–73] features of the NPMs. Instead, from the observed (NPM, adsorption property) values, our recommendation system jointly machine-learns latent representations of both the NPMs and adsorption properties that express (NPM, adsorption property) affinities, taking advantage of low rank structure in the material-property matrix. In some sense, our recommendation system is related to multi-task learning, but it (i) does not require a hand-crafted or machine-learned vector representation of the NPM and (ii) handles missing values in the target vectors associated with the NPMs.

Our work is not the first machine-learned recommendation system for use in the chemical sciences. Yuana et al. [74] imputed missing gas permeability in polymers. Sosnia et al. [75] developed a recommendation system for antiviral drugs by learning a low rank model of a compound-virus activity matrix. Seko et al. [76] and Hayashi et al. [77] used matrix/tensor factorization to predict stability based on composition and optimal processing conditions, respectively, of inorganic materials.

3 The material recommendation system

Here, we formulate the general problem of material-property matrix completion.

A material recommendation system jointly machine-learns, from observed (material, property) values, low-dimensional latent vector representations of the materials and properties that express (material, property) affinities. These learned representations allow us to (i) impute the missing (material, property) values and (ii) draw a map of the materials, wherein materials with similar properties congregate.

The data. We have observations of $A_{mp} \in \mathbb{R}$, the value of property p in material m , for $(m, p) \in \Omega \subset \{1, 2, \dots, M\} \times \{1, 2, \dots, P\}$, which defines Ω as the set of ordered pairs describing the entries in \mathbf{A} that are observed. That is, the material-property matrix $\mathbf{A} \in \mathbb{R}^{M \times P}$, whose entry (m, p) is A_{mp} , is not complete; some entries are missing ($|\Omega| < MP$).

The objective. The objective is to complete the material-property matrix by predicting the missing entries, A_{mp} for $(m, p) \in \{1, 2, \dots, M\} \times \{1, 2, \dots, P\} \setminus \Omega$.

The low-rank model. From an element perspective, the low-rank model assumes that each element of the matrix, A_{mp} , decomposes into

$$A_{mp} \approx \mathbf{m}_m^T \mathbf{p}_p + \mu_m, \tag{1}$$

where $\mathbf{m}_m \in \mathbb{R}^k$ and $\mathbf{p}_p \in \mathbb{R}^k$ are low-dimensional ($k < M, P$), latent vector representations of material m and property p , respectively, and $\mu_m \in \mathbb{R}$ is a bias for material m . The material-property interaction term, the dot product $\mathbf{m}_m^T \mathbf{p}_p$, represents the “affinity” (if positive) or “aversion” (if negative) of material m for property p . Geometrically, the interaction term is positive (negative) if \mathbf{m}_m and \mathbf{p}_p point in roughly the same (opposite) direction. The magnitude of the interaction term depends on both the angle between \mathbf{m}_m and \mathbf{p}_p and their norms. The material bias μ_m reflects variation of the values of the properties of material m independent of interactions; some materials may simply tend to have higher or lower values of the properties. See Koren et al. [21].

From a matrix perspective, the low rank model factorizes the material-property matrix \mathbf{A} as:

$$\mathbf{A} \approx \mathbf{M}^T \mathbf{P} + \boldsymbol{\mu} \mathbf{1}^T \quad (2)$$

with the columns of matrices $\mathbf{M} \in \mathbb{R}^{k \times M}$ and $\mathbf{P} \in \mathbb{R}^{k \times P}$ containing the latent representations of materials and properties, respectively; the entries of the column vector $\boldsymbol{\mu} \in \mathbb{R}^M$ containing the material biases; and $\mathbf{1} \in \mathbb{R}^P$ a column vector of ones. See Fig. 2. The dimensionality of the latent space, $k < M, P$, imposes the constraint $\text{rank}(\mathbf{M}^T \mathbf{P}) \leq k$, hence eqn. 2 is a low rank model / approximation of the matrix \mathbf{A} .




Figure 2: The low rank model of the matrix $\mathbf{A} \approx \mathbf{M}^T \mathbf{P} + \boldsymbol{\mu} \mathbf{1}^T$. The columns of \mathbf{M} and \mathbf{P} contain the latent representations of the M materials and P properties, respectively, which lie in a k -dimensional space. The vector $\boldsymbol{\mu}$ contains the M material biases.

The utility of the low rank model. The low rank model of the materials-property matrix is useful for two purposes [24].

(1) *Imputation of missing entries.* The decomposition in eqn. 1 holds for both observed and unobserved (material, property) values. Thus, once we learn \mathbf{M} , \mathbf{P} , and $\boldsymbol{\mu}$ from the *observed* entries, we can predict the unobserved entries, as is clear from eqn. 2.

(2) *Construction of a low-dimensional map of the materials and properties.* The rows of a fully observed version of \mathbf{A} , which lie in a P -dimensional vector space, can be viewed as feature vectors of the materials. In this view, each material is represented by a list of its properties. The set of latent vector representations of the materials, in the rows of \mathbf{M}^T , are embeddings/ compressions of the rows of \mathbf{A} into a lower ($k < P$) dimensional vector space. [24] Within this latent space, materials, represented by \mathbf{m}_m 's, that tend to have similar (dissimilar) properties congregate (separate). Thus, with the latent representations of the materials, the \mathbf{m}_m 's, we can (i) use clustering algorithms to group together materials with similar properties and (ii) visualize the scatter of the materials in the low-dimensional space to make a “map” of materials.

Similarly, the columns of a complete version of \mathbf{A} can be viewed as vector representations of the properties, and the columns of \mathbf{P} , the latent vector representations of the properties, are embeddings/compressions of them. Within this latent space, properties, represented by \mathbf{p}_p 's, that tend to take on similar (dissimilar) values in NPMs congregate (separate).

As a consequence of the dot product $\mathbf{m}_m^T \mathbf{p}_p$ in eqn. 1, the magnitude and directions of a pair of latent material and property vectors $(\mathbf{m}_m, \mathbf{p}_p)$, taken together, indicate the affinity/aversion for each other, since $\mathbf{m}_m^T \mathbf{p}_p = \|\mathbf{m}_m\|_2 \|\mathbf{p}_p\|_2 \cos \phi$, with ϕ the angle between \mathbf{m}_m and \mathbf{p}_p .

Machine-learning the low rank model. We learn the latent representations of the materials and properties and the material biases by balancing (i) the matching of the observed values of the matrix by the model given in eqn. 1 and (ii) the complexity of the latent vector representations, to avoid overfitting. Specifically, we aim to choose the \mathbf{M} , \mathbf{P} , and $\boldsymbol{\mu}$ that minimize the loss $\ell = \ell(\mathbf{M}, \mathbf{P}, \boldsymbol{\mu})$:

$$\ell(\mathbf{M}, \mathbf{P}, \boldsymbol{\mu}) = \sum_{(m,p) \in \Omega} [A_{mp} - (\mathbf{m}_m^T \mathbf{p}_p + \mu_m)]^2 + \lambda \left(\frac{1}{M} \sum_{m=1}^M \|\mathbf{m}_m\|_2^2 + \frac{1}{P} \sum_{p=1}^P \|\mathbf{p}_p\|_2^2 \right). \quad (3)$$

The first term is the approximation error, measured over all observed (m, p) pairs. The second term provides L2 regularization of the latent vector representations of the materials and properties to prevent overfitting and improve generalization, where $\lambda > 0$ is the regularization parameter. The sums are normalized by the number of elements in the sum to properly weigh regularization of the latent material and property vectors.

Either stochastic gradient descent or alternating minimization can be used to find the $(\mathbf{M}, \mathbf{P}, \boldsymbol{\mu})$ that minimize ℓ . The latter alternates between fixing \mathbf{M} and optimizing \mathbf{P} and fixing \mathbf{P} and optimizing \mathbf{M} . See Refs. [21, 24].

There are two hyperparameters in the low rank model: (1) $k \in \{0, 1, \dots, \min(M, P)\}$, the dimensionality of the vector space containing the latent representations of the materials and properties and (2) $\lambda \in [0, \infty)$, the regularization parameter that trades off prediction accuracy on the training data and the complexity of the latent vector representations.

4 Case study: a COF recommendation system

We now demonstrate a material recommendation system based on a low rank matrix model. Here, the materials are COFs, and the properties are the equilibrium uptakes of a variety of gases at different conditions, obtained from molecular simulations. The Julia code to reproduce all of our work is available at github.com/SimonEnsemble/material_recommendation_system.

4.1 The dataset

We leverage an open data set of simulated gas adsorption properties in $M = 560$ experimentally reported, structurally optimized, and porous COF materials [18, 25]. We selected $P = 16$ simulated adsorption properties of interest, comprised of the uptake [units: mmol/g] and Henry coefficients [mmol/(g·bar)] of a variety of gases—CH₄, H₂O, H₂S, Xe, Kr, CO₂, N₂, O₂ and H₂—at various conditions that apply to different gas storage and separation applications. We log₁₀-transformed the Henry coefficients because of the relatively long tail of their distributions. The resulting COF-adsorption-property matrix, $\mathbf{A}_{\text{complete}} \in \mathbb{R}^{560 \times 16}$ is fully observed, allowing us to study the effect of the fraction of missing entries on the performance of the low rank model.

Fig. 3 displays the distribution of the [standardized] properties and the pairwise relationships between the adsorption properties (see Fig. S1 for a pairwise correlation matrix). Some properties are strongly correlated, e.g., CH₄ uptake at (298 K, 65 bar) and H₂ uptake at (77 K, 100 bar), while others, such as Xe and H₂O Henry coefficients, are not. The low rank model exploits these correlations between properties to learn low-dimensional representations of the materials and properties.

4.2 Simulating the process of data collection

We simulate the stochastic process of incomplete data collection to construct an incomplete COF-adsorption-property matrix $\mathbf{A}^{(\theta)}$ (still $M = 560 \times P = 16$) with a fraction θ of missing entries. We construct $\mathbf{A}^{(\theta)}$ by (uniform) randomly sampling, without replacement, $(1 - \theta)MP$ entries to ablate (change to missing) from the MP entries of $\mathbf{A}_{\text{complete}}$. Fig. 4 visualizes a resulting incomplete COF-adsorption-property matrix $\mathbf{A}^{(0.4)}$ with a fraction $\theta = 0.4$ missing entries.

4.3 Standardization of adsorption properties

We standardize the adsorption properties (the columns of $\mathbf{A}^{(\theta)}$) to have mean zero and unit variance using only the observed training examples. Standardization accounts for the different scales of the different properties and prevents properties with a larger variance from dominating the loss function in eqn. 3. See Ref. [24] for theoretical arguments for standardization. The entries in Fig. 4 are standardized, hence the diverging colormap.

4.4 Training, hyperparameter tuning, and testing

We use `LowRankModels.jl` [24] in the Julia programming language [78] to train our low rank models of the form in eqn. 2. `LowRankModels.jl` implements an alternating proximal gradient descent [24] to minimize the loss in eqn. 3.

For training and hyperparameter (k, λ) tuning, we randomly partitioned the [simulated] observed entries of $\mathbf{A}^{(\theta)}$ into an 80/20 % training/validation set. The loss ℓ in eqn. 3 is minimized over the




Figure 3: The distribution of (diagonal) and pairwise relationships between (off-diagonal) the simulated gas adsorption properties of the COFs (data from Ref. [18]). Each point represents a COF. Each property was standardized to have zero mean and unit variance.

training set, while the validation set is used to select optimal hyperparameters. The remaining, [simulated] unobserved entries serve as test data to estimate the generalization error of the low rank model for matrix completion.

To determine the optimal hyperparameter tuple $(k_{opt}^{(\theta)}, \lambda_{opt}^{(\theta)})$ for a given $\mathbf{A}^{(\theta)}$, we perform a hyperparameter sweep over a (k, λ) grid, training one low rank model for each (k, λ) . We select $(k_{opt}^{(\theta)}, \lambda_{opt}^{(\theta)})$ as the hyperparameter tuple whose low rank model produces the lowest approximation error over




Figure 4: The 560×16 COF-adsorption-property matrix $\mathbf{A}^{(0.4)}$ with $\theta = 0.4$ the fraction of observed entries. Each column (adsorption property) is standardized to have zero mean and unit variance. Unobserved entries are gray. The value of the observed entries are depicted by the colors in the colorbar.

the validation set. The grid is the Cartesian product of (1) $k \in \{1, 2, \dots, 15\}$ and (2) 25 values of λ ranging from 10 to 1000 and evenly spaced on a log-scale.

The *deployment* low rank model is then a new low rank model (a) with hyperparameter tuple $(k_{opt}^{(\theta)}, \lambda_{opt}^{(\theta)})$ and (b) trained on all [simulated] observed entries (train + validation data) of $\mathbf{A}^{(\theta)}$. We evaluate the performance of the deployment model by comparing its predictions of the missing entries to the actual values of the missing entries that comprise the test data. N.b., the loss and performance evaluation is done on the standardized and, in the case of Henry coefficients, \log_{10} -transformed values.

4.5 Results for observed fraction $\theta = 0.4$

We now demonstrate the utility of a low rank model, using a particular instance of a COF-adsorption-property matrix with a fraction $\theta = 0.4$ observed entries (shown in Figs. 4), for (i) imputing missing entries and (ii) drawing of a map of COFs and adsorption properties. Figs. 5-8 all correspond to the same deployment low rank model trained using the instance of $\mathbf{A}^{(0.4)}$ shown in Fig. 4, where the hyperparameter sweep found $k_{opt}^{(0.4)} = 3$, $\lambda_{opt}^{(0.4)} = 215.44$.

4.5.1 Imputing missing entries

We judge the performance of the low rank model for imputing the missing entries of the COF-adsorption-property matrix $\mathbf{A}^{(0.4)}$ by comparing the predictions of the missing entries to the actual values in the test data set, composed of the simulated unobserved entries.

The parity plot in Fig. 5 shows the joint distribution of predicted and actual values of the (standardized and, in the case of Henry coefficients, \log_{10} -transformed) adsorption properties in the test data set—the simulated unobserved entries of $\mathbf{A}^{(0.4)}$. The density is greatest along the diagonal line of equality, indicating that the recommendation system is providing predictive value. The RMSE and Spearman’s rank correlation coefficient on the test data is 0.6 and 0.77, respectively.

The ultimate utility of the recommendation system is to rank COFs according to specific properties (for specific applications). Spearman’s rank (here, a ranking of COFs) correlation coefficient, ρ , between the prediction of a missing adsorption property by the deployment low rank model and its actual value (from the test set) is shown for each adsorption property in Fig. 6. With the exception of H_2O Henry coefficients, the recommendation system ranks the COFs according to their properties reasonably well, with $\rho > 0.6$. The relatively poor ranking of COFs by H_2O Henry coefficient is explained by its very weak correlation with the other properties (see Fig. S1).

As a baseline to judge the performance of our recommendation system, we also train and test (on the same data) a benchmark model that excludes the interaction term $\mathbf{m}_m^T \mathbf{p}_p$ in eqn. 1. This material bias model, equivalent to the low rank model in eqn. 2 when $k = 0$, gives $A_{mp} \approx \mu_m$ and only considers whether the COF in question tends to exhibit high or low values of the properties (reflected in μ_m) when predicting A_{mp} . By comparing the imputation performance of this material bias model




Figure 5: A parity plot showing the joint distribution, over the test data set (simulated unobserved entries in $\mathbf{A}^{(0.4)}$), of (y-axis) the predictions of the missing adsorption properties by the deployment low rank model and (x-axis) the actual value of the missing entries. The diagonal line represents perfect prediction.

with the $k > 0$ low rank model, we quantify the extent to which the interactions between the COFs and the gas adsorption properties—encoded in $\mathbf{m}_m^T \mathbf{p}_p$ terms for $k > 0$ —are useful in the recommendation system for imputing the missing values. For each adsorption property, the stars in Fig. 6 show Spearman’s rank correlation coefficients between the value of the missing property (from the test set) and the prediction of the missing property by the benchmark material bias model. Indeed, the interaction term enhances the ability of the recommendation system to rank COFs according to their adsorption properties, though by different margins depending on the property. O_2 adsorption at (298 K, 5 bar) and N_2 adsorption at (300 K, 0.001 bar) are the two properties where the interaction term is playing only a marginal role. Overall, this indicates that our recommendation system is (i) learning interactions between COFs and the adsorption properties and (ii) more likely to suggest high-performing COFs for an application than a simpler strategy that selects COFs purely based on how they perform on average (as in the material bias model).

4.5.2 The COF biases

The learned material bias of COF m , $\mu_m \in \mathbb{R}$, in eqn. 1 roughly describes the typical value of the (standardized) gas adsorption properties of COF m . Visualization of $\boldsymbol{\mu}$ can give us an idea of which COFs tend to exhibit the largest and smallest values of the gas adsorption properties. Fig. 7 visualizes the extremes of $\boldsymbol{\mu}$ from the deployment low rank model trained on $\mathbf{A}^{(0.4)}$ and displays the COF structures with the lowest and highest material biases. CCOF-2 (COF-LZU8) has the largest (smallest) μ_m , indicating that CCOF-2 (COF-LZU8) tends to exhibit the highest (lowest) values of the (standardized) gas adsorption properties among the COFs.




Figure 6: For each adsorption property, the height of the bar shows Spearman’s rank correlation coefficient ρ between the prediction of a missing adsorption property in $\mathbf{A}^{(0.4)}$ (in the test set) by the deployment low rank model and its actual value. For comparison, the stars show ρ for the benchmark, material offset model where $A_{mp} \approx \mu_m$ and the interaction term is excluded.

4.5.3 The learned map of COFs and gas adsorption properties

The learned latent representation of COF m , $\mathbf{m}_m \in \mathbb{R}^k$, encodes its adsorption properties into a compressed, low-dimensional vector. The locations of the COF representations in the latent space of COFs provide a “map” of the COFs. Within this map, COFs with similar (dissimilar) adsorption properties congregate (separate). Similarly, $\mathbf{p}_p \in \mathbb{R}^k$ is a latent representation of gas adsorption property p . Because the interaction term in eqn. 1 is $\mathbf{m}_m^T \mathbf{p}_p$, the latent property vectors \mathbf{p}_p indicate which regions of latent space tend to contain COFs with high values of those adsorption properties.

To visualize the map of COFs and adsorption properties, we resort to a dimension reduction method, Uniform Manifold Approximation and Projection (UMAP) [79], which embeds the latent representations of the COFs and properties, contained in the columns of \mathbf{M} and columns of \mathbf{P} respectively, into a 2D space. N.b. we apply UMAP on the horizontally concatenated matrix $\mathbf{M} \parallel \mathbf{P}$, as opposed to \mathbf{M} and \mathbf{P} separately, so that the latent representations of the material and property vectors are comparable. Fig. 8 shows the map of COFs and adsorption properties from the deployment low rank model for $\mathbf{A}^{(0.4)}$.

In the map of COFs in Figs. 8a and 8b, each point represents a COF, colored by (a) CH₄ adsorption at (298 K, 65 bar) and (b) H₂S Henry coefficient at 300 K. Indeed, nearby COFs in this map tend to exhibit similar values of the adsorption property: COFs with the highest CH₄ uptake at (298 K, 65 bar) are located in the top right corner, while COFs with the highest H₂S Henry coefficients are located in the top left corner. Fig. S4 shows the COF map colored by the other adsorption properties.

The map of adsorption properties is displayed in Fig. 8c. Comparing Figs. 8a and 8b with the location of the latent property vector of (a) CH₄ adsorption at (298 K, 65 bar) and (b) H₂S Henry coefficient at 300 K reveals that the latent vectors of COFs with large (small) values of these properties tend to




Figure 7: The ranked COF biases, $\{\mu_m\}$, for the $\theta = 0.4$ deployment low rank model. The COFs with the lowest and highest μ_m are shown, with the corresponding top and bottom three COF structures visualized below.

be oriented in the same (opposite) direction as the latent property vectors—consistent with the interaction term in eqn. 1 as the dot product $\mathbf{m}_m^T \mathbf{p}_p$. [As a cautionary note, UMAP does not preserve angles.]

In summary, Fig. 8 illustrates that the recommendation system machine-learns, from incomplete data, a meaningful map of COFs, wherein COFs with similar adsorption properties congregate.

4.6 The effect of observed fraction θ on performance

Because the COF-adsorption-property $\mathbf{A}_{\text{complete}}$ from Ref. [18] is in reality complete, we have the luxury of studying the impact of the fraction of observed entries, θ , on the performance of the recommendation system. This investigation is important to address the practical question: how complete must the COF-adsorption-property matrix be for the recommendation system to reliably rank COFs according to their adsorption properties?

For a fraction of observed values $\theta \in \{0.1, 0.2, \dots, 0.9\}$, we sampled an ensemble of COF-adsorption property matrices $\mathbf{A}^{(\theta)}$ (50 simulations of data collection for each θ). For each instance of $\mathbf{A}^{(\theta)}$, we

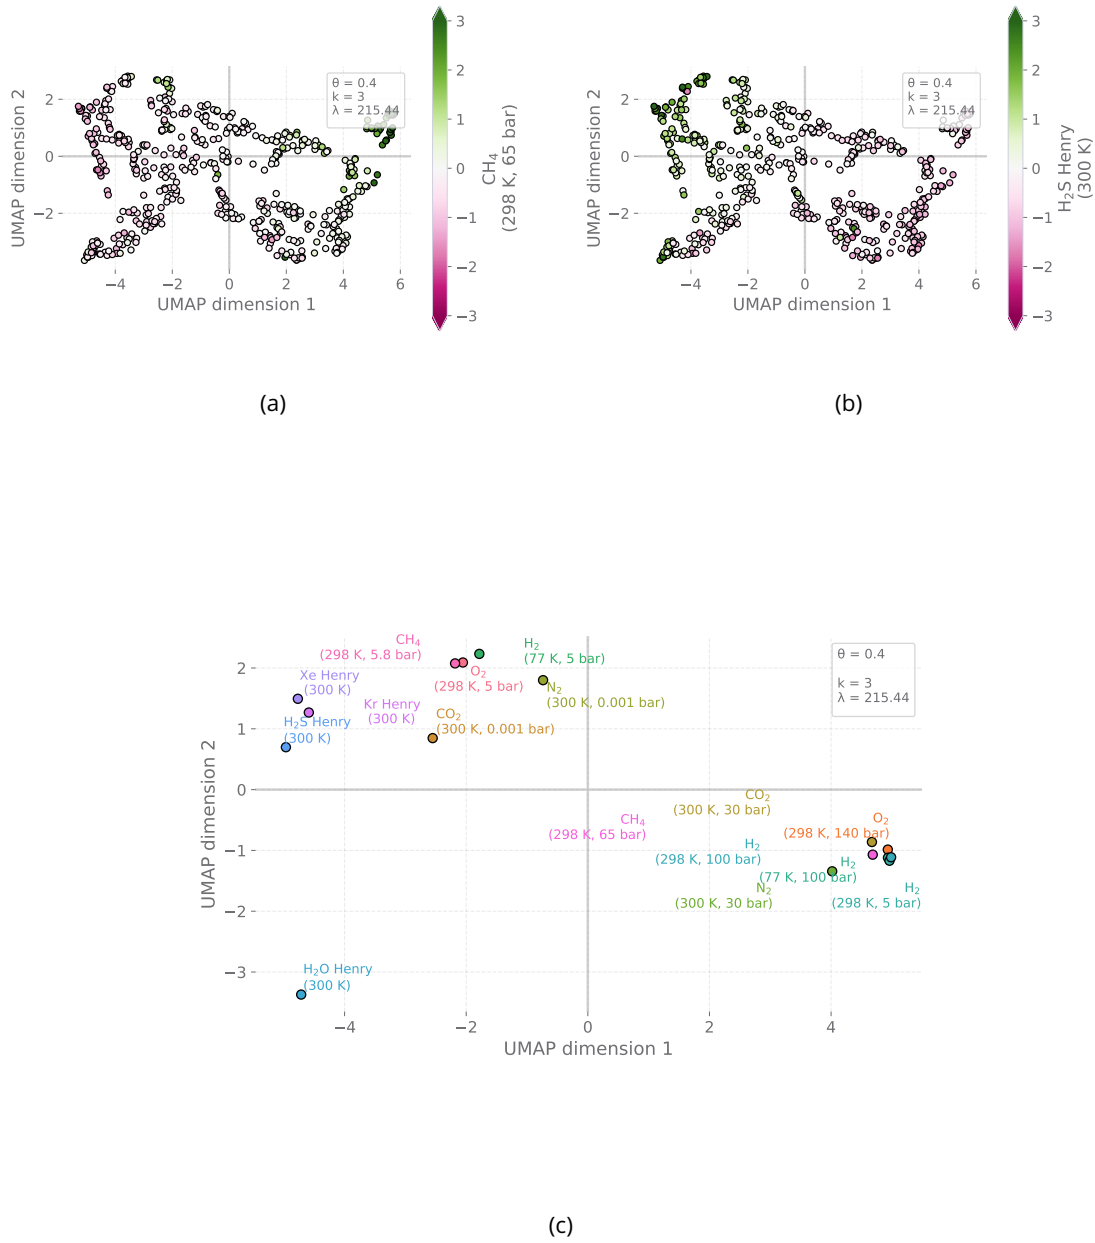


Figure 8: Learned map of COFs and gas adsorption properties: UMAP [79] embeddings of the latent representations of the (a, b) COFs, $\{\mathbf{m}_m\}$ and (b) adsorption properties, $\{\mathbf{p}_p\}$ into a 2D plane. (a, b) Each point represents a COF, colored by (a) CH_4 adsorption at (298 K, 65 bar), (b) H_2S Henry coefficients at 300 K. (c) Each point represents an adsorption property.

conducted a hyperparameter sweep using a training/validation split of the observed entries, retrained a deployment model on all observed entries, then tested the deployment model on the unobserved (missing) entries serving as test data. Fig. S2 shows the distribution (among the simulations of data collection) of optimal hyperparameters $(k_{\text{opt}}^{(\theta)}, \lambda_{\text{opt}}^{(\theta)})$ for each θ . Fig. 9 shows Spear-

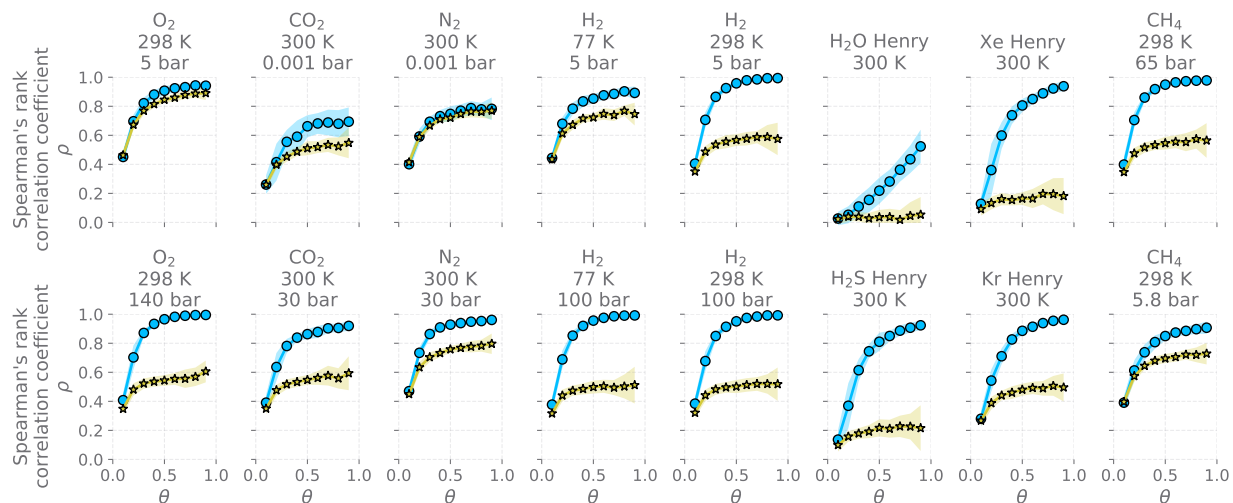


Figure 9: The effect of the fraction of observed values θ on the performance of the recommendation system for missing adsorption property imputation. Blue circles show the mean (over 50 simulations of data collection) Spearman’s rank correlation coefficient between the prediction of the missing adsorption property (from test data) and its true value, as a function of θ , and for each adsorption property. Yellow stars shows the mean Spearman’s rank correlation coefficient for the benchmark material bias model. Shaded bands signify the standard deviation.

man’s rank correlation coefficient, ρ , between the prediction of the missing adsorption property by the deployment low rank model and its actual value, for each adsorption property, as θ varies. The bands show the standard deviation over the 50 simulations of data collection. As in Fig. 6, ρ for the H₂O Henry coefficient is much lower than for other properties $\forall \theta$ owing to its poor correlation with the other properties. For the majority of the gas adsorption properties, the recommendation system ranks COFs according to the property reasonably well (i.e., a reasonably high ρ), until θ is reduced below $\theta = 0.4$, where the fraction of missing values is too large to provide accurate predictions owing to a paucity of training examples. In conclusion, at least 40 % of the values of the COF–adsorption-property matrix must be observed for the recommendation system to reliably rank COFs according to their adsorption properties. For comparison, ρ for the baseline material offset model is also shown in Fig. 6 as the dashed line; the interaction term provides significant predictive value, with the exception of N₂ adsorption at (300 K, 0.001 bar) and O₂ adsorption at (298 K, 5 bar).

5 Conclusion and Discussion

In materials science, we are often interested in many different properties of many different materials. The corresponding material-property matrix often, in practice, has many missing values, since every property of every material has not been measured. The idea of a material recommendation system is to leverage the observed (material, property) values to impute the missing ones. The (material, property) values are mathematically analogous to (item, customer) ratings in commercial

recommendation systems.

We demonstrated a COF recommendation system for different gas adsorption applications. Our COF-adsorption-property matrix was composed of the simulated uptake of several gases at different conditions in 560 COF structures by Ongari et al. [18]. We simulated the process of data observation by artificially introducing missing values into the matrix. The (simulated) unobserved entries served as test data to assess the performance of the data imputation by the recommendation system. To both (i) impute the missing adsorption properties and (ii) machine-learn a “map” of COFs, wherein COFs with similar adsorption properties congregate, we trained a low rank matrix model [24] of the COF-adsorption-property matrix that had missing entries. The recommendation system was able to rank COFs according to their adsorption properties reasonably well (Spearman’s rank correlation coefficient > 0.6), with the exception of water Henry coefficients. Moreover, coloring of the learned map of COFs by the adsorption properties indicated that, indeed, COFs with similar (dissimilar) adsorption properties clustered together (separated) in the map. The imputation performance of the recommendation system precipitously drops once the fraction of missing entries exceeds 60 %, though this figure does not necessarily generalize to other data sets.

We conclude that material recommendation systems, if sufficient training data is available, could be widely useful for leveraging measured properties of materials to fill in missing measurements. In turn, this could accelerate the matching of materials for specific applications.

The success of a recommendation system for NPMs is, however, predicated on structured, open databases of NPMs and their adsorption properties. One such database is the NIST/ARPA-E Database of Novel and Emerging Adsorbent Materials [68] (NIST-ISODB) that has collected and compiled gas adsorption measurements in NPMs from the literature, for both experimental and simulation sources, for a variety of gases at a wide range of conditions. We originally set out to develop a recommendation system using Henry coefficients extracted from this experimental data, but we found the resultant recommendation system was unable to reliably rank NPMs according to their adsorption properties. Particularly, we found the recommendation system with an interaction term included could not outperform the baseline material offset model. We propose four explanations for the poor performance of our recommendation system based on an NPM-adsorption-property matrix from NIST-ISODB; the explanations include both data-centric and model-centric concerns. First, the Henry coefficient matrix we constructed based on NIST-ISODB was only $\sim 20\%$ complete (i.e., too many missing values), which may have limited the success of recommendations for the remainder of the matrix. (Recall that the COF-adsorption-property matrix needed to be at least 40 % complete to satisfactorily make recommendations for the remainder of the matrix. This benchmark for the COF recommender system may not generalize to a broad set of NPMs and properties, but is nonetheless informative.) Second, the NPM-adsorption-property matrix may have included too much noise for a successful recommendation. It is known from meta-analyses of isotherms catalogued in NIST-ISODB [80, 81] that experimentally measured gas adsorption isotherms in NPMs exhibit high variance; this variance is ultimately manifested as noise that limits the success of the recommendation system. Third, the accuracy and reliability of Henry coefficients obtained from isotherms in NIST-ISODB are naturally limited by the source of data in the database itself. NIST-ISODB is primarily constructed from manual extraction of isotherm data from graphical figures in literature articles, since it is not common practice in the adsorption community to provide gas ad-

sorption measurements as raw tabular data in publications. (As of the time of writing this work, only 1.3 % of isotherms in NIST-ISODB were from tabular data sources.) Consequently, the adsorption isotherm data loses precision first when the data is plotted graphically and second by human error when the graphical figure is digitized back to numerical data. This loss of precision particularly affects the generation of Henry coefficients from figures, as those coefficients are especially dependent on low-pressure data, which is often difficult to extract from isotherms plotted on a linear pressure scale. Fourth, our low rank model in eqn. 1 is linear; a non-linear model may be able to capture relationships between the adsorption properties and achieve better performance. The third issue can be addressed by community adoption of the practice of releasing raw adsorption isotherm data in standardized, structured formats (cf. the CIF standard for crystallography [82]), which has been discussed previously [28, 83, 84].

We introduced missing entries in the (in reality, fully observed) COF-adsorption-property matrix by (uniform) randomly selecting entries to ablate. In practice, however, (i) some properties are more commonly measured than others, (ii) some materials are more commonly studied than others owing to e.g. ease of synthesis, and (iii) there are likely correlations between and temporal trends with the binary random variables that represent whether the (material, property) values are observed. To expand on (iii), for example, a material with a superior (inferior) value of a desired property may become popular (unpopular) for measurements of other properties. Future work entails (a) creating a model for the selection bias in selecting materials for measurements of properties and (b) accounting for the selection bias in the recommendation system [85].

Another interesting direction for future work is to determine what (material, property) measurements should be made next to most improve the recommendation system, in an active learning strategy [86].

As a remark, recommendation systems suffer from the *cold start* [20] problem: if a new material is reported, but none of its properties have been observed, the recommendation system is unable to make a prediction about any of its properties. To learn the latent material vector of this material, \mathbf{m}_{M+1} , using the loss in eqn. 3, we must have an observation of at least one property of the new material.

To (i) improve the performance of the recommendation system and (ii) alleviate the cold start problem, we propose to include structural and chemical properties of the materials that contribute to the prediction, in addition to the observed adsorption properties. For example, we could include in the model other information about the NPM structures, such as the void fraction, surface area, percent carbon atoms, etc. In the analogy with movie recommendation systems, this is analogous to including features about the movies, such as the genre, directors, year of production, and actors. These features could be added as additional (fully observed) columns in the material-property matrix, \mathbf{A} .

The material recommendation system is practically useful for recommending (i, application-led material search [1]) a material that optimizes a specific property or (ii, material-led application search [18]) an application for a given material. To motivate an experimental measurement in the lab, it may be necessary to quantify the uncertainty associated with a property imputed by the recommendation system. We remark that one could achieve this through bootstrapping and training an ensemble of recommendation systems on the bootstrap samples of observed (material, property)

values.

Acknowledgements

A.S., X.F., and C.M.S. acknowledge the National Science Foundation for support under grant No. 1920945.

Copyright and Disclaimer

Official contribution of the National Institute of Standards and Technology; not subject to copyright in the United States. Certain commercially available items may be identified in this paper. This identification does not imply recommendation by NIST, nor does it imply that it is the best available for the purposes described.

References

- [1] Anna G Slater and Andrew I Cooper. Porous materials. function-led design of new porous materials. *Science*, 348(6238):aaa8075–aaa8075, 2015.
- [2] Omar K. Farha, Ibrahim Eryazici, Nak Cheon Jeong, Brad G. Hauser, Christopher E. Wilmer, Amy A. Sarjeant, Randall Q. Snurr, SonBinh T. Nguyen, A. Özgür Yazaydın, and Joseph T. Hupp. Metal–organic framework materials with ultrahigh surface areas: Is the sky the limit? *Journal of the American Chemical Society*, 134(36):15016–15021, August 2012.
- [3] Russell E Morris and Paul S Wheatley. Gas storage in nanoporous materials. *Angewandte Chemie International Edition*, 47(27):4966–4981, 2008.
- [4] Jian-Rong Li, Ryan J Kuppler, and Hong-Cai Zhou. Selective gas adsorption and separation in metal–organic frameworks. *Chemical Society Reviews*, 38(5):1477–1504, 2009.
- [5] Alexander Schoedel, Zhe Ji, and Omar M Yaghi. The role of metal–organic frameworks in a carbon-neutral energy cycle. *Nature Energy*, 1(4):1–13, 2016.
- [6] Zoey R Herm, Eric D Bloch, and Jeffrey R Long. Hydrocarbon separations in metal–organic frameworks. *Chemistry of Materials*, 26(1):323–338, 2013.
- [7] Lauren E Kreno, Kirsty Leong, Omar K Farha, Mark Allendorf, Richard P Van Duyne, and Joseph T Hupp. Metal–organic framework materials as chemical sensors. *Chemical Reviews*, 112(2):1105–1125, 2012.
- [8] David Farrusseng, Sonia Aguado, and Catherine Pinel. Metal–organic frameworks: opportunities for catalysis. *Angewandte Chemie International Edition*, 48(41):7502–7513, 2009.

- [9] Hiroyasu Furukawa, Kyle E Cordova, Michael O’Keeffe, and Omar M Yaghi. The chemistry and applications of metal–organic frameworks. *Science*, 341(6149):1230444, 2013.
- [10] Christian S Diercks and Omar M Yaghi. The atom, the molecule, and the covalent organic framework. *Science*, 355(6328), 2017.
- [11] Weigang Lu, Daqiang Yuan, Dan Zhao, Christine Inge Schilling, Oliver Plietzsch, Thierry Muller, Stefan Brase, Johannes Guenther, Janet Blumel, Rajamani Krishna, Zhen Li, and Hong-Cai Zhou. Porous polymer networks: Synthesis, porosity, and applications in gas storage/separation. *Chemistry of Materials*, 22(21):5964–5972, November 2010.
- [12] Andrew I Cooper. Porous molecular solids and liquids. *ACS Central Science*, 3(6):544–553, 2017.
- [13] Eric J Gosselin, Casey A Rowland, and Eric D Bloch. Permanently microporous metal–organic polyhedra. *Chemical Reviews*, 2020.
- [14] David J Tranchemontagne, Zheng Ni, Michael O’Keeffe, and Omar M Yaghi. Reticular chemistry of metal–organic polyhedra. *Angewandte Chemie International Edition*, 47(28):5136–5147, 2008.
- [15] Sukhendu Mandal, Srinivasan Natarajan, Prabu Mani, and Asha Pankajakshan. Post-synthetic modification of metal–organic frameworks toward applications. *Advanced Functional Materials*, page 2006291, 2020.
- [16] Peter G Boyd, Yongjin Lee, and Berend Smit. Computational development of the nanoporous materials genome. *Nature Reviews Materials*, 2(8):1–15, 2017.
- [17] Connor W Coley. Defining and exploring chemical spaces. *Trends in Chemistry*, 2020.
- [18] Daniele Ongari, Leopold Talirz, and Berend Smit. Too many materials and too many applications: An experimental problem waiting for a computational solution. *ACS Central Science*, 6(11):1890–1900, 2020.
- [19] Paul Iacomi and Philip L Llewellyn. Data mining for binary separation materials in published adsorption isotherms. *Chemistry of Materials*, 32(3):982–991, 2020.
- [20] Charu C Aggarwal. *Recommender systems*. Springer, 2016.
- [21] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [22] Madeleine Udell and Alex Townsend. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science*, 1(1):144–160, 2019.
- [23] Madeleine Udell. Big data is low rank. *SIAG/OPT Views and News*, 2019.
- [24] Madeleine Udell, Corinne Horn, Reza Zadeh, and Stephen Boyd. Generalized low rank models. *arXiv preprint arXiv:1410.0342*, 2014.
- [25] Daniele Ongari, Aliaksandr V Yakutovich, Leopold Talirz, and Berend Smit. Building a consistent and reproducible database for adsorption evaluation in covalent–organic frameworks. *ACS Central Science*, 5(10):1663–1675, 2019.

- [26] Eun Hyun Cho, Qiang Lyu, and Li-Chiang Lin. Computational discovery of nanoporous materials for energy-and environment-related applications. *Molecular Simulation*, 45(14-15):1122–1147, 2019.
- [27] Jack D Evans, Guillaume Fraux, Romain Gaillac, Daniela Kohen, Fabien Trouselet, Jean-Mathieu Vanson, and François-Xavier Coudert. Computational chemistry methods for nanoporous materials. *Chemistry of Materials*, 29(1):199–212, 2016.
- [28] Arni Sturluson, Melanie T Huynh, Alec R Kaija, Caleb Laird, Sunghyun Yoon, Feier Hou, Zhenxing Feng, Christopher E Wilmer, Yamil J Colón, Yongchul G Chung, et al. The role of molecular modelling and simulation in the discovery and deployment of metal-organic frameworks for gas storage and separation. *Molecular Simulation*, 45(14-15):1082–1121, 2019.
- [29] Hilal Daglar and Seda Keskin. Recent advances, opportunities, and challenges in high-throughput computational screening of MOFs for gas separations. *Coordination Chemistry Reviews*, 422:213470, 2020.
- [30] Yongchul G. Chung, Diego A. Gómez-Gualdrón, Peng Li, Karson T. Leperi, Pravas Deria, Hongda Zhang, Nicolaas A. Vermeulen, J. Fraser Stoddart, Fengqi You, Joseph T. Hupp, Omar K. Farha, and Randall Q. Snurr. In silico discovery of metal-organic frameworks for precombustion CO₂ capture using a genetic algorithm. *Science Advances*, 2(10):e1600909, 2016.
- [31] Yi Bao, Richard L Martin, Cory M Simon, Maciej Haranczyk, Berend Smit, and Michael W Deem. In silico discovery of high deliverable capacity metal-organic frameworks. *The Journal of Physical Chemistry C*, 119(1):186–195, 2015.
- [32] Sean P. Collins, Thomas D. Daff, Sarah S. Piotrkowski, and Tom K. Woo. Materials design by evolutionary optimization of functional groups in metal-organic frameworks. *Science Advances*, 2(11), 2016.
- [33] Xiangyu Zhang, Kexin Zhang, and Yongjin Lee. Machine learning enabled tailor-made design of application-specific metal-organic frameworks. *ACS Applied Materials & Interfaces*, 12(1):734–743, 2019.
- [34] Michael Fernandez, Tom K Woo, Christopher E Wilmer, and Randall Q Snurr. Large-scale quantitative structure–property relationship (QSPR) analysis of methane storage in metal-organic frameworks. *The Journal of Physical Chemistry C*, 117(15):7681–7689, 2013.
- [35] Benjamin J Bucior, N Scott Bobbitt, Timur Islamoglu, Subhadip Goswami, Arun Gopalan, Taner Yildirim, Omar K Farha, Neda Bagheri, and Randall Q Snurr. Energy-based descriptors to rapidly predict hydrogen storage in metal-organic frameworks. *Molecular Systems Design & Engineering*, 4(1):162–174, 2019.
- [36] Eun Hyun Cho, Xuepeng Deng, Changlong Zou, and Li-Chiang Lin. Machine learning-aided computational study of metal-organic frameworks for sour gas sweetening. *The Journal of Physical Chemistry C*, 2020.

- [37] Cory M Simon, Rocio Mercado, Sondre K Schnell, Berend Smit, and Maciej Haranczyk. What are the best materials to separate a xenon/krypton mixture? *Chemistry of Materials*, 27(12):4459–4475, 2015.
- [38] Ryther Anderson, Jacob Rodgers, Edwin Argueta, Achay Biong, and Diego A. Gómez-Gualdrón. Role of pore chemistry and topology in the CO₂ capture capabilities of MOFs: From molecular simulation to machine learning. *Chemistry of Materials*, 30(18):6325–6337, August 2018.
- [39] Giorgos Borboudakis, Taxiarchis Stergiannakos, Maria Frysali, Emmanuel Klontzas, Ioannis Tsamardinos, and George E Froudakis. Chemically intuited, large-scale screening of mofs by machine learning techniques. *npj Computational Materials*, 3(1):40, 2017.
- [40] Maryam Pardakhti, Ehsan Moharrer, David Wanik, Steven L Suib, and Ranjan Srivastava. Machine learning using combined structural and chemical descriptors for prediction of methane adsorption performance of metal organic frameworks (MOFs). *ACS Combinatorial Science*, 19(10):640–645, 2017.
- [41] Eun Hyun Cho and Li-Chiang Lin. Nanoporous material recognition via 3D convolutional neural networks: Prediction of adsorption properties. *The Journal of Physical Chemistry Letters*, pages 2279–2285, March 2021.
- [42] Seyed Mohamad Moosavi, Kevin Maik Jablonka, and Berend Smit. The role of machine learning in the understanding and design of materials. *Journal of the American Chemical Society*, 142(48):20273–20287, 2020.
- [43] Steven Bennett, Andrew Tarzia, Martijn A Zwijnenburg, and Kim E Jelfs. Artificial intelligence applied to the prediction of organic materials. *Machine Learning in Chemistry*, 17:280, 2020.
- [44] Siwar Chibani and François-Xavier Coudert. Machine learning approaches for the prediction of materials properties. *APL Materials*, 8(8):080701, 2020.
- [45] Sanggyu Chong, Sangwon Lee, Baekjun Kim, and Jihan Kim. Applications of machine learning in metal-organic frameworks. *Coordination Chemistry Reviews*, 423:213487, 2020.
- [46] Zenan Shi, Wenyan Yang, Xiaomei Deng, Chengzhi Cai, Yaling Yan, Hong Liang, Zili Liu, and Zhiwei Qiao. Machine-learning-assisted high-throughput computational screening of high performance metal-organic frameworks. *Molecular Systems Design & Engineering*, 5(4):725–742, 2020.
- [47] Kevin Maik Jablonka, Daniele Ongari, Seyed Mohamad Moosavi, and Berend Smit. Big-data science in porous materials: Materials genomics and machine learning. *Chemical Reviews*, 2020.
- [48] Seyed Mohamad Moosavi, Aditya Nandy, Kevin Maik Jablonka, Daniele Ongari, Jon Paul Janet, Peter G. Boyd, Yongjin Lee, Berend Smit, and Heather J. Kulik. Understanding the diversity of the metal-organic framework ecosystem. *Nature Communications*, 11(1), August 2020.
- [49] Ruimin Ma, Yamil J. Colón, and Tengfei Luo. Transfer learning study of gas adsorption in metal-organic frameworks. *ACS Applied Materials & Interfaces*, 12(30):34041–34048, July 2020.

- [50] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272. PMLR, 2017.
- [51] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120(14):145301, 2018.
- [52] Ali Raza, Faaiq Waqar, Arni Sturluson, Cory Simon, and Xiaoli Fern. Towards explainable message passing networks for predicting carbon dioxide adsorption in metal-organic frameworks. *arXiv preprint arXiv:2012.03723*, 2020.
- [53] Arni Sturluson, Melanie T Huynh, Arthur HP York, and Cory M Simon. Eigencages: Learning a latent space of porous cage molecules. *ACS Central Science*, 4(12):1663–1676, 2018.
- [54] Thomas C Nicholas, Andrew L Goodwin, and Volker L Deringer. Understanding the geometric diversity of inorganic and hybrid frameworks through structural coarse-graining. *Chemical Science*, 46, 2020.
- [55] Yongjin Lee, Senja D. Barthel, Paweł Dłotko, S. Mohamad Moosavi, Kathryn Hess, and Berend Smit. Quantifying similarity of pore-geometry in nanoporous materials. *Nature Communications*, 8(1), May 2017.
- [56] Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.
- [57] Zhenpeng Yao, Benjamín Sánchez-Lengeling, N. Scott Bobbitt, Benjamin J. Bucior, Sai Govind Hari Kumar, Sean P. Collins, Thomas Burns, Tom K. Woo, Omar K. Farha, Randall Q. Snurr, and Alán Aspuru-Guzik. Inverse design of nanoporous crystalline reticular materials with deep generative models. *Nature Machine Intelligence*, January 2021.
- [58] R. Pollice, Gabriel dos Passos Gomes, Matteo Aldeghi, Riley J. Hickman, Mario Krenn, Cyrille Lavigne, Michael Lindner-D’Addario, AkshatKumar Nigam, Cher Tian Ser, Zhenpengand Yao, and Alan Aspuru-Guzik. Data-driven strategies for accelerated materials design. *Accounts of Chemical Research*, 54(4):849–860, 2021.
- [59] Guillaume Fraux, Siwar Chibani, and François-Xavier Coudert. Modelling of framework materials at multiple scales: current practices and open questions. *Philosophical Transactions of the Royal Society A*, 377(2149):20180220, 2019.
- [60] Leopold Talirz, Snehal Kumbhar, Elsa Passaro, Aliaksandr V. Yakutovich, Valeria Granata, Fernando Gargiulo, Marco Borelli, Martin Uhrin, Sebastiaan P. Huber, Spyros Zoupanos, Carl S. Adorf, Casper Welzel Andersen, Ole Schütt, Carlo A. Pignedoli, Daniele Passerone, Joost Van-deVondele, Thomas C. Schulthess, Berend Smit, Giovanni Pizzi, and Nicola Marzari. Materials cloud, a platform for open computational science. *Scientific Data*, 7(1), September 2020.
- [61] François-Xavier Coudert. Materials databases: the need for open, interoperable databases with standardized data and rich metadata. *Advanced Theory and Simulations*, 2(11):1900131, 2019.

- [62] Yongchul G Chung, Jeffrey Camp, Maciej Haranczyk, Benjamin J Sikora, Wojciech Bury, Vaiva Krungleviciute, Taner Yildirim, Omar K Farha, David S Sholl, and Randall Q Snurr. Computation-ready, experimental metal-organic frameworks: A tool to enable high-throughput screening of nanoporous crystals. *Chemistry of Materials*, 26(21):6185–6192, 2014.
- [63] Yongchul G Chung, Emmanuel Haldoupis, Benjamin J Bucior, Maciej Haranczyk, Seulchan Lee, Hongda Zhang, Konstantinos D Vogiatzis, Marija Milisavljevic, Sanliang Ling, Jeffrey S Camp, et al. Advances, updates, and analytics for the computation-ready, experimental metal-organic framework database: CoRE MOF 2019. *Journal of Chemical & Engineering Data*, 2019.
- [64] Minman Tong, Youshi Lan, Qingyuan Yang, and Chongli Zhong. Exploring the structure-property relationships of covalent organic frameworks for noble gas separations. *Chemical Engineering Science*, 168:456–464, 2017.
- [65] Peter G. Boyd, Arunraj Chidambaram, Enrique García-Díez, Christopher P. Ireland, Thomas D. Daff, Richard Bounds, Andrzej Gładysiak, Pascal Schouwink, Seyed Mohamad Moosavi, M. Mercedes Maroto-Valer, Jeffrey A. Reimer, Jorge A. R. Navarro, Tom K. Woo, Susana Garcia, Kyriakos C. Stylianou, and Berend Smit. Data-driven design of metal-organic frameworks for wet flue gas CO₂ capture. *Nature*, 576(7786):253–256, December 2019.
- [66] Peyman Z Moghadam, Aurelia Li, Seth B Wiggin, Andi Tao, Andrew GP Maloney, Peter A Wood, Suzanna C Ward, and David Fairen-Jimenez. Development of a cambridge structural database subset: a collection of metal-organic frameworks for past, present, and future. *Chemistry of Materials*, 29(7):2618–2625, 2017.
- [67] Rocío Mercado, Rueih-Sheng Fu, Aliaksandr V Yakutovich, Leopold Talirz, Maciej Haranczyk, and Berend Smit. In silico design of 2D and 3D covalent organic frameworks for methane storage applications. *Chemistry of Materials*, 30(15):5069–5086, 2018.
- [68] D.W. Siderius, V.K. Shen, R.D. Johnson III, and R.D. van Zee, editors. *NIST/ARPA-E Database of Novel and Emerging Adsorbent Materials*. National Institute of Standards and Technology, Gaithersburg, MD, 20899, 2014.
- [69] Andrew Rosen, Shaelyn Iyer, Debmalya Ray, Zhenpeng Yao, Alan Aspuru-Guzik, Laura Gagliardi, Justin Notestein, and Randall Q Snurr. Machine learning the quantum-chemical properties of metal-organic frameworks for accelerated materials discovery with a new electronic structure database. *ChemRxiv*, 2020.
- [70] Dalar Nazarian, Jeffrey S Camp, and David S Sholl. A comprehensive set of high-quality point charges for simulations of metal-organic frameworks. *Chemistry of Materials*, 28(3):785–793, 2016.
- [71] Thomas F. Willems, Chris H. Rycroft, Michael Kazi, Juan C. Meza, and Maciej Haranczyk. Algorithms and tools for high-throughput geometry-based analysis of crystalline porous materials. *Microporous and Mesoporous Materials*, 149(1):134–141, February 2012.

- [72] Christopher R Collins, Geoffrey J Gordon, O Anatole Von Lilienfeld, and David J Yaron. Constant size descriptors for accurate machine learning models of molecular properties. *The Journal of Chemical Physics*, 148(24):241718, 2018.
- [73] Aditi S. Krishnapriyan, Maciej Haranczyk, and Dmitriy Morozov. Topological descriptors help predict guest adsorption in nanoporous materials. *The Journal of Physical Chemistry C*, 124(17):9360–9368, April 2020.
- [74] Qi Yuana, Mariagiulia Longob, Aaron Thorntonc, Neil B McKeownd, Bibiana Comesaña-Gándarad, Johannes C Jansenb, and Kim E Jelfsa. Imputation of missing gas permeability data for polymer membranes using machine learning. *Journal of Membrane Science*, page 119207, 2021.
- [75] Ekaterina A Sosnina, Sergey Sosnin, Anastasia A Nikitina, Ivan Nazarov, Dmitry I Osolodkin, and Maxim V Fedorov. Recommender systems in antiviral drug discovery. *ACS Omega*, 5(25):15039–15051, 2020.
- [76] Atsuto Seko, Hiroyuki Hayashi, Hisashi Kashima, and Isao Tanaka. Matrix-and tensor-based recommender systems for the discovery of currently unknown inorganic compounds. *Physical Review Materials*, 2(1):013805, 2018.
- [77] Hiroyuki Hayashi, Katsuyuki Hayashi, Keita Kouzai, Atsuto Seko, and Isao Tanaka. Recommender system of successful processing conditions for new compounds based on a parallel experimental data set. *Chemistry of Materials*, 31(24):9984–9992, November 2019.
- [78] Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.
- [79] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [80] Jongwoo Park, Joshua D. Howe, and David S. Sholl. How reproducible are isotherm measurements in metal–organic frameworks? *Chemistry of Materials*, 29(24):10487–10495, 2017.
- [81] Lukas W. Bingel, Andrew Chen, Mayank Agrawal, and David S. Sholl. Experimentally verified alcohol adsorption isotherms in nanoporous materials from literature meta-analysis. *Journal of Chemical & Engineering Data*, 65(10):4970–4979, 2020.
- [82] S. R. Hall, F. H. Allen, and I. D. Brown. The crystallographic information file (CIF): a new standard archive file for crystallography. *Acta Crystallographica Section A*, 47(6):655–685, 1991.
- [83] N. Scott Bobbitt, Benjamin J. Bucior, Haoyuan Chen, Nathaniel Tracy-Amoroso, Zhao Li, Daniel W. Siderius, and Randall Q. Snurr. MOFdb: An accessible online database of computational adsorption data for nanoporous materials. *Journal of Chemical & Engineering Data*, 2021. In preparation.
- [84] Jack D Evans, Volodymyr Bon, Irena Senkovska, and Stefan Kaskel. A universal standard archive file for adsorption data. *ChemRxiv*, 2021.

- [85] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. Bias and debias in recommender system: A survey and future directions. *arXiv preprint arXiv:2010.03240*, 2020.
- [86] Neil Rubens, Dain Kaplan, and Masashi Sugiyama. Active learning in recommender systems. In P.B. Kantor, F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*, pages 735–767. Springer, 2011.