

Recommendation System to Predict Missing Adsorption Properties of Nanoporous Materials

Arni Sturluson, Ali Raza, Grant D. McConachie, Daniel W. Siderius, Xiaoli Z. Fern,* and Cory M. Simon*



Cite This: *Chem. Mater.* 2021, 33, 7203–7216



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

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 relevant adsorption properties in candidate NPMs, would enable the matching of NPMs with adsorption tasks. However, in practice, the NPM-property matrix is only partially observed (incomplete); many different properties of many different NPMs have not been measured. 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 product–customer ratings matrix to recommend products 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 constructed from simulated uptakes of CH₄, H₂O, H₂S, Xe, Kr, CO₂, N₂, O₂, and H₂ at various conditions. A low-rank model of the COF–adsorption-property matrix, fit to the observed (COF, adsorption property) values, provides (i) predictions of the missing (COF, adsorption property) values and (ii) a “map” of COFs, wherein COFs, represented as points, with similar (dissimilar) adsorption properties congregate (separate). The COF recommendation system is able to rank COFs reasonably well for most of the adsorption properties, but imputation performance diminishes precipitously when the fraction of missing entries exceeds 60%. The concepts in our COF recommendation system can be applied broadly to impute missing data pertaining to many different materials and properties.



1. INTRODUCTION

Nanoporous materials (NPMs)¹ often exhibit permanent porosity and possess large-area internal surfaces² decorated with functional groups. This enables them to (selectively) adsorb and concentrate gases in their pores.^{3–5} As a result, NPMs have applications in storing,⁵ separating,^{4,6} and sensing⁷ 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 compatible building blocks within many topologies, together with postsynthetic modifiability,¹⁵ makes the number of possible NPM structures extremely large.¹⁶

Thus, we have a large list of candidate NPMs and a list of their adsorption properties we wish to know for their many applications. If this NPM-property data table were complete, both searching for (i) the optimal NPM for a given application¹ and (ii) the optimal application for a given NPM¹⁷ would be trivial look-up problems. However, in

practice, the NPM-property data table, whether constructed from experimental adsorption measurements^{18,19} or molecular simulations of gas adsorption in libraries of NPMs,¹⁷ 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 NPMs (see Figure 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 in, or complete, the NPM-property matrix. A machine learning strategy to complete the NPM-property matrix is much less expensive and time consuming than experimentally measuring or computationally simulating these

Received: April 6, 2021

Published: September 7, 2021






Figure 1. Recommendation system for nanoporous materials (NPMs). In this toy NPM–adsorption-property matrix, entry (m, p) represents the value of adsorption property p of NPM m . Many entries are unobserved (“?”) because measurements are missing. The goal of our NPM recommendation system is to use the observed entries (values depicted by color) 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 make product recommendations tailored for specific customers, with NPM : adsorption property :: product : customer.

missing properties. The machine-completed NPM-property matrix is valuable because it can be used to direct higher fidelity but more expensive (experimental or simulated) measurements toward the most promising materials, thereby using less resources in the search for the optimal NPM for a given application.

Our hypothesis, which would permit accurate matrix completion, is that the NPM-property matrix exhibits a low-rank structure,^{20,21} owing to underlying structural and chemical similarities among both NPMs and gases that dictate their interactions. Specifically, we assume a low-rank structure where 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 jointly machine learned using the observed (NPM, property) values, then used to impute the missing ones.^{22,23}

Our goal to “fill in” the missing values in the NPM-property matrix—primarily to recommend NPMs for specific adsorption tasks—is analogous to the goal of a commercial recommendation system that recommends products to customers. For example, consider a movie recommendation system at Netflix.²³ Movie ratings by Netflix users are stored in a movie-user rating matrix (rows: movies, columns: users, entries: ratings).²³ The movie-user ratings matrix is incomplete; most entries are missing because (i) each user 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 leverages the observed (movie, user) ratings (perhaps, in addition to features of the movies and users) to impute the missing ones.²⁴ The machine-completed movie-user rating matrix is then used to make user-specific recommendations of movies. Thus, the (material, property) values in our material recommendation system are analogous to (product, customer) ratings in commercial recommendation systems.

Herein, we demonstrate a prototype recommendation system, based on a low-rank matrix model,^{22,23} that recommends COFs for various gas adsorption tasks. The COF–gas-adsorption-property matrices pertain to 572 experimentally reported COFs²⁵ and the simulated adsorption of CH₄, H₂O, H₂S, Xe, Kr, CO₂, N₂, O₂, and H₂ in those COFs at various conditions¹⁷ relevant to different gas storage and separation applications. Advantageously, this COF–gas-adsorption-property matrix is in reality complete, allowing us to ablate different fractions of the entries and investigate how imputation performance depends on the fraction of missing values. From the observed (COF, gas adsorption) values, we machine-learn a low-rank model of the COF–gas-adsorption-property matrix, yielding low-dimensional latent vector representations of both COFs and adsorption properties. This low-rank model provides: (i) predictions of the missing (COF, gas adsorption property) values and (ii) a “map,” wherein COFs are represented as points and COFs with similar (dissimilar) adsorption properties congregate (separate). Such a map of COFs is useful for experimental design to explore COF space and for the optimization of promising but still suboptimal “lead”-COFs.

1.1. Review of Previous Work. Machine learning plays an important role in the discovery and deployment of NPMs.^{26–33} Supervised machine learning models have been widely used to predict the adsorption properties of NPMs^{34–43} from vectors of hand-crafted structural features^{44,45} or graph representations.^{46,47} Unsupervised machine learning methods have been used to embed NPMs into a low-dimensional “material space”⁴⁸ and cluster together NPMs with similar structures.^{49–52} Genetic algorithms,^{53–56} Monte Carlo tree search,⁵⁷ and Bayesian optimization^{58,59} have been used to more efficiently search for the NPM(s) with an optimal adsorption property. Finally, recently⁶⁰ an autoencoder enabled inverse design^{61,62} of NPMs, where one specifies a desired adsorption property, and the machine learning model generates a NPM structure with that property. To enable machine learning approaches to NPM discovery, several open, structured databases^{63–65} of (i) crystal structure models of NPMs,^{25,66–70} (ii) simulated^{17,69,71,72} and experimentally measured¹⁹ adsorption properties of NPMs, and (iii) electronic properties of NPMs^{73,74} have been curated. Text mining and natural language processing could be used to extract data and knowledge from the literature for machine learning studies as well.^{75–77}

Our material recommendation system deviates from previous data-driven approaches to predict properties of NPMs by: (i) as a latent variable model, embedding materials into a latent space, negating the need for explicit, hand-crafted features of NPMs, and (ii) performing multitask prediction while (a) transferring knowledge between tasks and (b) handling missing values in the target vectors associated with NPMs. Loosely related, meta-learning has been used to predict an adsorption property of materials at different conditions by learning an intermediate representation of the material based only on available adsorption data.⁷⁸ N.b., recommendation systems have been built for use in chemical sciences to impute missing gas permeabilities in polymers,⁷⁹ antiviral activities of molecules,⁸⁰ and stabilities of inorganic materials.^{81,82}

2. 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 missing (material, property) values and (ii) draw a map of the materials, wherein materials with similar properties congregate.

2.1. Material–Property Matrix. Suppose we have M candidate materials with P properties of interest. Entry (m, p) of the $M \times P$ material–property matrix $\mathbf{A}$, $A_{mp} \in \mathbb{R}$, represents the value of property p of material m .

2.2. Observations (Data). We have observations of A_{mp} 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 $\mathbf{A}$ is not complete; some entries are missing ($|\Omega| < MP$).

2.3. Objective. The objective is to leverage the observations $\{A_{mp}\}_{(m,p) \in \Omega}$ to complete the material–property matrix, i.e., to predict the missing entries, $\{A_{mp}\}_{(m,p) \in \{1, 2, \dots, M\} \times \{1, 2, \dots, P\} \setminus \Omega}$.

2.4. Low-Rank Matrix Model. From an element perspective, the low-rank model^{22,23} assumes that each element of the matrix, A_{mp} , decomposes into

$$A_{mp} \approx \mathbf{m}_m^T \mathbf{p}_p + \mu_m \quad (1)$$

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 (i) depends on both the norm of the vectors $\mathbf{m}_m$ and $\mathbf{p}_p$ and the angle between them and (ii) is positive (negative) if the vectors point in roughly the same (opposite) directions. 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.

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 material biases; and $\mathbf{1} \in \mathbb{R}^P$ a column vector of ones (see Figure 2). The dimensionality of the latent space, $k < M, P$, imposes the constraint $\text{rank}(\mathbf{M}^T \mathbf{P}) \leq k$; hence, eq 2 is a low-rank approximation of the matrix $\mathbf{A}$.

2.5. Utility of the Low-Rank Model. The low-rank model of the material–property matrix is useful for two purposes.²²

(1) *Imputation of missing entries.* The decomposition in eq 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 eq 2.

(2) *Construction of a low-dimensional map of 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




Figure 2. Low-rank model of the material–property 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 M materials and P properties, respectively, which lie in a k -dimensional space. The vector $\boldsymbol{\mu}$ contains M material biases. Entry (m, p) of $\mathbf{A}$ is modeled as $A_{mp} \approx \mathbf{m}_m^T \mathbf{p}_p + \mu_m$.

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.²² Within this latent space, materials, represented by $\{\mathbf{m}_m\}$, that tend to have similar (dissimilar) properties congregate (separate). Using dimension reduction techniques, we can visualize the scatter of materials in the low-dimensional space to draw a map of materials. The latent representations of materials, $\{\mathbf{m}_m\}$, and the map that visualizes them are useful for: (i) grouping together/organizing materials with similar properties, (ii) lead optimization, where we search the map for materials nearby a “lead” material with a good but still suboptimal property, (iii) selecting diverse materials to efficiently explore material space in an experimental design, and (iv) training supervised machine learning models for other prediction tasks, as $\mathbf{m}_m$ is a feature vector for material m .

Similarly, the columns of a complete version of $\mathbf{A}$ can be viewed as vector representations of 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 $\{\mathbf{p}_p\}$, that tend to take on similar (dissimilar) values in NPMs congregate (separate).

The geometric interpretation of the dot product $\mathbf{m}_m^T \mathbf{p}_p$ in eq 1 makes a comparison between the map of materials and the map of properties useful. The magnitudes of and the angle between a pair of latent material and property vectors ($\mathbf{m}_m, \mathbf{p}_p$), taken together, indicate the affinity/aversion of material m with property p .

2.6. Machine-Learning the Low-Rank Model. We learn the latent representations of materials and properties and material biases from the observed (material, property) values by balancing (i) the matching of the observed entries of the matrix by the model given in eq 1 and (ii) the simplicity of latent vector representations, to avoid overfitting. Specifically, we aim to choose $\mathbf{M}$, $\mathbf{P}$, and $\boldsymbol{\mu}$ that minimize the loss function:

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

The first term is the approximation error, measured over all observations. The second term provides L2 regularization of the latent vector representations of materials and properties to prevent overfitting and improve generalization, where $\lambda > 0$ is the regularization parameter. The sums are normalized to properly weigh regularization of the latent material and property vectors.

Stochastic gradient descent or alternating minimization are commonly used to find a $(\mathbf{M}, \mathbf{P}, \boldsymbol{\mu})$ that (locally²²) minimize l .^{22,23} In alternating minimization, we alternate between optimizing $\mathbf{M}$ with $\mathbf{P}$ fixed and optimizing $\mathbf{P}$ with $\mathbf{M}$ fixed.

Two hyperparameters are involved in fitting a low-rank matrix model to the observed entries of $\mathbf{A}$: (1) $k \in \{0, 1, \dots, \min(M, P)\}$, the dimensionality of the vector space containing the latent representations of materials and properties and (2) $\lambda \in [0, \infty)$, the regularization parameter that trades off prediction accuracy on the training data and the simplicity of the latent vector representations.

3. CASE STUDY: A COF RECOMMENDATION SYSTEM

We now demonstrate a material recommendation system based on a low-rank matrix model. Here, materials are COFs, and 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.

3.1. Data Set. We construct the COF–adsorption-property matrix using an open data set (v9 on Materials Cloud⁶⁴) of simulated gas adsorption properties in $M = 572$ experimentally reported, porous COF materials.^{17,25} We selected $P = 16$ simulated adsorption properties—the uptake [units: mmol/g] and Henry coefficients [mmol/(g·bar)] of a variety of gases (CH_4 , H_2O , H_2S , Xe, Kr, CO_2 , N_2 , O_2 , and H_2) at various conditions pertinent to gas storage and separation applications (see Table 1). We $\log_{10}$ -transformed Henry coefficients because of the relatively long tail of their distributions. The resulting COF–adsorption-property matrix $\mathbf{A}_{\text{complete}} \in \mathbb{R}^{572 \times 16}$ is fully observed, allowing us to study how imputation performance of the low-rank model depends on the completeness of the matrix.

Figure 3 displays the distribution of and pairwise relationships between [standardized] adsorption properties, and Figure S1 displays a pairwise correlation matrix of properties. Some properties are strongly correlated, e.g., CH_4 uptake at (298 K, 65 bar) and O_2 uptake at (298 K, 140 bar), while others, such as H_2 uptake at (77 K, 5 bar) and H_2S Henry coefficients at 300 K, are not. The low-rank model exploits these correlations between properties to learn low-dimensional representations of materials and properties.

3.2. Methods. **3.2.1. 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 $M = 572 \times P = 16$) with a fraction $\theta \in [0, 1]$ of observed entries. We construct $\mathbf{A}^{(\theta)}$ by (uniform) randomly sampling, without replacement, $(1 - \theta)MP$ entries

Table 1. List of Equilibrium Adsorption Properties Included in Our COF Recommendation System

adsorption property	thermodynamic condition	units
O_2 uptake	298 K, 5 bar	mmol/g
O_2 uptake	298 K, 140 bar	mmol/g
CO_2 uptake	300 K, 0.001 bar	mmol/g
CO_2 uptake	300 K, 30 bar	mmol/g
N_2 uptake	300 K, 0.001 bar	mmol/g
N_2 uptake	300 K, 30 bar	mmol/g
H_2 uptake	77 K, 5 bar	mmol/g
H_2 uptake	77 K, 100 bar	mmol/g
H_2 uptake	298 K, 5 bar	mmol/g
H_2 uptake	298 K, 100 bar	mmol/g
CH_4 uptake	298 K, 65 bar	mmol/g
CH_4 uptake	298 K, 5.8 bar	mmol/g
H_2O Henry coefficient	300 K	mmol/(g·bar)
H_2S Henry coefficient	300 K	mmol/(g·bar)
Xe Henry coefficient	300 K	mmol/(g·bar)
Kr Henry coefficient	300 K	mmol/(g·bar)

to ablate (change to missing) from the MP entries of $\mathbf{A}_{\text{complete}}$. Figure S2 in the Supporting Information visualizes an instance of an incomplete COF–adsorption-property matrix $\mathbf{A}^{(\theta=0.4)}$.

3.2.2. Standardization of Adsorption Properties. In accordance with standard practice,²² we z-score standardize each adsorption property, i.e., we subtract from each column of $\mathbf{A}^{(\theta)}$ the mean of the corresponding adsorption property and then divide by its standard deviation (each computed using only the observed, training examples). We z-score standardize the adsorption properties for two reasons. First, standardization prevents properties with a larger range/variance from dominating the loss function in eq 3. Second, z-score standardization centers the latent vectors at the origin and is consistent with underlying model assumptions, e.g., a probabilistic interpretation of (closely related) principal component analysis⁸³ is that the data matrix is observed with normally distributed errors, justifying z-score standardization as opposed to Min–Max normalization.²²

3.2.3. Training, Hyperparameter Tuning, and Testing. Given an incomplete COF–adsorption-property matrix $\mathbf{A}^{(\theta)}$, we outline how we arrive at a low-rank model for deployment and test its imputation performance. Naturally, the [simulated] observed entries of $\mathbf{A}^{(\theta)}$ are used for training and hyperparameter tuning, while [simulated, but actually known] unobserved entries are used as test data to evaluate the generalization error of the final deployment low-rank model.

3.2.4. Fitting. To fit a low-rank matrix model to training observations, we use `LowRankModels.jl`²² in the Julia programming language.⁸⁴ `LowRankModels.jl` implements alternating proximal gradient descent²² to minimize the loss in eq 3 over the training observations.

3.2.5. Hyperparameter Selection. We aim to determine the optimal hyperparameters $(k_{\text{opt}}, \lambda_{\text{opt}})$ for the low-rank model. We randomly partition the [simulated] observed entries of $\mathbf{A}^{(\theta)}$ into an 80/20% training/validation set. We then fit a set of low-rank models, each with different hyperparameters, to the training data. The optimal hyperparameters follow from those of the model that gives the lowest imputation error (RMSE) over the validation set. The hyperparameter search is conducted over a regular 2D grid of (1) $k \in \{1, 2, \dots, 15\}$ and




Figure 3. Distribution of (diagonal) and pairwise relationships between (off-diagonal; each point represents a COF) the simulated gas adsorption properties of COFs (data from ref 17). Each property was z-score standardized to have zero mean and unit variance.

(2) 25 values of λ ranging from 10 to 1000 and evenly spaced on a log scale.

3.2.6. Deployment of Low-Rank Model and Testing. Finally, the *deployment* low-rank model is a new low-rank model fit to all [simulated] observed entries in $\mathbf{A}^{(\theta)}$ with hyperparameters $(k_{\text{opt}}^{(\theta)}, \lambda_{\text{opt}}^{(\theta)})$. 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 metrics are computed on the standardized and, in the case of Henry coefficients, $\log_{10}$ -transformed values.

3.3. Results for Observed Fraction $\theta = 0.4$. We now demonstrate the utility of a low-rank matrix model, using a particular instance of an incomplete COF–adsorption-

property matrix $\mathbf{A}^{(\theta=0.4)}$ (so 60% of the entries are missing), for (i) imputing missing (COF, adsorption property) values and (ii) drawing of a map of COFs and adsorption properties. Figures 4–6 all pertain to the same deployment low-rank model trained using the instance of $\mathbf{A}^{(\theta=0.4)}$ visualized in Figure S2. Our hyperparameter sweep found $k_{\text{opt}}^{(0.4)} = 9$, $\lambda_{\text{opt}}^{(0.4)} = 100$.

3.3.1. Imputing the Missing Entries. We judge the performance of the deployment low-rank matrix model for imputing the missing entries of the COF–adsorption-property matrix $\mathbf{A}^{(\theta=0.4)}$ 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 Figure 4a shows the joint distribution of the predictions for and actual values of the missing (COF, adsorption property) entries in the test data set. The density is




Figure 4. Performance of the deployment low-rank model in the task of imputing the missing (COF, adsorption property) values in $\mathbf{A}^{(\theta = 0.4)}$ comprising the test set. (a) Parity plot showing the joint distribution of predicted and actual missing (COF, adsorption property) values. The diagonal line represents perfect prediction. (b) Bar plot showing Spearman's rank correlation coefficient ρ between the predicted and actual missing (COF, adsorption property) values, grouped by adsorption property. For comparison, the stars show ρ for the benchmark, material bias model where $A_{mp} \approx \mu_m$ (interaction term excluded).

greatest along the diagonal line of equality. The mean absolute error (MAE), root mean square error (RMSE), and coefficient of determination (R^2) are 0.3, 0.52, and 0.71, respectively. The magnitude of the RMSE and MAE are directly interpretable as units of a standard deviation in the adsorption property because they are computed on z-score standardized properties. These metrics indicate that the low-rank model imputes the missing entries of the matrix with reasonable accuracy.

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 (COF, adsorption property) value by the deployment low-rank model and its actual value (from the test set), grouped by adsorption property, is shown in Figure 4b (blue bars). If the recommendation system were to randomly rank COFs, ρ would be zero, and the search for the optimal COF orchestrated by the recommendation system would be

equivalent to a random, trial-and-error search; a perfectly accurate ranking of the COFs would result in $\rho = 1$. With the exception of H₂O Henry coefficients, the recommendation system ranks COFs according to their properties reasonably well, with $\rho > 0.6$. The relatively poor ranking of COFs by H₂O Henry coefficients is explained by its very weak correlation with the other properties (see Figure S1).

Figure S3 shows the correlation between and the distributions of material biases $\{\mu_m\}$ and interaction terms $\{\mathbf{m}_m^T \mathbf{p}_p\}$, for the unobserved entries, in the deployment low-rank model. The distributions are centered around approximately zero and exhibit a similar range, indicating that both play a role in the prediction of the missing adsorption properties.

3.3.2. Comparing Imputation Performance to a Benchmark Model $A_{mp} \approx \mu_m$. As a baseline to judge the imputation performance of our recommendation system, we also train and




Figure 5. Ranked COF biases, $\{\mu_m\}$, from the deployment low-rank model. Only the COFs with the lowest and highest μ_m are shown. The bottom row visualizes the three top- and bottom-ranked COF structures.

test (on the same data) a benchmark model that excludes the interaction term $\mathbf{m}_m^T \mathbf{p}_p$ in eq 1 ($k = 0$). This material bias model $A_{mp} \approx \mu_m$ [naively] considers only 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 $k = 0$ material bias model with the $k > 0$ low-rank model, we quantify the extent to which interactions between COFs and 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 Figure 4b show Spearman's rank correlation coefficients between the predictions of the missing (COF, adsorption property) values made by the benchmark material bias model $A_{mp} \approx \mu_m$ and the actual values. Indeed, the interaction term enhances the ability of the recommendation system to rank COFs according to each adsorption property, 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 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 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).

Unsurprisingly, material biases μ_m in the $k > 0$ low-rank model in eq 1 are strongly correlated with material biases in the benchmark material bias model $A_{mp} \approx \mu_m$ (see Figure S4).

3.3.3. COF Biases. The learned material bias of COF m , μ_m in eq 1, roughly describes the typical value of the (standardized) gas adsorption properties of COF m . Figure 5 shows a partial, sorted bar plot of the material biases $\{\mu_m\}$ of COFs and

displays COF structures with the lowest and highest material biases. Py-1P-quasi-AB (COF-LZU8) has the largest (smallest) μ_m , indicating that Py-1P-quasi-AB (COF-LZU8) tends to exhibit the highest (lowest) values of (standardized) gas adsorption properties among the COFs. Given a new gas adsorption task for which we seek a COF with maximal uptake, the high material bias μ_m of Py-1P-quasi-AB makes it a good candidate for measurements, in the absence of any other information; in the analogy of movie recommendation systems, Py-1P-quasi-AB is like a movie that is widely liked. Figure S3 shows the distribution of $\{\mu_m\}$.

3.3.4. Learned Map of COFs and Gas Adsorption Properties. We now visualize and analyze the learned latent vector representations of COFs, $\{\mathbf{m}_m\}$, and of the adsorption properties, $\{\mathbf{p}_p\}$, in the deployment low-rank matrix model of $\mathbf{A}^{(\theta=0.4)}$. We resort to principal component analysis (PCA) to project $\{\mathbf{m}_m\}$ and $\{\mathbf{p}_p\}$, contained in the columns of $\mathbf{M}$ and $\mathbf{P}$, respectively, onto the 2D subspace spanned by the first two principal components of $\mathbf{M} \parallel \mathbf{P}$, where $\parallel$ denotes horizontal concatenation. The dimension reduction incurred a relative reconstruction error of 63% based on the Frobenius norm.

The resulting map of COFs in Figure 6a visualizes the organization of COF vectors $\{\mathbf{m}_m\}$ in the latent space. Because the learned latent representation of COF m , $\mathbf{m}_m$, encodes its adsorption properties, COFs with similar (dissimilar) adsorption properties are expected to congregate (separate) in the map. To illustrate, each COF vector in Figure 6a is colored by the (left) CH_4 adsorption at (298 K, 65 bar), (middle) H_2S Henry coefficient at 300 K, and (right) H_2O Henry coefficient at 300 K. Indeed, nearby COFs in the map tend to exhibit similar values of a given adsorption property: COFs with the highest CH_4 uptake at (298 K, 65 bar), H_2S Henry coefficients at 300 K, and H_2O Henry coefficients at 300 K, respectively, tend to lie on the bottom left, top, and top




Figure 6. Learned map of COFs and gas adsorption properties: the latent representations of (a) COFs, $\{\mathbf{m}_m\}$, and (b) a subset of the adsorption properties, $\{\mathbf{p}_p\}$, projected onto the 2D subspace spanned by the first two principal components (PCs) of $\mathbf{M}||\mathbf{P}$. (a) Each point represents a COF, colored by (left) CH₄ adsorption at (298 K, 65 bar), (middle) H₂S Henry coefficients at 300 K, and (right) H₂O Henry coefficients at 300 K (Outlier: Py-IP-quasi-AB). (b) Each point represents an adsorption property.

right of the latent COF space. Figure S7 shows the COF map colored by the remaining adsorption properties.

Figure S8 shows the complete map of adsorption properties—i.e., the organization of the adsorption property vectors $\{\mathbf{p}_p\}$ in the latent space. Because the interaction term in eq 1 is the dot product $\mathbf{m}_m^T \mathbf{p}_p$, the location of the latent property vector $\mathbf{p}_p$ indicates the region of latent space that tends to contain COFs with high values of adsorption property p . To illustrate, compare (i) the location of the latent property vectors $\{\mathbf{p}_p\}$ pertaining to the CH₄ adsorption at (298 K, 65 bar), H₂S Henry coefficient at 300 K, and H₂O Henry coefficient at 300 K in Figure 6b and (ii) the location of the COFs with the highest and lowest values of these properties in Figure 6a. The COFs whose latent vectors are oriented in the same (opposite) direction of the latent vector of a property tend to have large (small) values of that property.

In summary, Figure 6 illustrates that a low-rank model of an incomplete COF–adsorption-property matrix machine learns a

map of COFs, wherein COFs with similar adsorption properties congregate. These latent representations were learned from the observed values in an incomplete COF–adsorption-property matrix. The map of COFs, wherein proximity implies similarity of adsorption properties, is practically useful for: (1) lead optimization, where we search the latent space for nearest neighbors of a lead COF with good but insufficient performance, (2) selecting diverse sets of COFs in an experimental design strategy to efficiently explore COF space, and (3) building supervised machine learning models to predict other properties of COFs, where the latent representations of COFs can serve as feature vectors.

In movie recommendation systems, the latent variables of movies, contained in the vector representations of movies learned via a low-rank model from the observed (movie, user) ratings, may correspond to interpretable, intuitive features of movies such as genre or orientation toward children; however, they may also be uninterpretable.²³ It would be interesting to




Figure 7. Effect of the fraction of observed values θ on the performance of the recommendation system for imputing missing (COF, adsorption property) values. Blue circles show the mean (over 50 simulations of data collection) Spearman's rank correlation coefficient between the prediction of the missing (COF, adsorption property) value and its true value, as a function of θ , grouped by adsorption property. Yellow stars show the mean Spearman's rank correlation coefficient for the benchmark material bias model. Shaded bands signify the standard deviation.

inspect the relationship between the learned latent variables in the vectors $\{\mathbf{m}_m\}$ and the structures of COFs they represent, which were not explicitly input to the recommendation system.

3.4. Effect of Observed Fraction θ on Performance.

How complete must the COF–adsorption-property matrix be for the recommendation system to reliably rank COFs according to their adsorption properties? Because the COF–adsorption-property matrix is in reality complete, we have the luxury of studying the impact of the fraction of observed entries, θ , on the imputation performance of the recommendation system.

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 conducted a hyperparameter sweep using a training/validation split of the observed entries, trained a deployment model on all observed entries using the optimal hyperparameters, and then tested the imputation performance of the deployment model on the unobserved entries serving as test data. Figure S5 shows the distribution (among the simulations of data collection) of optimal hyperparameters ($k_{\text{opt}}^{(\theta)}$, $\lambda_{\text{opt}}^{(\theta)}$) for each θ . Figure 7 shows Spearman's rank correlation coefficient, ρ , between the prediction of the missing (COF, adsorption property) value by the deployment low-rank model and its actual value, grouped by adsorption property, as θ varies. The bands show the standard deviation over the 50 simulations of data collection. As in Figure 4b, ρ for the H₂O Henry coefficient is unsatisfactory $\forall \theta$ owing to its poor correlation with the other properties. For the majority of gas adsorption properties, the recommendation system ranks COFs according to the property reasonably well (i.e., a reasonably high $\rho > 0.7$), until θ is reduced below $\theta = 0.4$. With such a paucity of training examples for $\theta < 0.4$, the low-rank model cannot learn useful latent representations of COFs and adsorption properties, resulting in diminished imputation performance. For comparison, ρ for the baseline material bias model is also shown in Figure 4b; 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).

In conclusion, the imputation performance of the low-rank matrix model rapidly diminishes when more than 60% of the entries in the COF–adsorption-property matrix are missing. As the fraction of observed entries θ in the matrix increases, the recommendation system ranks COFs according to their adsorption properties more reliably.

3.5. Uncertainty Quantification. Suppose our material recommendation system predicts some missing (material, property) value to be optimal for an adsorption-based engineering application. To motivate an experimental measurement in the lab, we may wish to quantify the uncertainty associated with this prediction. One approach to quantify the uncertainties associated with the (material, property) values imputed by the low-rank model in eq 2 is through bootstrap resampling of the observed entries.

We demonstrate uncertainty quantification via bootstrapping for our COF recommendation system based on an instance of an incomplete COF–adsorption-property matrix, $\mathbf{A}^{(\theta=0.4)}$. First, we select 100 bootstrap samples of the observed (COF, adsorption property) values (bootstrap sample = random sample of the observations Ω , with replacement, of size $|\Omega|$). Then, we fit a low-rank model to each bootstrap sample of observations, giving us an ensemble of 100 different low-rank models. The point prediction of a missing (material, property) value follows from the mean prediction among the ensemble of low-rank models; the associated uncertainty follows from the standard deviation of the predictions. Figure S9a shows bootstrap confidence intervals overlaid on the true values of the missing (COF, adsorption property) values. The width of each bootstrap confidence interval reflects the uncertainty in the prediction. To assess whether the uncertainty estimates via bootstrap resampling indeed capture uncertainty, we show the sharpness and calibration/honesty of the uncertainty estimates⁸⁵ in Figure S9b and S9c. Indeed, the uncertainty estimates are satisfactorily honest, demonstrating that bootstrap resampling is an easy-to-implement, but computationally

expensive, means to quantify uncertainty in a recommendation system.

4. DISCUSSION

4.1. Challenges. The success of a recommendation system for NPMs is predicated on structured, open, and high-fidelity databases of NPMs and their adsorption properties.^{86–88} The NIST/ARPA-E Database of Novel and Emerging Adsorbent Materials¹⁹ (NIST-ISODB) is a collection of compiled gas adsorption measurements in NPMs from the literature, from both experimental and simulation sources, for a variety of gases at a wide range of conditions. We initially set out to develop a recommendation system using Henry coefficients extracted from experimental data in the NIST-ISODB, but we found the resultant recommendation system to perform poorly. We discuss this further and provide possible explanations in Section S7.

Our material recommendation system suffers from the *cold start*²⁴ problem. Suppose a new material is reported, but none of its properties have been observed. Then, the model in eq 1 is unable to make a prediction about any of this material's properties because we do not have data to learn its latent representation, $\mathbf{m}_{M+1}$.

Some material properties of interest may be derived from properties that appear in the material–property matrix. Owing to error propagation, the ranking by the recommendation system according to a derived property may be less accurate than according to a property that directly appears in a column of the matrix. We illustrate this in Section S9 for dilute Xe/Kr selectivity in COFs, which follows from the ratio of Xe and Kr Henry coefficients that appear in two separate columns. For this reason, it may be beneficial to include the derived property of interest as a separate column in the matrix; however, if it is derived from many other properties, this column will be sparse.

Likely, new (material, property) observations, perhaps pertaining to new materials and new properties, will be continually acquired. To update the recommendation system in light of newly acquired data, we do not need to retrain the low-rank model from scratch. The loss in eq 3 (with Ω augmented) can be minimized on-line with, e.g., alternating minimization. Given new observations pertaining to existing materials and properties, the current low-rank model serves as a warm start for minimizing the loss, and a few more iterations of alternating minimization will update the latent vectors of materials and properties and material biases. Given new observations pertaining to a new material (property), we can learn the latent vector of the new material (property) by treating the latent vectors of properties (materials) as static, much like conducting an iteration of alternating minimization.²² That said, the run time of training a new low-rank model from scratch is likely to be insignificant given the typical size of (material, property) data sets. For example, in this work, the run time to fit a low-rank model with LowRankModels.jl, on the 572×16 COF–adsorption–property matrix, is on the order of seconds.

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 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. In Section S8, we show the imputation performance of a low-rank model fit to an incomplete COF–adsorption–property matrix, where entries were observed in a bias manner; 10% of the COFs were popular for measurements and 10 % of the COFs were unpopular. Unsurprisingly, the imputation performance of the recommendation system was enhanced and diminished among the popular and unpopular COFs, respectively.

Below, we propose future research directions to address each of these challenges and improve the imputation performance of a material recommendation system.

4.2. Future Research Directions. **4.2.1. Structural and Chemical Features of Materials.** Most importantly, we propose to build a recommendation system that leverages structural and chemical features of materials, in addition to the observed adsorption properties, to impute the missing entries in the material–property matrix. A strength of the low-rank model in eq 1 is that vector representations of materials are not needed to impute the missing entries of the material–property matrix; rather, the vector representations of materials are learned from their observed properties. However, a recommendation system that uses chemical and structural features of materials to learn structure–property relationships will boost imputation performance and solve the cold start problem for new materials or properties.

A list of hand-crafted features of materials, such as the void fraction, surface area, pore size, percent carbon atoms, etc., could be input to the recommendation system for it to learn their relationships with the properties. In the analogy with movie recommendation systems, this is like including features about movies, such as the genre, directors, year of production, and actors. A simple approach to include hand-crafted features of materials is to append each feature as a (fully observed) column to the material–property matrix. However, then latent material vectors will encode the structural features of materials in addition to (more relevant to applications) properties.

4.2.2. Address Selection Bias. To characterize selection bias, it would be interesting to construct a temporal model for the process by which materials and properties are selected for experimental investigation. More, we can work to debias a recommendation system trained on experimental data with selection bias.⁸⁹

4.2.3. Active Learning and Bayesian Optimization in a Low-Rank Model Framework. In an active learning strategy, we aim to decide which missing (material, property) values to observe/measure to reduce the uncertainty in predictions by the recommendation system.^{90,91} In Bayesian optimization,⁹² we aim to find the material with the optimal property in a column of the material–property matrix using the fewest experiments—by revealing the fewest missing entries in the column. The property of interest could be a property currently covered by the recommendation system or a new, unseen property. Both active learning and Bayesian optimization will require a matrix completion method that quantifies uncertainty in its predictions.

4.2.4. Uncertainty Quantification. We demonstrated bootstrap resampling of observations for quantifying uncertainty in the predictions of the recommendation system. However, training many low-rank matrix models, one for each bootstrap sample, is computationally expensive. More advanced,

probabilistic matrix completion methods are designed to quantify uncertainties in the imputed entries of the matrix.^{93,94}

4.2.4.. Nonlinear Models. The low-rank model in eq 1 models property p of material m as a linear model in terms of the material vector input: $A_{mp} \approx f_p(\mathbf{m}_m) = \mathbf{p}_p^T \mathbf{m}_m + \mu_m$ (view $\mathbf{p}_p$ as parameters of the function $f_p(\cdot)$). In contrast, a nonlinear version of matrix factorization⁹⁵ that treats $f_p(\cdot)$ as a nonlinear function may grant us more expressiveness for capturing complex relationships between materials and adsorption properties, thereby improving imputation performance.

5. CONCLUSIONS

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 entries 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. Interestingly, the (material, property) values in a material recommendation system are mathematically analogous to (product, customer) ratings in commercial recommendation systems. A material recommendation system is useful for recommending (i, application-led material search¹) a material that optimizes a specific property or (ii, material-led application search¹⁷) an optimal application for a given material.

A material recommendation system constituted by a low-rank model of the incomplete material–property matrix, fit to the observed (material, property) values, (i) provides predictions of the missing entries in the material–property matrix and (ii) generates a map of materials, wherein materials with similar properties congregate.

We demonstrated a recommendation system that recommends COFs for different gas adsorption tasks. We constructed incomplete COF–adsorption-property matrices from 16 simulated gas adsorption properties of 572 COF structures from Ongari et al.¹⁷ Then, we trained low-rank models²² of the incomplete COF–adsorption-property matrices and (i) assessed their performance on the task of predicting the missing entries and (ii) inspected their learned map of COFs. Given that fewer than 60 % of the entries of the matrix were missing, the recommendation system was able to rank COFs according to their (missing) adsorption properties reasonably well (Spearman's rank correlation coefficient > 0.6), with the exception of H₂O Henry coefficients. Imputation performance diminishes rapidly when more than 60% of the entries in the material–property matrix are missing. Though the 60% figure does not necessarily generalize to other data sets, this demonstrates that the success of a material recommendation system is predicated on having a sufficient amount of training data. Finally, we drew the map of COFs and colored each COF vector by the adsorption properties to find that, indeed, COFs with similar (dissimilar) adsorption properties clustered together (separated) in the map.

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

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.chemmater.1c01201>.

Supporting Information contains correlations between COF adsorption properties, a visualization of the COF–adsorption-property matrix, a correlation between and distributions of material biases and interaction terms, the distribution of optimal hyperparameters among the low-rank models, the latent space of COFs colored by other adsorption properties, the latent space of properties, uncertainty quantification via bootstrap resampling, preliminary work on the NIST-ISODB, and simulations of selection bias in data collection (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Xiaoli Z. Fern – School of Electrical Engineering and Computer Science, Oregon State University, Corvallis 97331 Oregon, United States; Email: xiaoli.fern@oregonstate.edu

Cory M. Simon – School of Chemical, Biological, and Environmental Engineering, Oregon State University, Corvallis 97331 Oregon, United States; orcid.org/0000-0002-8181-9178; Email: cory.simon@oregonstate.edu

Authors

Arni Sturluson – School of Chemical, Biological, and Environmental Engineering, Oregon State University, Corvallis 97331 Oregon, United States

Ali Raza – School of Electrical Engineering and Computer Science, Oregon State University, Corvallis 97331 Oregon, United States

Grant D. McConachie – School of Chemical, Biological, and Environmental Engineering, Oregon State University, Corvallis 97331 Oregon, United States

Daniel W. Siderius – National Institute of Standards and Technology, Chemical Sciences Division, Gaithersburg 20899 Maryland, United States; orcid.org/0000-0002-6260-7727

Complete contact information is available at: <https://pubs.acs.org/doi/10.1021/acs.chemmater.1c01201>

Notes

The authors declare no competing financial interest. 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 a recommendation by NIST, nor does it imply that it is the best available for the purposes described.

■ ACKNOWLEDGMENTS

A.S., X.F., and C.M.S. acknowledge the National Science Foundation for support under grant no. 1920945. The authors thank Madeleine Udell for tips on `LowRankModels.jl` and Nikolas Christian Achatz and Tristan McInteer Gavin for querying the latest COF data from Materials Cloud.

■ REFERENCES

(1) Slater, A. G.; Cooper, A. I. Porous materials function-led design of new porous materials. *Science* **2015**, 348, aaa8075–aaa8075.

- (2) Farha, O. K.; Eryazici, I.; Jeong, N. C.; Hauser, B. G.; Wilmer, C. E.; Sarjeant, A. A.; Snurr, R.; Nguyen, S. T.; Yazaydin, A. Ö.; Hupp, J. T. Metal–Organic Framework Materials with Ultrahigh Surface Areas: Is the Sky the Limit? *J. Am. Chem. Soc.* **August 2012**, *134*, 15016–15021.
- (3) Morris, R. E.; Wheatley, P. S. Gas storage in nanoporous materials. *Angew. Chem., Int. Ed.* **2008**, *47*, 4966–4981.
- (4) Li, J.-R.; Kuppler, R. J.; Zhou, H.-C. Selective gas adsorption and separation in metal-organic frameworks. *Chem. Soc. Rev.* **2009**, *38*, 1477–1504.
- (5) Schoedel, A.; Ji, Z.; Yaghi, O. M. The role of metal-organic frameworks in a carbon-neutral energy cycle. *Nat. Energy* **2016**, *1*, No. 16034.
- (6) Herm, Z. R.; Bloch, E. D.; Long, J. R. Hydrocarbon separations in metal-organic frameworks. *Chem. Mater.* **2013**, *26*, 323–338.
- (7) Kreno, L. E.; Leong, K.; Farha, O. K.; Allendorf, M.; Van Duyn, R. P.; Hupp, J. T. Metal-organic framework materials as chemical sensors. *Chem. Rev.* **2012**, *112*, 1105–1125.
- (8) Furrusseng, D.; Aguado, S.; Pinel, C. Metal-organic frameworks: opportunities for catalysis. *Angew. Chem., Int. Ed.* **2009**, *48*, 7502–7513.
- (9) Furukawa, H.; Cordova, K. E.; O’Keeffe, M.; Yaghi, O. M. The chemistry and applications of metal-organic frameworks. *Science* **2013**, *341*, No. 1230444.
- (10) Diercks, C. S.; Yaghi, O. M. The atom, the molecule, and the covalent organic framework. *Science* **2017**, *355*, No. eaal1585.
- (11) Lu, W.; Yuan, D.; Zhao, D.; Schilling, C.; Plietzsch, O.; Muller, T.; Brase, S.; Guenther, J.; Blumel, J.; Krishna, R.; Li, Z.; Zhou, H.-C. Porous Polymer Networks: Synthesis, Porosity, and Applications in Gas Storage/Separation. *Chem. Mater.* **November 2010**, *22*, 5964–5972.
- (12) Cooper, A. I. Porous molecular solids and liquids. *ACS Cent. Sci.* **2017**, *3*, 544–553.
- (13) Gosselin, A. J.; Rowland, C. A.; Bloch, E. D. Permanently microporous metal-organic polyhedra. *Chem. Rev.* **2020**, *120*, 8987–9014.
- (14) Tranchemontagne, D. J.; Ni, Z.; O’Keeffe, M.; Yaghi, O. M. Reticular chemistry of metal-organic polyhedra. *Angew. Chem., Int. Ed.* **2008**, *47*, 5136–5147.
- (15) Mandal, S.; Natarajan, S.; Mani, P.; Pankajakshan, A. Post-Synthetic Modification of Metal–Organic Frameworks Toward Applications. *Adv. Funct. Mater.* **2020**, *31*, No. 2006291.
- (16) Boyd, P. G.; Lee, Y.; Smit, B. Computational development of the nanoporous materials genome. *Nat. Rev. Mater.* **2017**, *2*, No. 17037.
- (17) Ongari, D.; Talirz, L.; Smit, B. Too Many Materials and Too Many Applications: An Experimental Problem Waiting for a Computational Solution. *ACS Cent. Sci.* **2020**, *6*, 1890–1900.
- (18) Iacomini, Paul.; L Llewellyn, Philip. Data mining for binary separation materials in published adsorption isotherms. *Chem. Mater.* **2020**, *32*, 982–991.
- (19) NIST/ARPA-E Database of Novel and Emerging Adsorbent Materials; Siderius, D. W.; Shen, V. K.; Johnson, R. D., III; van Zee, R. D., Eds.; National Institute of Standards and Technology: Gaithersburg, MD, 2014; p 20899.
- (20) Udell, M.; Townsend, A. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science* **2019**, *1*, 144–160.
- (21) Udell, M. Big Data is Low Rank *SIAG/OPT Views and News* **2019**, *27*.
- (22) Udell, M.; Horn, C.; Zadeh, R.; Boyd, S. Generalized Low Rank Models. *Found. Trends® in Mach. Learn.* **2016**, *9*, 1–118.
- (23) Koren, Y.; Bell, R.; Volinsky, C. Matrix factorization techniques for recommender systems. *Computer* **2009**, *42*, 30–37.
- (24) Aggarwal, C. C. *Recommender Systems*. Springer, 2016.
- (25) Ongari, D.; Yakutovich, A. V.; Talirz, L.; Smit, B. Building a Consistent and Reproducible Database for Adsorption Evaluation in Covalent-Organic Frameworks. *ACS Cent. Sci.* **2019**, *5*, 1663–1675.
- (26) Moosavi, S. M.; Jablonka, K. M.; Smit, B. The Role of Machine Learning in the Understanding and Design of Materials. *J. Am. Chem. Soc.* **2020**, *142*, 20273–20287.
- (27) Bennett, S.; Tarzia, A.; Zwiijnenburg, M. A.; Jelfs, K. E. Artificial Intelligence Applied to the Prediction of Organic Materials. In *Machine Learning in Chemistry*; 2020; Vol. 17, p. 280.
- (28) Chibani, S.; Coudert, F.-X. Machine learning approaches for the prediction of materials properties. *APL Mater.* **2020**, *8*, No. 080701.
- (29) Chong, Sanggyu.; Lee, Sangwon.; Kim, Baekjun.; Kim, Jihan. Applications of machine learning in metal-organic frameworks. *Coord. Chem. Rev.* **2020**, *423*, 213487.
- (30) Shi, Z.; Yang, W.; Deng, X.; Cai, C.; Yan, Y.; Liang, H.; Liu, Z.; Qiao, Z. Machine-learning-assisted high-throughput computational screening of high performance metal-organic frameworks. *Mol. Syst. Des. Eng.* **2020**, *5*, 725–742.
- (31) Jablonka, K. M.; Ongari, D.; Mohamad Moosavi, S.; Smit, B. Big-data science in porous materials: materials genomics and machine learning. *Chem. Rev.* **2020**, *120*, 8066–8129.
- (32) Mohamad Moosavi, S.; Nandy, A.; Maik Jablonka, K.; Ongari, D.; Paul Janet, J.; G Boyd, P.; Lee, Y.; Smit, B.; Kulik, H. Understanding the diversity of the metal-organic framework ecosystem. *Nat. Commun.* **2020**, *11*, No. 4068.
- (33) Mukherjee, K.; Colón, Y. J. Machine learning and descriptor selection for the computational discovery of metal-organic frameworks. *Mol. Simul.* **April 2021**, *47*, 857–877.
- (34) Fernandez, M.; Woo, T. K.; Wilmer, C. E.; Snurr, R. Q. Large-scale quantitative structure-property relationship (QSPR) analysis of methane storage in metal-organic frameworks. *J. Phys. Chem. C* **2013**, *117*, 7681–7689.
- (35) Bucior, B. J.; Bobbitt, N. S.; Islamoglu, T.; Goswami, S.; Gopalan, A.; Yildirim, T.; K Farha, O.; Bagheri, N.; Snurr, R. Q. Energy-based descriptors to rapidly predict hydrogen storage in metal-organic frameworks. *Mol. Syst. Des. Eng.* **2019**, *4*, 162–174.
- (36) Cho, E. H.; Deng, X.; Zou, C.; Lin, L.-C. Machine learning-aided computational study of metal-organic frameworks for sour gas sweetening. *J. Phys. Chem. C* **2020**, *124*, 27580–27591.
- (37) Simon, C. M.; Mercado, R.; Schnell, S. K.; Smit, B.; Haranczyk, M. What are the best materials to separate a xenon/krypton mixture? *Chem. Mater.* **2015**, *27*, 4459–4475.
- (38) Anderson, R.; Rodgers, J.; Argueta, E.; Biong, A.; Gómez-Gualdrón, D. A. Role of Pore Chemistry and Topology in the CO₂ Capture Capabilities of MOFs: From Molecular Simulation to Machine Learning. *Chem. Mater.* **2018**, *30*, 6325–6337.
- (39) Borboudakis, G.; Stergiannakos, T.; Frysali, M.; Klontzas, E.; Tsamardinos, I.; Froudakis, G. E. Chemically intuited, large-scale screening of MOFs by machine learning techniques. *npj Comput. Mater.* **2017**, *3*, No. 40.
- (40) Pardakhti, M.; Moharreri, E.; Wanik, D.; Suib, S. L.; Srivastava, R. Machine learning using combined structural and chemical descriptors for prediction of methane adsorption performance of metal organic frameworks (MOFs). *ACS Comb. Sci.* **2017**, *19*, 640–645.
- (41) Cho, E. H.; Lin, L.-C. Nanoporous Material Recognition via 3D Convolutional Neural Networks: Prediction of Adsorption Properties. *J. Phys. Chem. Lett.* **2021**, 2279–2285.
- (42) Ma, R.; Colón, Y. J.; Luo, T. Transfer Learning Study of Gas Adsorption in Metal–Organic Frameworks. *ACS Appl. Mater. Interfaces* **July 2020**, *12*, 34041–34048.
- (43) Ahmed, A.; Siegel, D. J. Predicting hydrogen storage in MOFs via machine learning. *Patterns* **2021**, *2*, No. 100291.
- (44) Willems, T. F.; Rycroft, C. H.; Kazi, M.; Meza, J. C.; Haranczyk, M. Algorithms and tools for high-throughput geometry-based analysis of crystalline porous materials. *Microporous Mesoporous Mater.* **2012**, *149*, 134–141.
- (45) Krishnapriyan, A. S.; Haranczyk, M.; Morozov, D. Topological Descriptors Help Predict Guest Adsorption in Nanoporous Materials. *J. Phys. Chem. C* **2020**, *124*, 9360–9368.

- (46) Ohno, H.; Mukae, Y. Machine Learning Approach for Prediction and Search: Application to Methane Storage in a Metal–Organic Framework. *J. Phys. Chem. C* **2016**, *120*, 23963–23968.
- (47) Raza, A.; Waqar, F.; Sturluson, A.; Simon, C.; Fern, X. Towards Explainable Message Passing Networks for Predicting Carbon Dioxide Adsorption in Metal-organic Frameworks. Machine Learning for Molecules Workshop at NeurIPS 2020, 2020.
- (48) Coley, C. W. Defining and Exploring Chemical Spaces. *Trends Chem.* **2021**, *3*, 133–145.
- (49) Sturluson, A.; Huynh, M.; York, A. H. P.; Simon, C. M. Eigencages: Learning a latent space of porous cage molecules. *ACS Cent. Sci.* **2018**, *4*, 1663–1676.
- (50) Nicholas, T. C.; Goodwin, A. L.; Deringer, V. L. Understanding the Geometric Diversity of Inorganic and Hybrid Frameworks through Structural Coarse-Graining. *Chem. Sci.* **2020**, *11*, 12580–12587.
- (51) Lee, Y.; Barthel, S. D.; Dlotko, P.; Mohamad Moosavi, S.; Hess, K.; Smit, B. Quantifying similarity of pore-geometry in nanoporous materials. *Nat. Commun.* **2017**, *8*, No. 15396.
- (52) Nicholas, T. C.; Alexandrov, E. V.; Blatov, V. A.; Shevchenko, A. P.; Proserpio, D. M.; Goodwin, A. L.; Deringer, V. L. Visualization and Quantification of Geometric Diversity in Metal-Organic Frameworks, 2021.
- (53) Chung, Y. G.; Gómez-Gualdrón, D. A.; Li, P.; Leperi, K. T.; Deria, P.; Zhang, H.; Vermeulen, N.; Stoddart, J. F.; You, F.; Hupp, J. T.; Farha, O. K.; Snurr, R. Q. In silico discovery of metal-organic frameworks for precombustion CO₂ capture using a genetic algorithm. *Sci. Adv.* **2016**, *2*, No. e1600909.
- (54) Bao, Y.; Martin, R. L.; Simon, C. M.; Haranczyk, M.; Smit, B.; Deem, M. W. In silico discovery of high deliverable capacity metal-organic frameworks. *J. Phys. Chem. C* **2015**, *119*, 186–195.
- (55) Collins, S. P.; Daff, T. D.; Piotrkowski, S. S.; Woo, T. K. Materials design by evolutionary optimization of functional groups in metal-organic frameworks. *Sci. Adv.* **2016**, *2*, No. e1600954.
- (56) Beauregard, N.; Pardakhti, M.; Srivastava, R. In Silico Evolution of High-Performing Metal Organic Frameworks for Methane Adsorption. *J. Chem. Inf. Model.* **2021**, 3232–3239.
- (57) Zhang, X.; Zhang, K.; Lee, Y. Machine Learning Enabled Tailor-Made Design of Application-Specific Metal-Organic Frameworks. *ACS Appl. Mater. Interfaces* **2019**, *12*, 734–743.
- (58) Donval, G.; Hand, C.; Hook, J.; Dupont, E.; Sabaté Landman, M.; Freitag, M.; Lennox, M.; Düren, T. Autonomous Exploration and Identification of High Performing Adsorbents using Active Learning *ChemRxiv*, 2021.
- (59) Deshwal, A.; Simon, C.; Doppa, J. R. Bayesian Optimization of Nanoporous Materials, *ChemRxiv*, 2021.
- (60) Yao, Z.; Sánchez-Lengeling, B.; Bobbitt, N. S.; Bucior, B. J.; Hari Kumar, S. G.; Collins, S. P.; Burns, T.; Woo, T. K.; Farha, O. K.; Snurr, R. Q.; et al. Inverse design of nanoporous crystalline reticular materials with deep generative models. *Nat. Mach. Intell.* **2021**, *3*, 76–86.
- (61) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **2018**, *361*, 360–365.
- (62) Pollice, R.; dos Passos Gomes, G.; Aldeghi, M.; Hickman, R. J.; Krenn, M.; Lavigne, C.; Lindner-D'Addario, M.; Kumar Nigam, A.; Tian Ser, C.; Yao, Z.; Aspuru-Guzik, A. Data-Driven Strategies for Accelerated Materials Design. *Acc. Chem. Res.* **2021**, *54*, 849–860.
- (63) Fraux, G.; Chibani, S.; Coudert, F.-X. Modelling of framework materials at multiple scales: current practices and open questions. *Philos. Trans. R. Soc., A* **2019**, *377*, No. 20180220.
- (64) Talirz, L.; Kumbhar, S.; Passaro, E.; Yakutovich, A. V.; Granata, V.; Gargiulo, F.; Borelli, M.; Uhrin, M.; Huber, S. P.; Zoupanos, S.; Adorf, C. S.; Andersen, C. W.; Schütt, O.; Pignedoli, C. A.; Passerone, D.; Vondele, J. V.; Schulthess, T. C.; Smit, B.; Pizzi, G.; Marzari, N. Materials Cloud, a platform for open computational science. *Sci. Data* **2020**, *7*, No. 299.
- (65) Coudert, F.-X. Materials databases: the need for open, interoperable databases with standardized data and rich metadata. *Adv. Theory Simulat.* **2019**, *2*, No. 1900131.
- (66) Chung, Y. G.; Camp, J.; Haranczyk, M.; Sikora, B. J.; Bury, W.; Krungleviciute, V.; Yildirim, T.; Farha, O. K.; Sholl, D. S.; Snurr, R. Q. Computation-ready, experimental metal-organic frameworks: A tool to enable high-throughput screening of nanoporous crystals. *Chem. Mater.* **2014**, *26*, 6185–6192.
- (67) Chung, Y. G.; Haldoupis, E.; Bucior, B. J.; Haranczyk, M.; Lee, S.; Zhang, H.; Vogiatzis, K. D.; Milisavljevic, M.; Ling, S.; Camp, J. S.; Slater, B.; Siepmann, J. I.; Sholl, D. S.; Snurr, R. Q. Advances, Updates, and Analytics for the Computation-Ready, Experimental Metal–Organic Framework Database: CoRE MOF 2019. *J. Chem. Eng. Data* **2019**, *64*, 5985–5998.
- (68) Tong, M.; Lan, Y.; Yang, Q.; Zhong, C. Exploring the structure-property relationships of covalent organic frameworks for noble gas separations. *Chem. Eng. Sci.* **2017**, *168*, 456–464.
- (69) Boyd, P. G.; Chidambaram, A.; García-Díez, E.; Ireland, C. P.; Daff, T. D.; Bounds, R.; Gladysiak, A.; Schouwink, P.; Mohamad Moosavi, S.; Maroto-Valer, M. M.; Reimer, J. A.; Navarro, J. A. R.; Woo, T. K.; Garcia, S.; Stylianou, K. C.; Smit, B. Data-driven design of metal–organic frameworks for wet flue gas CO₂ capture. *Nature* **2019**, *576*, 253–256.
- (70) Moghadam, P. Z.; Li, A.; Wiggin, S. B.; Tao, A.; Maloney, A. G. P.; Wood, P. A.; Ward, S. C.; Fairen-Jimenez, D. Development of a Cambridge Structural Database subset: a collection of metal–organic frameworks for past, present, and future. *Chem. Mater.* **2017**, *29*, 2618–2625.
- (71) Mercado, R.; Fu, R.-S.; Yakutovich, A. V.; Talirz, L.; Haranczyk, M.; Smit, B. In silico design of 2D and 3D covalent organic frameworks for methane storage applications. *Chem. Mater.* **2018**, *30*, 5069–5086.
- (72) Bucior, B. J.; Rosen, A. S.; Haranczyk, M.; Yao, Z.; Ziebel, M. E.; Farha, O. K.; Hupp, J. T.; Siepmann, J. I.; Aspuru-Guzik, A.; Snurr, R. Q. Identification Schemes for Metal–Organic Frameworks To Enable Rapid Search and Cheminformatics Analysis. *Cryst. Growth Des.* **2019**, *19*, 6682–6697.
- (73) Rosen, A.; Iyer, S. M.; Ray, D. M.; Yao, Z.; Aspuru-Guzik, A.; Gagliardi, L.; Notestein, J. M.; Snurr, R. Q. Machine learning the quantum-chemical properties of metal–organic frameworks for accelerated materials discovery. *Matter* **2021**, *4*, 1578–1597.
- (74) Nazarian, D.; Camp, J. S.; Sholl, D. S. A Comprehensive Set of High-Quality Point Charges for Simulations of Metal–Organic Frameworks. *Chem. Mater.* **2018**, *30*, 785–793. DOI: 10.1021/acs.chemmater.5b03836.
- (75) Park, S.; Kim, B.; Choi, S.; Boyd, P. G.; Smit, B.; Kim, J. Text mining metal-organic framework papers. *J. Chem. Inf. Model.* **2018**, *58*, 244–251.
- (76) Tshitoyan, V.; Dagdelen, J.; Weston, L.; Dunn, A.; Rong, Z.; Kononova, O.; Persson, K. A.; Ceder, G.; Jain, A. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature* **2019**, *571*, 95–98.
- (77) Spangler, S.; Wilkins, A. D.; Bachman, B. J.; Nagarajan, M.; Dayaram, T.; Haas, P.; Regenbogen, S.; Pickering, C.; Comer, A.; Myers, J. et al. *Automated Hypothesis Generation Based on Mining Scientific Literature*, Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2014; pp 1877–1886.
- (78) Sun, Y.; DeJaco, R. F.; Li, Z.; Tang, D.; Glante, S.; Sholl, D. S.; Colina, C. M.; Snurr, R. Q.; Thommes, Matthias.; Hartmann, Martin.; Siepmann, J. I. Fingerprinting diverse nanoporous materials for optimal hydrogen storage conditions using meta-learning. *Science Advances* **2021**, *7*, eabg3983.
- (79) Yuana, Q.; Longob, M.; Thornton, A.; McKeown, N. B.; Comesana-Gándarad, B.; Jansen, J. C.; Jelfs, K. E. Imputation of Missing Gas Permeability Data for Polymer Membranes using Machine Learning. *J. Membr. Sci.* **2021**, *627*, No. 119207.

- (80) Sosnina, E. A.; Sosnin, S.; Nikitina, A. A.; Nazarov, I.; Osolodkin, D. I.; Fedorov, M. V. Recommender systems in antiviral drug discovery. *ACS Omega* **2020**, *5*, 15039–15051.
- (81) Seko, A.; Hayashi, H.; Kashima, H.; Tanaka, I. Matrix-and tensor-based recommender systems for the discovery of currently unknown inorganic compounds. *Phys. Rev. Mater.* **2018**, *2*, No. 013805.
- (82) Hayashi, H.; Hayashi, K.; Kouzai, K.; Seko, A.; Tanaka, I. Recommender System of Successful Processing Conditions for New Compounds Based on a Parallel Experimental Data Set. *Chem. Mater.* **2019**, *31*, 9984–9992.
- (83) Tipping, M.; Bishop, C. M. Probabilistic principal component analysis. *J. Royal Stat. Soc.: Series B* **1999**, *61*, 611–622.
- (84) Bezanson, J.; Edelman, A.; Karpinski, S.; Shah, V. B. Julia: A fresh approach to numerical computing. *SIAM Rev.* **2017**, *59*, 65–98.
- (85) Tran, K.; Neiswanger, W.; Yoon, J.; Zhang, Q.; Xing, E.; Ulissi, Z. W. Methods for comparing uncertainty quantifications for material property predictions. *Mach. Learning: Sci. Technol.* **2020**, *1*, No. 025006.
- (86) Sturluson, A.; Huynh, M.; Kaija, A. R.; Laird, C.; Yoon, S.; Hou, F.; Feng, Z.; Wilmer, C. E.; Colón, Y.; Chung, Y. G.; et al. The role of molecular modelling and simulation in the discovery and deployment of metal-organic frameworks for gas storage and separation. *Mol. Simul.* **2019**, *45*, 1082–1121.
- (87) Evans, J. D.; Bon, V.; Senkovska, I.; Kaskel, S. A universal standard archive file for adsorption data. *Langmuir* **2021**, *37*, 4222–4226.
- (88) Park, J.; Howe, J. D.; Sholl, D. S. How Reproducible Are Isotherm Measurements in Metal-Organic Frameworks? *Chem. Mater.* **2017**, *29*, 10487–10495.
- (89) Chen, J.; Dong, H.; Wang, X.; Feng, F.; Wang, M.; He, X. Bias and Debias in Recommender System: A Survey and Future Directions. *arXiv preprint arXiv:2010.03240*, 2020.
- (90) Rubens, N.; Kaplan, D.; Sugiyama, M. Active Learning in Recommender Systems. In *Recommender Systems Handbook*; Kantor, P. B.; Ricci, F.; Rokach, L.; Shapira, B., Eds.; Springer, 2011; pp 735–767.
- (91) Yang, C.; Akimoto, Y.; Won Kim, D.; Udell, M. In *Oboe: Collaborative filtering for Automl Model Selection*, Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019; pp 1173–1183.
- (92) Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* **2015**, *104*, 148–175.
- (93) Chen, Y.; Fan, J.; Ma, C.; Yan, Y. Inference and uncertainty quantification for noisy matrix completion. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 22931–22937.
- (94) Zhao, Y.; Udell, M. Matrix Completion with Quantified Uncertainty through Low Rank Gaussian Copula. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- (95) Lawrence, D. N.; Urtasun, R. In *Non-linear Matrix Factorization With Gaussian Processes*, Proceedings of the 26th Annual International Conference on Machine Learning, 2009; pp 601–608.