MOFs to enhance green NH₃ synthesis in plasma reactors: hierarchical computational screening enhanced by iterative machine learning

Tsung-Wei Liu, Fernando Fajardo-Rojas, Sumaya Addish, Enrique Martinez, Diego A. Gomez-Gualdron*

Department of Chemical and Biological Engineering, Colorado School of Mines, 1601 Illinois St, Golden CO 80401, USA

* Corresponding author: dgomezgualdron@mines.edu

ABSTRACT

Plasma reactors are promising to decarbonize the production of NH₃, but their NH₃ energy yields need to improve to facilitate their broad adoption. Two emerging strategies to reduce energy inefficiencies aim to protect the freshly formed NH₃ from destruction by the plasma by leveraging NH₃ adsorption properties of porous materials as either catalyst supports or as membranes. As metal-organic frameworks (MOFs) are promising porous materials for adsorption-based applications, we performed large-scale computational screening of 13,460 MOFs to study their potential for the abovementioned uses. To reduce computational cost by ~10-fold, we developed a generalizable hierarchical MOF screening strategy that starts with the selection of a 200-MOF set based on NH₃ adsorption Henry's constants, for which the relevant performance metrics are calculated via molecular simulation. This set is used to "initialize" a machine learning (ML) model that predicts the relevant metrics in the whole MOF database, in turn guiding the selection of additional promising MOFs to be evaluated via molecular simulation. The ML model is then iteratively refined leveraging the emerging molecular simulation data from the MOFs selected at each iteration from the ML predictions themselves. From evaluation of only ~10 % of the database, for each use (catalyst support or membrane), 20 extant MOFs were holistically assessed and proposed for experimental testing based on desirable adsorption properties as well as complementary properties (e.g., high thermal decomposition temperature, constituted by earth abundant metals, etc). Data-driven material design guidelines also emerged from the screening. For instance, a pore diameter of ~10 Å and a heat of adsorption of ~90 kJ/mol were found beneficial for the catalyst support use. On the other hand, for the membrane-based strategy, a pore diameter of ~2.75 Å, and a heat of adsorption of ~80 kJ/mol were found beneficial. The presence of V was found beneficial for both uses.

KEYWORDS: *metal-organic frameworks, GCMC simulations, dielectric discharge barrier, shielding protection, membrane reactor*, *gradient boosted machines, MOF histograms*

1. INTRODUCTION

NH₃ is a crucial chemical in our society primarily due to its role in fertilizer production. Over 200 million tons of NH₃ are produced yearly.[1] Moreover, NH₃ demand could increase in the future if it were to be exploited as an energy vector.[2] For instance, NH₃ could be used *i)* as a means to store H₂,[3] which could be released from NH₃ via cracking, and then used in H₂ fuel cells, whose only byproduct is H₂O, *ii)* directly in NH₃ fuel cells,[4] in which the only byproducts would be N₂ and H₂, *iii)* even directly in internal combustion engines[5], (ideally) only emitting N₂ and H₂O as byproducts. However, the current CO₂ emissions tied to NH₃ production (ca. 2% of the world's annual CO₂ emissions)[6] puts climate and food sustainability goals at odds with each other, and jeopardizes future utilization of NH₃ as a "clean" energy vector. Thus, it is apparent that achieving carbon-free NH₃ synthesis stands as one of the most important needs of our society.

Most CO₂ emissions associated with current NH₃ synthesis arise from CH₄ reforming that is done *in situ* at NH₃ production plants to obtain the H₂ feedstock for the NH₃ forming reaction (i.e. N₂ + 3H₂ = 2NH₃).[6] Thus, the key to carbon-free NH₃ synthesis is the switch to "green" H₂ as feedstock, which is likely to be produced in distributed fashion from H₂O electrolysis.[1] Accordingly, coupling NH₃ synthesis with (intermittently available) green H₂ demands distributed synthesis of NH₃ in easy turn-on/turn-off reactors,[7] thereby requiring NH₃ synthesis at low pressure.[8] To avoid high pressure requirements, one needs to synthesize NH₃ at sufficiently low temperature to overcome equilibrium constraints. Thus, achieving sufficiently fast N₂ feed decomposition (and subsequent hydrogenation) at low temperature is key, for which alternatives such as plasma-assisted catalysis are being explored.[9]

In plasma-assisted catalytic NH₃ synthesis, high-energy electrons in a non-thermal plasma collide with gas species, either dissociating them or vibrationally (or electronically) exciting them. One common way to start a plasma is by creating a discharge between two electrodes separated by insulating media, as in a dielectric barrier

discharge (DBD) reactor.[10] Although numerous species form in the plasma, vibrationally excited N_2

($N_2(v)$), and N and H radicals are currently thought as the most important species interacting with the catalyst. Either $N_2(v)$ or N radicals (depending on catalyst nitrophilicity) facilitate the formation of surface-bound N*, whereas H radicals facilitate its subsequent hydrogenation via Eley-Rideal (ER) reactions.[11] This plasma-assisted approach significantly boosts NH_3 production at low temperature (e.g., 400 K) and ambient pressure relative to thermal catalysis.[9] However, with a maximum NH_3 energy yield of 36 g_{NH3}/kWh reported to date,[10] at least a 3-fold improvement is needed to reach the 100 g_{NH3}/kWh some estimate is needed for economic feasibility.[11]

In earlier work, the reaction energy for the elementary ER step H· + $HNNH_2$ → $HNNH_3$ ($\Delta E_{rxn|r20}$) was found to empirically correlate with overall NH_3 formation rate in various plasma reactors.[7] Using the most optimistic of these correlations (i.e., TOF_{NH3} = -0.404 $\Delta E_{rxn|r20}$ - 0.55, R^2 = 0.9), NH_3 formation rate at constant plasma power could increase ~10 % for every ~0.10 eV reduction in $\Delta E_{rxn|r20}$. Given the correlation between $\Delta E_{rxn|r20}$ and N binding energy (E_N) to the catalyst ($\Delta E_{rxn|r20}$ = -0.1885 E_N - 2.64 , R^2 = 0.85), the above ~10% increase in NH_3 formation rate would correspond to an increase of 0.53 eV in E_N. Therefore, a 3-fold
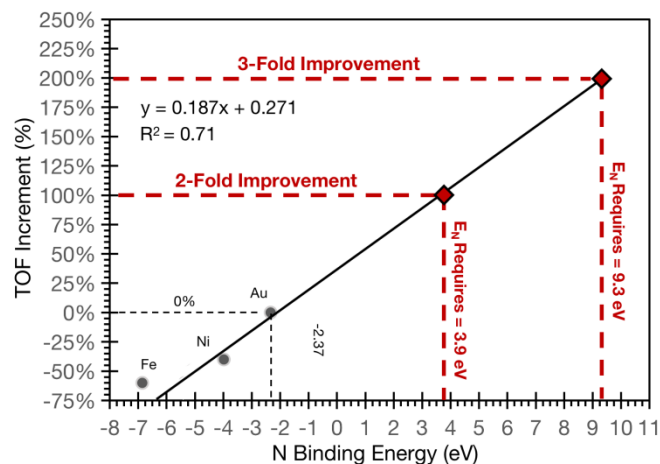

Figure 1. Percentage increment of NH_3 production turnover frequency (TOF) as a function of the N binding energy (E_N), published in a previous work.[7,11] Au is the best reported catalyst among the investigated metals, therefore used as the reference (i.e., 0% TOF increment) in this study. Red dashed lines show the estimated requirement in E_N (vertical) and its corresponding percentage TOF increment (horizontally) to achieve the required 3-fold increase in NH_3 energy yield if only the catalyst material is modified.

improvement on NH_3 energy yield solely changing the catalyst composition may require an unrealistic change in E_N beyond 9 eV over the current best catalyst for plasma-assisted NH_3 synthesis (**Fig. 1**).

On the other hand, it has been realized that additional reactor inefficiencies paradoxically arise from the same mechanisms that facilitate NH_3 formation at low temperature under plasma. Namely, the collisions with high energy electrons that can facilitate dissociation of reactants can also destroy a significant fraction of the freshly formed NH_3. For instance, simulations by vant' Veer et al.[12] suggests that between 50% and 90% of the NH_3 produced in a plasma reactor is decomposed at different stages during NH_3 synthesis aided by pulsed plasma. Thus, to complement catalysts development strategies there is now growing interest in strategies to reduce the exposure of the formed NH_3 to the plasma. Strategies to reduce NH_3 exposure to the plasma include *i)* packing the reactor bed with a highly porous catalyst, within which formed NH_3 tends to be retained through adsorption (**Fig. 2a**), and *ii)* equipping the reactor with a porous membrane that selectively adsorbs the formed NH_3 and permeates it out of the reactor (**Fig. 2b**).

Both strategies present pros and cons. For instance, the plasma reactor strategy *i* (PRS-I) can provide more active sites per reactor volume, but could unfavorably alter the discharge behavior, while also requiring cyclic operation to periodically recover the NH_3. The plasma reactor strategy *ii* (PRS-II), on the other hand, allows continuous reactor operation, but limits the number of active sites per reactor volume. In proof-of-concept for PRS-I, Gorky et al. showed a 3-fold increase in NH_3 synthesis rate using mesoporous silica SBA-15 as a catalyst, compared to fumed SiO_2.[13] Rouwenhorst et al. displayed a 2-fold enhancement in NH_3 energy yield by using zeolite-4A as an adsorbent.[14] On the other hand, as proof-of-concept for PRS-II, Gorky et al. employed CC3 as a membrane in a plasma reactor, which facilitated the removal of NH_3, achieving a 5-fold improvement in the NH_3 synthesis rate using as a baseline a reactor without catalyst.[15] Mizushima et al. utilized an aluminum tubular membrane-like catalyst to achieve a 2-fold energy efficiency improvement compared to a reactor without catalyst as well.[16]

For adsorption-driven applications—as PRS-I and PRS-II are—metal-organic frameworks (MOFs) appear as enticing material prospects due to the possibility to fine-tuning their chemistry and porous architecture to engender the required adsorption behavior.[17] MOF tunability arises because MOFs are made of interconnected metal-based nodes and organic linkers, whose combinatorics can yield an overwhelming number of porous crystal networks of distinct architecture and/or chemistry.[18–20] Indicative of the potential applicability of MOFs in plasma technologies, Shah et al. used Ni-MOF-74 to improve the NH_3 yield under different plasma powers,[21] whereas as Yang et al has shown the potential of MOFs as catalyst supports.[22] Concurrently, interest in MOF NH_3 adsorption properties has increased recently, with a growing number of experiments reporting NH_3 adsorption at 1 bar, albeit usually at near room temperature. MOFs tested for the latter include UiO-66 (16.9 $mmol_{NH3}$/g_{MOF}),[23]

Figure 2. Strategies to protect NH_3 from destruction by the plasma in dielectric barrier discharge (DBD) reactors. a) Plasma reactor strategy I (PRS-I): the DBD reactor is filled with porous adsorbent particles between the inner and outer electrodes (catalyst sites can be supported on or embedded within the particles). NH_3 is adsorbed/retained within the pores so that it is shielded from the plasma. b) Plasma reactor strategy II (PRS-II): the DBD reactor uses a metallic mesh as inner electrode over a porous membrane. NH_3 selectively adsorbs into and diffuses through the membrane leaving the reactor with little exposure to the plasma. Success of both PRS-I and PRS-II depends on material adsorption properties.

MFM-300(Sc) (19.5 $mmol_{NH3}/g_{MOF}$),[24] LiCl-MIL-53 (33.9 $mmol_{NH3}/g_{MOF}$),[25] among many others.[26,27] However, the overwhelming number of extant and hypothesized MOFs makes impractical a comprehensive evaluation of these materials for a given application. A bottleneck that in other MOF applications has been usually alleviated by large-scale computational studies,[28–30] which are however yet to be performed for PRS-I and PRS-II.

As NH_3 adsorption is not as affordable to calculate as the adsorption of small nonpolar molecules (e.g., CH_4, H_2, Xe, or Kr) that other MOF screening studies have focused on,[31,32] we decided to set up a hierarchical screening workflow. We thus first set out to calculate the NH_3 adsorption Henry's constant (K_{NH3}) as a first filter to identify 200 promising MOFs for PRS-I and PRS-II based on adsorption affinity for NH_3—from this point on referred to as the MOF_{200KH} set—which were then to have their adsorption-based performance metrics fully evaluated using molecular simulations of adsorption. Note that the use of an early Henry's constant-based filter is rather established in computational hierarchical screening of adsorbents, with earlier examples including the screening of MOFs for hexane isomer separation capabilities,[33] zeolites for separation of alcohols from aqueous solutions,[34] and MOFs for CO_2 capture from air.[35]

However, suspecting limitations in the ability of the NH_3 adsorption Henry's constant to identify the actual 200 most promising MOFs in the database, we developed and introduced a new additional screening stage to improve the quality of MOFs that ultimately emerges from screening, while still avoiding to evaluate the whole MOF database. Specifically, we iteratively trained a machine learning (ML) model (*vide infra*), and used the model iteratively to reassess the database and identify MOFs that may outperform the MOF_{200KH} set that had emerged from the early Henry's constant-based screening filter. Using the additional ML-aided stage, for each plasma reactor strategy, we ultimately obtained a new set of 200 promising MOFs (from this point on referred to as MOF_{200ML} sets). These sets presented significantly higher average performance than the original set emerging from the Henry's constant-based screening filter. Thus, the present work also presents a methodological innovation that could be applied to other MOF screenings.

Upon completion of the screening, we proceeded to propose MOF designs that could be experimentally tested to implement the two plasma reactor strategies discussed in this work. As successful material development is a multi-objective problem that goes beyond the maximization of a specific metric, additional considerations (some requiring calculation of additional properties) were made to propose these designs holistically. Furthermore, recognizing that even the synthesis and experimental testing of carefully selected candidate MOFs is not guaranteed, we also aimed to use the obtained screening data to derive robust data-driven material design rules for the described plasma reactor strategies. These rules are based on data-driven structure-property relationships that naturally emerge as a byproduct of large-scale material screening studies,[36,37] and could have impact beyond MOFs. In other words, guidelines such as optimal pore sizes, heats of adsorption, and favorable chemical functionalities/moieties should translate relatively well to other classes of porous materials that could also be used to facilitate realizing the plasma reactor design strategies presented in this work.

2. COMPUTATIONAL METHODS
2.1. MOF database. The hybrid database used in this work incorporates 13,460 MOFs, with ~75% of MOFs coming from the 2019 CoRE MOF database,[38] and ~25% of MOFs generated in previous work using our topology-based crystal constructor code ToBaCCo-3.0.[18] CoRE MOFs are extant MOFs known for their high, but non-systematic, structural and chemical diversity, and their bias towards small pores.[39] In contrast, ToBaCCo MOFs are hypothetical MOF prototypes known for their moderate, but systematic, chemical and structural diversity, and their bias toward medium to large pores.[39] The complementarity of these two MOF sources make the hybrid database highly appealing for screening. Textural properties of MOFs in the hybrid database (i.e., void fraction, surface area, and pore sizes) were estimated using the code zeo++.[40] Zeo++ utilizes two spherical probe radii for pore structure analysis: one for identifying accessible pores via a percolation algorithm and one for Monte Carlo sampling to assess the pore structure. Both radii were set at 1.3 Å to match the kinetic diameter of NH_3.

2.2. Molecular simulations. Molecular simulations were conducted using the RASPA-2.0 code.[41] The Widom insertion method, with at least 10,000 random insertion moves, was applied to determine adsorption Henry's constants at the desired temperature.[42] Grand canonical Monte Carlo (GCMC) simulations were used to estimate adsorption loadings. Each GCMC simulation was conducted with 10,000 equilibration cycles followed by 10,000 production cycles. The number of Monte Carlo moves per cycle was equal to the number of molecules in the simulation box, with a minimum of 20 moves. These moves included molecule insertion/deletion, translation, and rotation, with additional molecule swap moves for mixtures. Molecule-molecule and molecule-MOF interactions were simulated using the Lennard-Jones (LJ) and Coulomb potentials,[43] with cutoffs respectively set at 12.8 Å and 12.0 Å. Additionally, for Coulomb interactions beyond 12.0 Å, Ewald summation was used.[44,45] NH_3 and N_2 molecules were assigned LJ parameters and charges based on the TraPPE force field (which was parameterized to reproduce VLE curves),[46,47] while H_2 parameters were taken from a well-known study on dispersion attraction-driven H_2 sorption by Darkrim and Levesque.[48] For all MOF atoms, LJ parameters were sourced from the Dreiding force field, unless they were unavailable (as for some metals), in which case the universal force field parameters were borrowed.[49,50] LJ parameters for cross-interactions were obtained using Lorentz-Berthelot mixing rules.[51] Charges for MOF atoms were assigned based on the best method available for each MOF subset (higher similarity to DFT calculated charges in periodic MOF unit cells). Thus, atomic charges in ToBaCCo MOFs were previously determined using the MBBB method, which relies on density functional theory (DFT) calculations on MOF building blocks.[28] And atomic charges in CoRE MOFs were assigned using the PACMOF method, which is a machine learning model developed by Snurr and colleagues trained on DFT calculations on complete MOFs unit cells.[52] Recently, Liu and Luan highlighted the superior accuracy of PACMOF over other rapid charge assignment techniques.[53]

2.3. Performance Metrics.

Plasma reactor strategy *i* (PRS-I). Plasma discharges cannot occur within nanopores.[54] Thus, the basic material feature needed to implement the NH_3 shielding protection strategy (**Fig. 2a**) is the ability to retain as much NH_3 as possible within the material pores. We proposed to simply use the NH_3 adsorption loading N_{NH3} at 400 K and 1 bar as the performance metric for PRS-I. The temperature and pressure are selected based on typical DBD reactor operation parameters.[7,12] As adsorption selectivity is not expected to be decisive for this strategy, we perform our simulations using a pure NH_3 gas phase. We hypothesize that the benefits of a porous material with higher N_{NH3} are two-fold. Higher N_{NH3} as an indication of *i)* a higher affinity for NH_3 that should be beneficial for the retention of NH_3 within the material pores, and *ii)* a higher NH_3-holding capacity that may extend the reactor operation time elapsed between NH_3 removal periods.

Plasma reactor strategy *ii* (PRS-II). This strategy incorporates a membrane in the plasma reactor. Thus, an important material feature for PRS-II (**Fig. 2b**) is the ability to rapidly (high permeability $\mathscr{P}$) flush out the product NH_3 without flushing the reactants N_2 and H_2 (high permselectivity α). Under the solution-diffusion theory precepts, NH_3 permeability $\mathscr{P}_{NH3}$ is:

$$\mathscr{P}_{NH3} = N_{NH3} \text{ x } D_{NH3} / P_{NH3} \quad (1)$$

where D_{NH3} and P_{NH3} are the diffusion coefficient and partial pressure of NH_3, respectively. Permselectivity, on the other hand, is the product of adsorption selectivity (S_{NH3}):

$$S_{NH3} = \frac{N_{NH3}/y_{NH3}}{(N_{N2} + N_{H2})/(y_{N2} + y_{H2})} \qquad 2)$$

(where N_i and y_i represent the adsorption loading and gas phase molar fraction of species i, respectively) and diffusion selectivity ζ_{NH3}:

$$\zeta_{NH3} = D_{NH3}/D_{m,N2,H2} \qquad (3)$$

where $D_{m,N2,H2}$ is the mean diffusion coefficient of N_2 and H_2. There is usually a tradeoff between permeability and permselectivity,[55] so that a performance metric accounting for both properties can be the product $\mathscr{P}_{NH3}$ x α_{NH3}. The (less computationally expensive) adsorption-dependent contribution to that metric is "$S_{NH3} \times N_{NH3}$," which we propose as the first performance metric for screening. To this end, we calculated adsorption loadings at 400 K and 1 bar, again informed by typical DBD reactor operation parameters.[7,12] For the gas phase composition, we considered several scenarios, but focused on a case where the plasma reactor is fed H_2 and N_2 at stoichiometric ratio (3:1), and 10% conversion was achieved (similar to a single pass through the Haber Bosch process)[56]. The above resulted in a composition corresponding to 6% NH_3, 23% N_2 and 71% H_2. Note, however, that consideration of other feed ratios and conversion scenarios seems to result in similar MOF rankings (**Fig. S3**).

The (more computationally expensive) diffusion-dependent contribution to the $\mathscr{P}_{NH3}$ x α_{NH3} product is $D_{NH3}{}^2/D_{m,N2,H2}$, for which we originally anticipated performing a second screening step (adhering to the hierarchical screening philosophy) based on diffusion coefficient calculations. However, this step was modified upon learning that the diffusion-limiting pore diameter (DLPD) among numerous promising MOFs based on the adsorption-based metric ranged between 2.61 and 2.86 Å. Thus, the second screening step instead focused on identifying MOFs with DLPD values between the kinetic diameter of NH_3 (2.6 Å) and H_2/N_2 (2.9 Å/3.6 Å). This step effectively considers size selectivity (i.e., molecular sieving), which is an extreme case of diffusion selectivity.[57] Through the approach above the MOFs ultimately presented as candidates for PRS-II combine both favorable adsorption and diffusion properties.

2.4. Machine learning basics.
The search workflow was developed in Python 3.9, comprising the machine learning (ML) model itself and the iterative

learning pipeline. For the ML model, we used the tree ensemble algorithm from XGBoost 2.1.0 as it is expected to outperform other algorithms when working with small datasets. The mean absolute error of the prediction (MAE) was used as the loss function, along with elastic net regularization (ENR) to help combat overfitting. This regularization adds a penalty for adding leaves and with penalties for having large leaf weights based on the summation of the absolute (Ridge regression type) and squares (Lasso regression type) of these weights. Hyperparameter tuning was done via grid search conducted on the model trained with an initial training set of 200 MOFs (MOF_{200KH} set). The hyperparameters include a tree depth of 11, a learning rate of 0.05, and values for L1 and L2 of 6 and 2, respectively. The iterative learning pipeline was written leveraging libraries in Scipy 1.14, Pandas 2.2.2, and Sklearn 1.5.0, and primarily focuses on continuously selecting new MOFs to grow the training data based on rankings based on ML predictions, while ensuring that previously selected MOFs are excluded from the rankings. Additional details about the pipeline are discussed in Section 3.2.

3. RESULTS AND DISCUSSION

3.1. Efficacy of Henry's constant-based filter. We started by assessing the ability of the NH_3 adsorption constant, K_{NH3}, to broadly rank MOFs for PRS-I and PRS-II. To this end, 200 MOFs were randomly selected from the hybrid database (from now on referred to as the MOF_{200R} set) and had both K_{NH3} and the corresponding performance metrics calculated. Namely, N_{NH3} from pure gas adsorption for PRS-I and S_{NH3} x N_{NH3} from gas mixture (NH_3, N_2, and H_2) adsorption for PRS-II, respectively. Then, three rankings were created for MOFs in the MOF_{200R} set based on K_{NH3}, N_{NH3} (for PRS-I) and S_{NH3} x N_{NH3} (for PRS-II), respectively. Within the MOF_{200R} set, we observed the MOF rankings based on K_{NH3} and N_{NH3} (**Fig. 3a**) to correlate well with each other (Spearman ranking correlation coefficient, SRRC, equal to 0.95). A similar correlation was observed for the rankings based on K_{NH3} and S_{NH3} x N_{NH3} (**Fig. 3b**, SRRC = 0.97). These results reassured us of the ability of K_{NH3} to broadly discriminate between two MOFs with high disparity in performance, and thus be an acceptable early filter for hierarchical screening. However, despite the high SRCC values, the relatively scattered points around the parity line indicate that K_{NH3} may not be as effective in discriminating between two MOFs that exhibit closer performance. Especially in the region corresponding to the most highly ranked MOFs (e.g., **Fig. 3c**, bottom left).
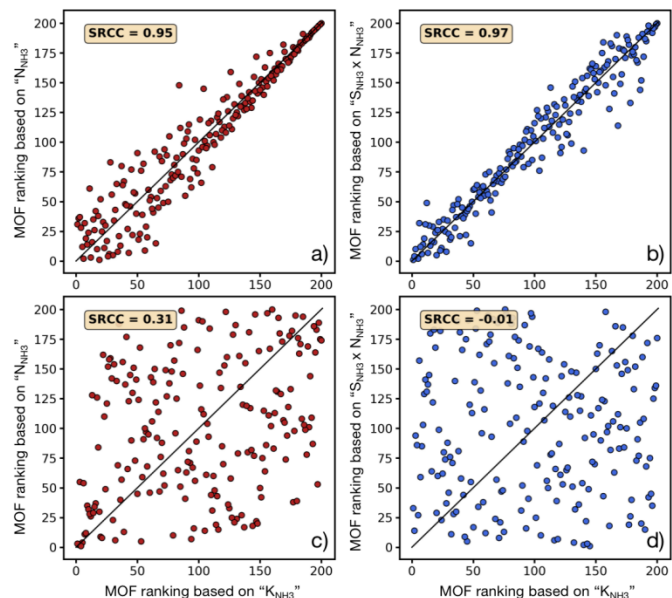

Figure 3. Parity plots comparing MOF rankings based on the Henry's constant (K_{HNH3}) and rankings based on N_{NH3} (PRS-I) (a,c) and S_{NH3} x N_{NH3} (PRS-II) (b,d). a-b) Comparison of ranking for 200 randomly chosen MOFs (MOF_{200R} set). c-d) Comparison of ranking for the 200 MOFs with the highest K_{NH3} values (MOF_{200KH} set emerging from the K_{NH3} filter in traditional hierarchical screening). Spearman ranking correlation coefficient (SRCC) is shown in each plot. K_{NH3} was found capable of broadly discriminating between really good and really bad MOFs (a-b), but not of finely discriminating between MOFs with similar performance (c,d).

Thus, we calculated K_{NH3} for the whole MOF database and selected the 200 MOFs with the highest K_{NH3} values (i.e., the previously mentioned MOF_{200KH} set), as done in standard hierarchical screening. We then calculated the corresponding performance metrics for the MOF_{200KH} sets, and again created three rankings based on K_{NH3}, N_{NH3} (for PRS-I) and S_{NH3} x N_{NH3} (for PRS-II), respectively. Within MOF_{200KH} sets, no correlation was observed between the ranking based on K_{NH3} and the rankings based on N_{NH3} and S_{NH3} x N_{NH3}, respectively (**Fig. 3c,d**), reinforcing our concern about the shortcomings of K_{NH3} to discriminate between MOFs that are not dramatically different in performance. Based on these observations, we hypothesized that the MOF_{200KH} set would capture *some* good-performing MOFs but was unlikely to closely represent the *actual* top-200 MOFs for PRS-I and PRS-II in the 13,460-MOF database. Thus, we decided to use machine learning to find a new set of 200 MOFs (i.e., the previously mentioned MOF_{200ML} set) that more closely reflects the actual top-200 MOFs for PRS-I and PRS-II applications.

3.2. Iterative machine learning-aided search workflow.

The workflow for the ML-aided search stage is summarized in **Fig. 4**. The overarching goal of our strategy was to improve the MOF_{200KH} set that emerged from traditional hierarchical screening, while still avoiding running molecular simulations on the full 13,460-MOF database. The fundamental idea was to start by initializing ML models to predict N_{NH3} (for PRS-I) and S_{NH3} x N_{NH3} (for PRS-II), respectively, leveraging the molecular

simulation data already generated for the MOF_{200R} and MOF_{200KH} sets. Specifically, the MOF_{200R} set was consistently used as a validation set and the MOF_{200KH} set was used as the *initial* training set. We expected this dataset selection strategy to make the initial training set imbalanced toward high performance MOFs, probably making the initial ML model more accurate to predict the metric in high-performing MOFs than in low-performing MOFs. However, we considered this acceptable since the goal of the model was to accurately identify the high-performing MOFs, not necessarily to correctly predict the metric in all 13,460 MOFs.

The initial ML models were then used to make a first prediction of N_{NH3} (for PRS-I) and S_{NH3} x N_{NH3} (for PRS-II), respectively, on the remaining 13,060 MOFs in the database for which molecular simulations had not been done yet. Out of this prediction, the 50 highest-performing MOFs according to the ML models for PRS-I and PRS-II, respectively, were identified and then their metrics evaluated via molecular simulation. We then *i)* identified the new top-200 MOFs by examining all MOFs evaluated with molecular simulation (this set constitutes the initial MOF_{200ML} set), and *ii)* "grew" the training dataset by adding the new molecular simulation data, and used it to train new ML models to predict N_{NH3} (for PRS-I) and S_{NH3} x N_{NH3} (for PRS-II), respectively.

The above was the end of the first iteration cycle, and subsequent cycles continued by sequentially *i)* identifying the 50 highest performing MOFs for PRS-I and PRS-II according to the predictions of the respective ML models trained in the preceding cycle, among MOFs not yet evaluated by molecular simulations, and then evaluating them with molecular simulations, *ii)* updating the MOF_{200ML} set by identifying the new top-200 MOFs among *all* MOFs evaluated with

here is a gradient boosting machine. In the flow diagram, *M* stands for the metric evaluated for each application.

molecular simulation up to that iteration, *iii)* training new ML models to predict the metric for PRS-I and PRS-II, respectively, using the larger training dataset resulting from the addition of new simulation data to the data collected in previous iterations. As we will detail later, to track the progress of the search, we determine the differences between the MOF_{200ML} set at a given iteration and *i)* the MOF_{200KH} set and *ii)* the MOF_{200ML} from the preceding iteration.

3.3. MOF representation by histograms. Adsorption loadings depend on the adsorption energies of all potential adsorption sites. Thus, towards predicting adsorption loadings via ML, some authors have opted to pre-calculate the adsorption energy of a given species along a grid of adsorption sites within a MOF unit cell, and feed that information to a ML model in the form of adsorption energy histograms.[58] However, energy histograms are a MOF representation that requires molecular modeling expertise and is largely specific to the species for which the histogram was built. Looking for a more accessible and generalizable MOF representation, and recognizing that adsorption energies in potential adsorption sites largely depend on the force field parameters (i.e., σ and ε LJ parameters and charge q) of the corresponding nearest MOF atoms (and the distance to them), we opted to develop 2D histograms based on the latter information as input for our ML model (**Fig. 5**).

The first step in generating the histograms is the generation of an evenly spaced grid (here using a 1 Å spacing between neighbor grid points). The grid symmetry matches the symmetry of the MOF unit cell (**Fig. 5a**). For instance, cubic and trigonal unit cells result in grids of cubic and trigonal symmetry, respectively. The second step is identifying the nearest atom to each grid point (considering periodic boundary conditions) (**Fig. 5b**). To make the identification efficient, instead of linear search (which involves measuring the distance of a given grid point to every MOF atom), we used space partitioning by KD trees, which are commonly used in nearest neighbor searches in many fields.[59–61] Here, KD trees partition the unit cell space into subregions by iteratively bisecting the original subregion in which an atom resides into two new subregions, each containing half the atoms residing in the original subregion. To identify the nearest neighbor atoms for a given grid point, first the subregion in which the grid point resides is determined. Then, only the distances to the atoms also residing in that subregion are calculated, and the nearest atom among them is determined.

Once the nearest atom to a given grid point is determined, the grid point is described by its distance d to the nearest atom, and the LJ parameters (σ and ε) and charge (q)
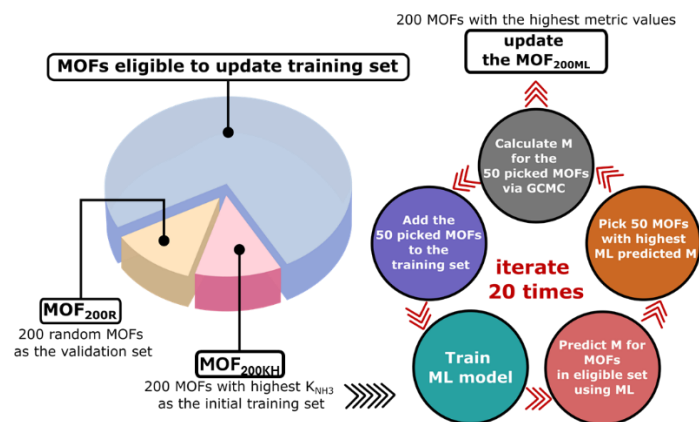

Figure 4. The workflow of the iterative machine learning (ML) pipeline. The pie chart illustrates the division of the dataset used in this work. A flow diagram is linked to this pie chart illustrating the iterative process of enhancing prediction accuracy from the ML model. This process starts with the MOF_{200KH} set and finalizes with the selection of the MOF_{200ML} set by updating the initial set with MOFs in the eligible set. The ML model integrated into the pipeline
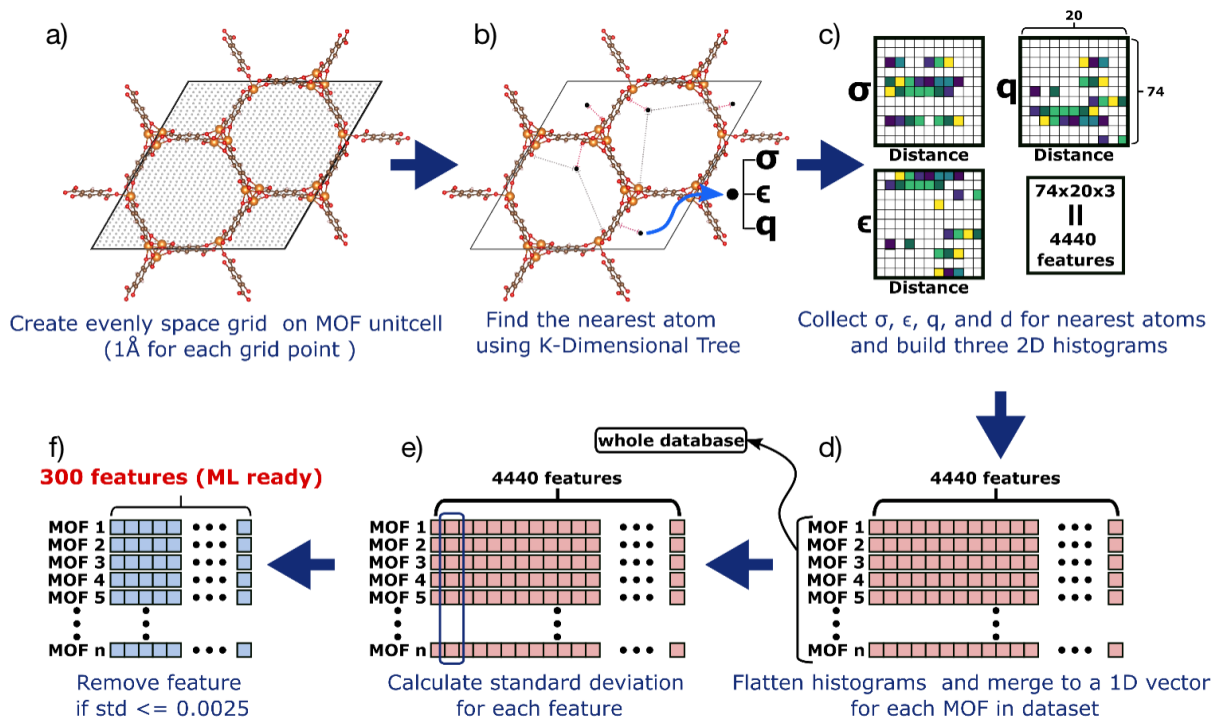
6

Figure 5. Steps to convert MOF structures into 2D histograms and then into machine-readable one-dimensional vector representations. a) An evenly spaced grid with points separated by 1 Å is created within the MOF unit cell. The symmetry of the grid is consistent with the symmetry of the unit cell. b) A K-Dimensional tree algorithm is used to find the nearest MOF atom to each grid point, and the LJ parameters (σ and ε) and charge (q) of this nearest atom are assigned to the grid point. c) The data collected of the interaction parameters (σ, ε and q) is collapsed in three 2D histograms per MOF as a function of distance (all distances higher than 20 Å are placed in the last bin). d) Each one of the three 2D histograms per MOF is "flattened" and concatenated resulting in a vector with 4440 features. e) Standard deviations for each feature across MOFs were evaluated to discard those with negligible differences along the whole dataset (using as threshold a standard deviation per feature <= 0.0025) leading to f) a MOF representation with 300 features per MOF to be used as input to the ML model.

of said nearest atom. Then the frequency (counts) with which grid points with particular (σ, d), (ε, d), and (q, d) combinations respectively occur is determined and used to build three 2D histograms. One for each kind of combination. Thus, the number assigned to each 2D histogram bin corresponds to the number of grid points falling in that particular bin divided by the total number of grid points. Each histogram features 1480 bins arranged in a 74 x 20 grid (**Fig. 5c**). The grid size is the result of dividing d into 20 bins spanning the 0-20 Å range, and σ, ε and q into 74 bins spanning the 0 to 4 Å, 0 to 350 K, and -3 to 3 e ranges, respectively (see details in Section S3).

3.4. From MOF histograms to MOF feature vectors. Visually, the resulting 2D histogram appears as a characteristic heatmap encoding the adsorption environment in a given MOF. However, as our envisioned ML model (gradient boosting machines) require a 1D vector as input, each 2D histogram was "flattened" by stacking their columns back-to-back into a 1,480 x 1 vector. Then, the three resulting vectors were stacked back-to-back into a 4,440 x 1 vector (**Fig. 5d**). Then, to reduce the 4,400 x 1 feature vector, we sought to eliminate features that were expected to correlate poorly with the metrics to be predicted, due to the similar values of these features across all MOFs. Accordingly, the standard deviations for each feature across all 13,460 MOFs were calculated (**Fig. 5e**), and features with standard deviations less or equal to 0.0025 were

eliminated. Upon completion of this exercise, a 300 x 1 feature vector emerged as a representation of the MOF for ML (**Fig. 5f**). As apparent from **Fig. S4** most of the eliminated features correspond to bins for distances beyond 6 – 10 Å (depending on σ, ε, q) and/or for q values more negative than –1.03 e and/or σ values higher than 4.4 Å.

3.5. Evolution of ML-aided MOF search. We can think of the differences between the evolving MOF$_{200ML}$ set and the MOF$_{200KH}$ set as indicative of the extent to which the ML model helps identify higher-performing MOFs over traditional hierarchical screening. These differences can be quantified based on the changes that occur relative to the MOF$_{200KH}$ to make the MOF$_{200ML}$ set in each iteration. As it turns out, by the 20th iteration of the ML-aided search about 80% of the MOFs in the MOF$_{200KH}$ were replaced (**Fig. 6a**) with higher performing MOFs, and changes between consecutive MOF$_{200ML}$ sets were minimal. The number of MOFs in the MOF$_{200KH}$ that were demoted in ranking in each iteration is lower during early and late iterations, and higher during middle iterations (**Fig. S5a,b**). The lower demotion rate during early iterations is probably due to the early ML models being less accurate (**Fig. S6**). While during late iterations the lower demotion rate is probably due to the MOF$_{200ML}$ set converging towards the actual top-200 MOFs for each scenario, respectively.

The overlap between the *actual* top-200 MOFs and the MOF_{200ML} set cannot be determined without evaluating all MOFs in the database with molecular simulation. However, there are indications that the overlap is large such as the ranking changes between consecutive MOF_{200ML} sets approaching zero by the 20^{th} iteration (**Fig. S5c-d**). Regardless of the above, by the 20^{th} iteration, the MOF_{200ML} set clearly outperforms the MOF_{200KH} set as evinced by the change in the mean value of N_{NH3} ($\log(S_{NH3} \times N_{NH3})$) from 7.36 mol/kg (4.87 log(mol/kg)) in the MOF_{200KH} set to 13.66 mol/kg (5.81 log(mol/kg)) in the MOF_{200ML} set, as shown in **Fig. 6b** (**Fig. 6c**). Thus, it is clear that the added ML-based search stage identified more promising MOFs than the standard hierarchical screening did. Remarkably, the significant improvement over the MOF_{200KH} set was accomplished with an ML model that ultimately only required molecular simulation data for ~10% of the MOF database.
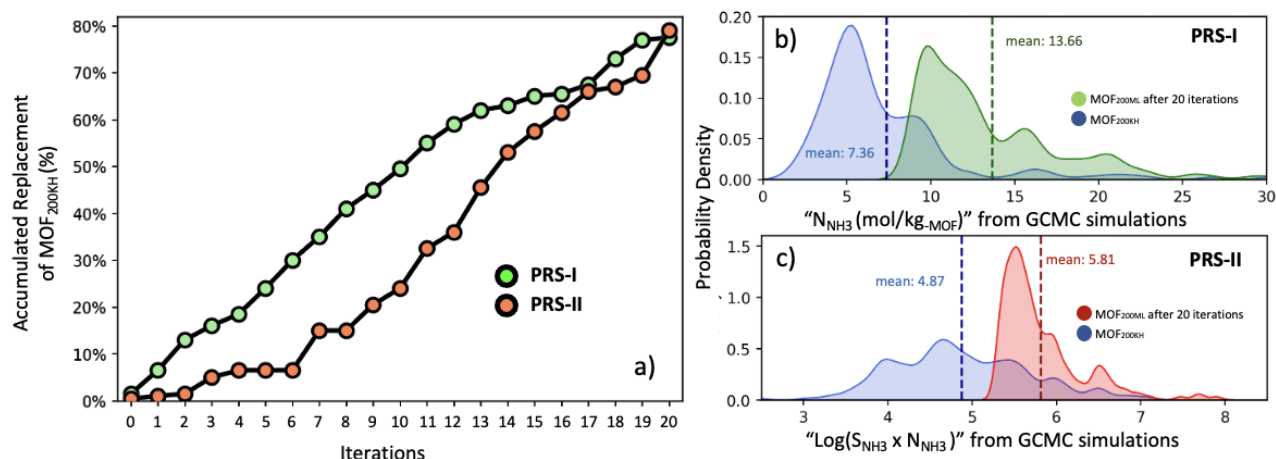
The creation of the MOF_{200ML} set is facilitated by the evolving ML model, whose evolving accuracy can be examined from the perspective of the training and validation sets. From the training set perspective, the prediction mean absolute error (MAE) for the PRS-I (log of the PRS-II) metric changed from 4.39 mol/kg (2.68 log(mol/kg)) to 0.44 mol/kg (0.1 log(mol/kg)). From the validation set perspective, the prediction error for the PRS-I (log of the PRS-II) metric changed from 2.05 mol/kg (1.73 log(mol/kg)) to 1.42 mol/kg (1.12 log(mol/kg)) (**Fig. S8-S9**). As a reference, note that the mean of the PRS-I (log of the PRS-II) metric in the final training set is 13.66 mol/kg (5.81 log(mol/kg)). Inspection of the learning curves (**Fig. S10-S11**) show that the differences in MAE between training and validation sets are not due to overfitting. Instead, these differences can be explained by the anticipated imbalance mismatch between the validation set (MOF_{200R}), which is random, and the training set, which is

Figure 6. ML-aided MOF search performance across iterations: a) Percentage of MOFs in MOF_{200KH} set (200 initial MOFs chosen based on the Henry constant, K_{NH3}) replaced by MOFs with higher metric values during each iteration. b-c) Shift in probability density of MOF performance metric values for PRS-I and for the logarithm of the metric for PRS-II, respectively, between zero (MOF_{200KH} set) and twenty iterations (final MOF_{200ML} set). All quantities calculated with N_{NH3} mol/kg.

biased toward high-performing MOFs. As noted earlier, we deemed this anticipated imbalance to be acceptable, under the condition that the models could sufficiently predict the metrics for the highest-performing MOFs, (**Fig. S7**), ensuring that their ability to correctly rank MOFs (especially high-performing ones) was sufficiently high as shown in this work (see Spearman ranking correlation coefficients, SRCC, in **Fig. S8d, S9d**).

3.6. Data-driven MOF design guidelines. A byproduct of computational material screening is the emergence of data-driven structure-property relationships (SPR) that can be exploited as guidelines to design materials for the applications of choice. These data-driven guidelines allow extracting value out of screening studies regardless of material candidate identification and experimental testing, which sometimes may not occur due to synthesis, stability, cost, or other unforeseen challenges. Herein, these relationships could be applicable beyond MOFs—as long as the adsorption mechanism is based

on physisorption—and emerged efficiently as we only ran molecular simulations for a fraction of the MOF database. This fraction includes the 200 random MOFs in the MOF_{200R} set, which allows us to glimpse the overall shape of the relationship, whereas the bias of the remaining examined MOFs toward high performance allows us to define well the relevant region where the metric of interest is optimized.

Dimensionality reduction techniques such as t-SNE clearly shows that optimal materials for PRS-I and PRS-II tend not to overlap and demand different properties of the MOFs (**Fig. S12**). For instance, PRS-I tends to demand more porosity than PRS-II. Thus, we examine structure-property relationships separately. Still, note that there are 29 MOFs appearing in the MOF_{200ML} sets for both PRS-I and PRS-II (these sets are provided as supplementary info). Specifically, for PRS-I, materials with an average pore diameter ~10 Å are better suited to maximize the performance metric (N_{NH3}) at the conditions that NH_3 is synthesized in typical DBD reactors (**Fig. 7a**). This

pore diameter is almost four times the kinetic diameter of NH_3 (2.6 Å), reflecting that an optimal material for PRS-I needs sufficient space to accommodate and retain NH_3 molecules. As pore size correlates inversely with heat of adsorption, there is a tradeoff between providing space for NH_3 and attracting NH_3 enough to retain it. This tradeoff seems to result in an optimal heat of adsorption of 90 kJ/mol (**Fig. 7b**) for PRS-I.

For PRS-II, materials with average pore diameters ~2.75 Å are better suited to maximize the performance metric (S_{NH3} x N_{NH3}) at typical plasma-assisted NH_3 synthesis conditions (**Fig. 7c**). This optimal average pore diameter is close to the kinetic diameter of NH_3 and is a manifestation of the known peak in adsorption selectivity S when the adsorbent pore diameter is closer to the diameter of the adsorbate of interest than to that of the other adsorbates.[62] The closeness of the optimal pore diameter for PRS-II to the kinetic diameter of NH_3 also suggests the stronger dominance of S in the S_{NH3} x N_{NH3} product. The optimal heat of adsorption for PRS-II is 80 kJ/mol (**Fig. 7d**). Due to the overall correlation between

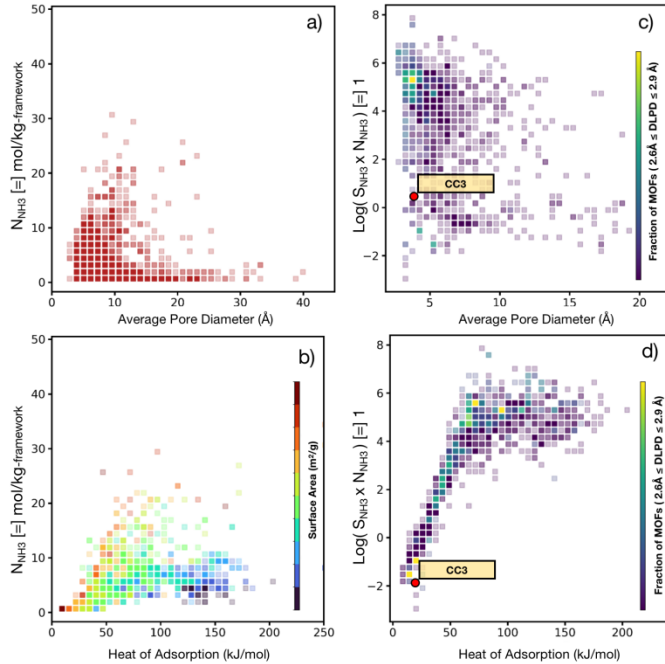leveraged in our selection of MOFs proposed for experimental testing later on.

While aiming for a material design that has specific pore dimensions is a tangible task, aiming for a certain heat of adsorption is somewhat more abstract. Partly because the heat of adsorption correlates inversely with pore diameter but also depends on chemistry. Thus, to examine the possibility of chemistry-based design rules we conducted one-sided t-tests to identify statistically significant differences between the compositions of MOFs in the final MOF_{200ML} set and the whole database (**Table 1**). For PRS-I, we find that elements that are significantly more abundant in the MOF_{200ML} set, at least with a p-value threshold of 0.05, are H, C, and V thus are deemed to potentially boost the N_{NH3} metric for a MOF of a given pore structure. By contrast, with the same p-value threshold elements such as N, O, F, Si, S, Cl, and Br are found to be

Table 1. Most statistically significant cases of higher element abundance (+ case) or lower element abundance (- case) in the final MOF_{200ML} set than in the whole database. Significance was assessed by the p-value from one-sided t-tests comparing the average element percentages in the MOF in each set.

	PRS-I			PRS-II	
Element	log p-value (+)	log p-value (-)	Element	log p-value (+)	log p-value (-)
H	-6.8	-	O	-46.9	-
V	-1.8	-	P	-6.1	-
C	-1.3	-	V	-3.8	-
Li	-0.9	-	Mo	-2.9	-
Cu	-0.9	-	Na	-2.5	-
Y	-0.9	-	Sr	-1.8	-
Pd	-0.6	-	Y	-1.6	-
In	-0.6	-	K	-1.5	-
Mo	-0.6	-	Ru	-1.2	-
Se	-0.6	-	Se	-1	-
Br	-	-69	Zn	-	-55.5
Cl	-	-22.7	Br	-	-35.7
Cr	-	-22.7	C	-	-26
Al	-	-18	Cr	-	-25.4
S	-	-11.6	H	-	-20.6
Zr	-	-7.4	Cu	-	-18.5
Mn	-	-6.6	F	-	-15.8
Cd	-	-5.9	Zr	-	-9.8
Co	-	-5.2	Fe	-	-9.4
Zn	-	-4.1	N	-	-5.6
F	-	-3.5	Ni	-	-5.6
Fe	-	-3.4	Co	-	-3.1
Si	-	-2.3	Cl	-	-3.1
Mg	-	-2.3	Mn	-	-1.7
O	-	-1.7	S	-	-1.4
Ni	-	-1.6	Mg	-	-1.3

Figure 7. The values of the metric for PRS-I (N_{NH3}) and for the logarithm of the metric for PRS-II (log(S_{NH3}xN_{NH3})) from 1,200 MOFs evaluated via molecules simulation (200 MOFs from the MOF_{200R} set, and 1,000 MOFs that at one point or another belonged to the MOF_{200ML} set) plotted against the average value of the pore size distribution (a and c) and the heat of adsorption (b and d). The color of each bin represents the average surface area in (b) and the fraction of MOFs in the database with a DLPD between 2.6 and 2.9 Å in (c and d). The percent transparency of each bin corresponds to the number of MOFs each bin contains: 25%, 50%, 75% and 100% for bins with fewer than 2 MOFs, 2-4 MOFs, 4-6 MOFs, more than 6 MOFs. Red points reference simulated values for CC3, which was experimentally tested for PRS-II in the literature.[15,63]

average pore diameter and diffusion-limiting pore window diameter (DLPD), a considerable number of MOFs with high S_{NH3} x N_{NH3} product also possess a DLPD that lies in between the kinetic diameter of NH_3 (2.6 Å) and H_2 and N_2 (2.9 Å and 3.6 Å, respectively), as shown in **Fig. 7c,d**. A fact that will be

N	-	-1.4	Cd	-	-1.3

relatively less abundant in the MOF$_{200ML}$ set for PRS-I, and thus are considered to be potentially detrimental for N_{NH3}. For PRS-II, on the other hand, we find O, P and elements like Na, K, V, Se, Sr, Y, and Mo to be more abundant in the final MOF$_{200ML}$ set than in the database and are good candidates to boost S_{NH3} x N_{NH3}. By contrast, metals such as Mg, Cr, Mn, Fe, Co, Ni, Cu, Zn, Zr, Cd, and elements such as H, C, N, F, S, Cl, and Br are less abundant in the MOF$_{200ML}$ set than in the entire dataset, suggesting these elements to typically hamper S_{NH3} x N_{NH3}. To be sure, the applicability of the above observations is contingent on the absence of chemical reactions during NH$_3$ adsorption.

We discarded the possibility that the seemingly more (less) favorable elements for PRS-I and/or PRS-II just happened to appear in MOFs featuring textural properties closer (farther) to the optimal values discussed in **Fig. 7**. To achieve this, for all MOFs featuring a given element, we created a vector containing the normalized average of largest pore diameter (LPD), diffusion limiting pore diameter (DLPD), gravimetric surface area (GSA), void fraction (V$_f$) (so that each vector component can only vary between 0 and 1), and calculated their distance to the vector containing the normalized average of the above properties in the corresponding MOF$_{200ML}$ sets. For each element X, we call this Δd$_{textural-X}$. We observed no clear direct or inverse correlation
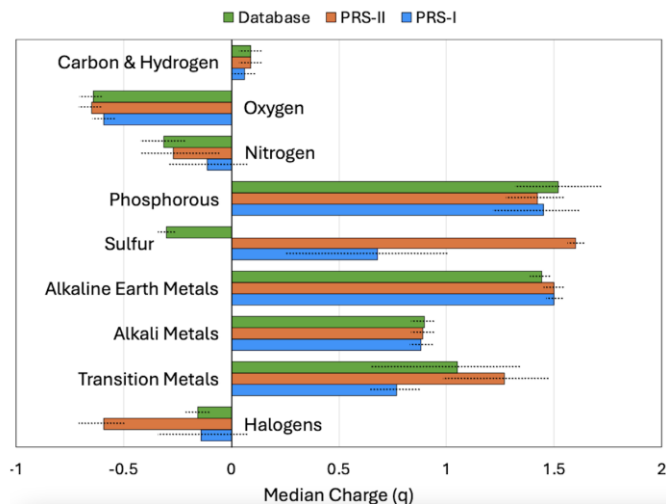

Figure 8. Comparison of the median charge (q) of different element groups in the whole database (green), the MOF$_{200ML}$ set for N_{NH3} (blue) and MOF$_{200ML}$ set for S_{NH3} x N_{NH3} (orange). Each bar indicates the corresponding median, and the dashed line indicates the corresponding median absolute error.

between Δd$_{textural-X}$ and whether an element was seemingly more or less favorable for PRS-I and/or PRS-II (**Fig. S15**). Accordingly, we proceeded to examine potential links between element properties and MOF performance, using median element charges (**Fig. 8**) as reference—given the expected role played by electrostatics during NH$_3$ adsorption—as well as visual inspection of MOFs.

With almost neutral charge, C and H having above (below) average presence in the best MOFs for PRS-I (PRS-II) is probably unrelated to their charge. But as C and H constitute the backbone of linkers, their above (below) average presence likely stems from the longer (shorter) organic linkers required to meet the larger (smaller) ideal pore diameter for PRS-I (PRS-II). Linker length requirements are also probably the primary driver of the below (above) average presence of O in the best MOFs for PRS-I (PRS-II), as O atoms are typically found capping linkers and/or embedded in the inorganic nodes. Halogens (F, Cl, Br), which typically appear as part of linker functionalization, have negative charges and below average presence in the best MOFs for either PRS-I or PRS-II. This observation suggests negative charges to be not as well suited to attract NH$_3$ as large positive charges. This notion is strengthened with S, who is typically bonded to C and has a negative charge across the database, but is bonded to O and has a positive charge in the best MOFs for PRS-I and PRS-II. Thus, its below average presence in the best MOFs reflects S rarely featuring the better suited positive charge.

Note, however, that when accompanied with a nearby positively charged element, elements with negative charges can create a pattern that reinforces the attraction to NH$_3$ (see Section 3.8). For instance, O, which was found to typically have a nearby H in a node hydroxyl or in a linker, can create such a pattern. On the other hand, elements with large positive charges and high accessibility seem more capable of attracting NH$_3$ on their own (see Section 3.8). For instance, V, which is typically a node constituent, emerges as an element with above average abundance for both PRS-I and PRS-II due to its high accessibility and high positive charge (median q ~1.5e, **Fig. S16**). Indeed, although other node-constituent metals such as Zr may have higher charges (median q ~ 2.2e, **Fig. S16**), they have below average presence in the best MOFs because they seldom present the necessary accessibility. In contrast, other node-constituent metals such as Cu, may tend to have high accessibility but may also have lower charges (median q ~1.0e, **Fig. S16**), leading to average or below average presence of this element in the best MOFs. Overall, the below, around or above average presence of an element in the best MOFs is tied to its charge and accessibility, with largely positive accessible sites being better suited for NH$_3$ adsorption. Note that the exact charges assigned to MOF atoms somewhat varies with the charge assignment method. The above information could be used as an additional early filter to pre-select a smaller set of MOFs on which to run simulations if a different charge assignment method is used.

3.7. MOFs proposed for experimental testing. We sought to suggest candidate MOFs holistically for experimental testing for PRS-I and PRS-II, going beyond simply listing the MOFs with the very highest performance metric values. We followed a rubric where, given two MOFs of similar performance to choose from for either PRS-I or PRS-II, we prefer a *i)* a previously synthesized MOF over a hypothesized one, *ii)* a MOF featuring only one metal type over one featured multiple metal types, *iii)* a MOF featuring an earth-abundant metal over

one featuring scarce or rare earth metals, *iv)* the MOF featuring the lowest NH₃ heat of adsorption, *v)* the MOF with the highest expected thermal stability. For PRS-II, we further limited our selection to MOFs with DLPD between 2.6 and 2.9 Å, as this range is conducive to very high diffusion selectivity for NH₃ over H₂ and N₂ via molecular sieving. The latter complements NH₃ adsorption selectivity and capacity (i.e., S_{NH_3} x N_{NH}) to facilitate fast and selective NH₃ permeance for PRS-II MOFs.

According to the above, while metric values for all MOFs in the final MOF_{200ML} sets are provided as Supplementary Information, here we suggest three (twenty in SI) specific MOFs for PRS-I and PRS-II, respectively, in **Fig. 9** (**Table S3-S4**). In considering the properties of the suggested MOFs, keep in mind that partly due to our additional (post-screening) selection criteria, their properties are close but not exactly equal to the optimal property values brought up in the discussion around **Fig. 7** (e.g., we traded off somewhat smaller performance metric values for more manageable heats of adsorption). These MOFs also feature chemical formulas that may not always feature elements that *tend* to be favorable or may feature elements that *tend not* to be favorable for PRS-I or PRS-II per the discussion around **Table 1**. This situation is just an indication that while here we suggest promising candidates for testing for PRS-I and PRS-II from the 13,460-MOF hybrid database, there are likely more optimal combinations of MOF chemistry and structure that were not encountered among the studied materials.

The MOF_{200ML} sets are completely dominated by structures in the CoRE MOF database, hence the MOFs suggested herein for experimental testing are all previously synthesized MOFs, whose original publication, presenting their synthesis recipe, is provided in **Table S3-S4**. The thermal stability of the MOFs was estimated using an ML model developed by Nandy *et al.*,[64] and made available in the MOFSimplify website.[65] This model consists of a trained artificial neural network that uses revised autocorrelation (RAC) descriptors and geometric features of the MOFs as input,[64,65] and was trained using reported thermogravimetric analysis (TGA) data for 3,132 MOFs (64% training, 16% validation, and 20% testing), with a mean absolute error (MAE) of 47 K. All MOFs in **Fig. 9** (and **Table S3-S4**) are predicted to have thermal decomposition temperatures (T_{Decom}) exceeding 500K, which is more than twice the MAE of the model over the relevant plasma reactor operating temperature of 400K. Thus, the anticipated ease of synthesis and high thermal stability of the proposed MOFs makes them highly attractive for experimental testing.

3.8. Adsorption mechanism in MOFs proposed for experimental testing. We conclude this work with an investigation of the underlying NH₃ adsorption mechanism in the proposed MOF candidates for PRS-I and PRS-II. To this end, we obtained the radial distribution functions (RDFs) of N and H atoms in NH₃ with respect of key atoms of each MOF, and simulated the adsorption isotherms at 400 K at the temperature and compositions relevant for each application but varying the total pressure from 0.01 to 10 bar (**Fig. S17**). We observed that a different NH₃ adsorption mechanism operates in MOFs optimal for PRS-I than those optimal for PRS-II (**Fig. S18**).
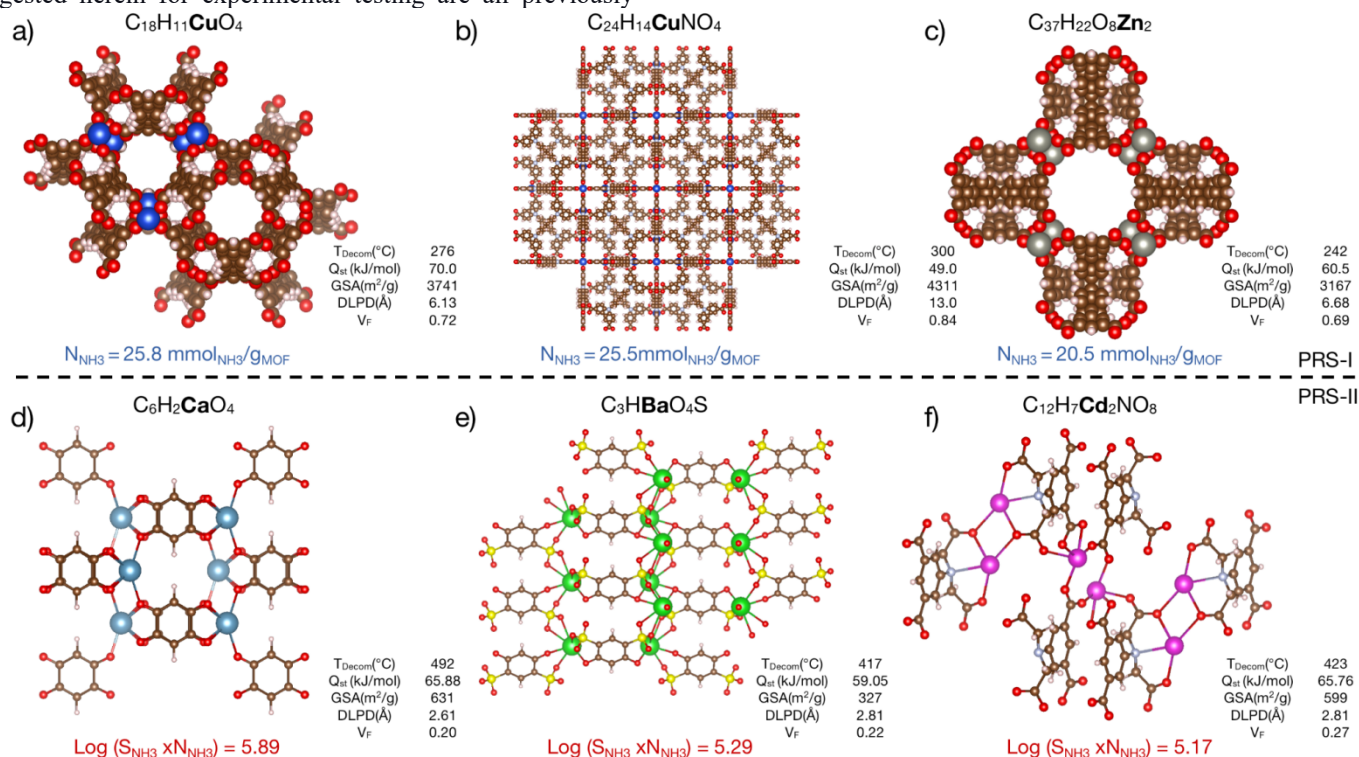

	a) C₁₈H₁₁**Cu**O₄	b) C₂₄H₁₄**Cu**NO₄	c) C₃₇H₂₂O₈**Zn₂**
T_{Decom}(°C)	276	300	242
Q_{st} (kJ/mol)	70.0	49.0	60.5
GSA(m²/g)	3741	4311	3167
DLPD(Å)	6.13	13.0	6.68
V_F	0.72	0.84	0.69

N_{NH_3} = 25.8 mmol$_{NH_3}$/g$_{MOF}$ N_{NH_3} = 25.5 mmol$_{NH_3}$/g$_{MOF}$ N_{NH_3} = 20.5 mmol$_{NH_3}$/g$_{MOF}$ PRS-I

PRS-II

	d) C₆H₂**Ca**O₄	e) C₃H**Ba**O₄S	f) C₁₂H₇**Cd₂**NO₈
T_{Decom}(°C)	492	417	423
Q_{st} (kJ/mol)	65.88	59.05	65.76
GSA(m²/g)	631	327	599
DLPD(Å)	2.61	2.81	2.81
V_F	0.20	0.22	0.27

Log (S_{NH_3} xN_{NH_3}) = 5.89   Log (S_{NH_3} xN_{NH_3}) = 5.29   Log (S_{NH_3} xN_{NH_3}) = 5.17

Figure 9. The top three holistically chosen MOFs for PRS-I (top) and PRS-II (bottom), which meet the criteria of being made with a single (relatively common) metal, having moderate heats of adsorption (Q_{st}), high thermal stability (T_{decom}). The CSD codes are CESFIQ (a), ACOCOM (b), LEQBUG (c), ALATAK (d), ZEDRAD (e), and FOKZIQ (f). T_{decom}, Q_{st}, GSA, DLPD, and V_f represent MOF decomposition temperature, heat of adsorption, gravimetric surface area, diffusion limiting pore diameter, and void fraction, respectively, for each MOF.

Generally, MOFs for PRS-II tend to operate closer to saturation than MOFs for PRS-I (**Fig. 10a,c**), which is consistent with the typically larger pores required for the latter application. For PRS-I, the H atoms in NH_3 tend to be closer to the MOF O atoms, whereas the N atom in NH_3 tends to be closer to MOF H atoms (**Fig. 10b**). This shows an orientation effect on NH_3 driven by the opposite charge of H (N) in NH_3 with respect to the charge of the MOF O (H atoms), with the H_{NH3}-O_{MOF} interaction likely corresponding to O···H···N hydrogen bonds. The further located and not particularly prominent peaks for pair interactions involving the MOF metal suggests a secondary role for the metal. The latter is consistent with the significantly lower than average transition metal charge in top MOFs for PRS-I (**Fig. 8**) and only one metal (V) making the list of elements seeming to favor performance for PRS-I (**Table 1**).

For PRS-II, the orientation effect in NH_3 is preserved with respect to the O and H atoms in the MOF. However, contrary to the PRS-I cases, it is observed that the N atom of NH_3 is closer to the metal than the H atoms in NH_3 are. Besides, the N-metal and H-metal peaks also present a distinctive sharpness not found in the PRS-I cases (**Fig. 10d**, **Fig. S18**). Accordingly, it is likely that NH_3 sorption in optimal MOFs for PRS-II is mainly driven by strong electrostatic interactions between negatively charged N in NH_3 and the MOF metal. The weak peak for H atoms in NH_3 and H atoms in the MOF, which are of equal sign charge, can be explained as a consequence of the smaller size of H atoms, the MOF operating close to saturation, and the N of NH_3 attempting to get as close as possible to the MOF metal atoms. The latter is consistent with the optimal metal charge for PRS-II being generally higher than for PRS-I, and the larger number of metals making the list of elements to seemingly favor performance for PRS-II (**Table 1**).

between N (solid lines) and H (dashed lines) atoms in NH_3 and key atom types in a representative MOF for b) PRS-I (CSD code: CESFIQ) and d) PRS-II (CSD code: ALATAK). Plots show that for PRS-I MOFs NH_3 sorption is driven by the interaction of MOF O (H) and with the H (N) of NH_3. For PRS-II, NH_3 sorption is driven by the interaction of MOF metal atoms with the N of NH_3.

CONCLUSIONS

A MOF search strategy leveraging an iteratively trained ML model was successfully developed to find MOFs with promising adsorption properties to potentially enhance plasma reactor performance for NH_3 synthesis. The MOFs found through this method had better predicted performance metrics than those of the MOFs that would have been selected using traditional hierarchical screening, which usually leverages calculation of adsorption Henry's constants. The iteratively trained ML model was found to successfully drive the MOF search despite only learning from an imbalanced training dataset that was biased toward "high performance" MOFs. The resulting MOF search required full examination (i.e., with molecular simulation) of only ~10% of a 13,460-MOF hybrid database, from which forty MOFs were proposed for experimental testing for two separate plasma reactor enhancement strategies (twenty for each case). One strategy leverages MOFs as catalyst supports while the other leverages MOFs as membranes. Selection of these MOFs went beyond sole examination of the adsorption metrics and was holistically made considering factors such as expected ease of synthesizability (e.g., having an existing recipe, being made of easily accessible metals, etc), thermal stability, among others. For instance, as part of the holistic selection for MOFs for plasma membrane reactors, the calculation of diffusion coefficients as screening step was bypassed by using an additional selection criterion: a diffusion limiting pore window size between the kinetic diameter of NH_3 and N_2/H_2. The molecular simulation generated data was shown to yield useful data-driven structure-property relationships, which could be leveraged to guide design porous materials for the plasma reactor enhancement strategies described here, even if said materials may not necessarily be MOFs. The search strategy presented here could be applied to MOF search for other applications, or even had the same principles applied to other materials and properties.

SUPPLEMENTARY INFORMATION

The supplementary information is available free of charge. Molecular simulation details; MOF rankings across different feed ratios and conversions; details of the 2D histogram algorithm; ML model performance; dimensionality reduction of the MOF textural space; structure-property and elemental composition relationships; top 20 (per application), holistically
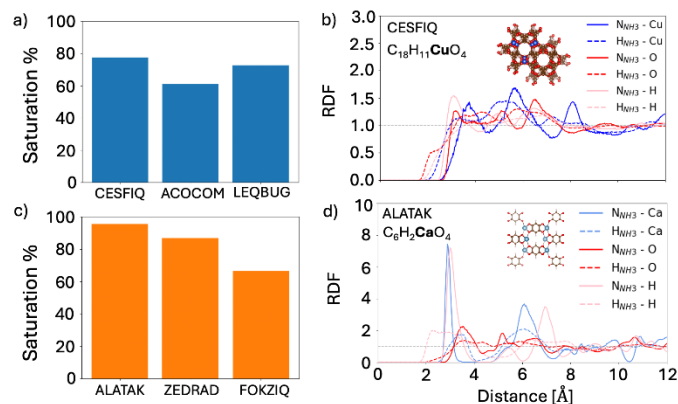

Figure 10. Saturation of the selected MOFs under operation conditions (400 K and 1 bar) for a) PRS-I, and c) PRS-II. Radial distribution function (RDF)

chosen, promising MOFs for experimental testing. Spreadsheets with data for MOF$_{200KH}$ and MOF$_{200ML}$ sets.

REFERENCES

(1) Ghavam, S.; Vahdati, M.; Wilson, I. A. G.; Styring, P. Sustainable ammonia production processes. *Front. Energy Res.* **2021**, *9*. 580808.

(2) Salmon, N.; Bañares-Alcántara, R. Green ammonia as a spatial energy vector: a review. *Sustainable Energy Fuels* **2021**, *5*, 2814–2839.

(3) Giddey, S.; Badwal, S. P. S.; Munnings, C.; Dolan, M. Ammonia as a renewable energy transportation media. *ACS Sustain. Chem. Eng.* **2017**, *5*, 10231–10239.

(4) Afif, A.; Radenahmad, N.; Cheok, Q.; Shams, S.; Kim, J. H.; Azad, A. K. Ammonia-fed fuel cells: a comprehensive review. *Renew. Sustain. Energy Rev* **2016**, *60*, 822–835.

(5) Chiong, M.-C.; Chong, C. T.; Ng, J.-H.; Mashruk, S.; Chong, W. W. F.; Samiran, N. A.; Mong, G. R.; Valera-Medina, A. Advancements of combustion technologies in the ammonia-fuelled engines. *Energy Conversion and Management* **2021**, *244*, 114460.

(6) Li, C.; Wang, T.; Gong, J. Alternative strategies toward sustainable ammonia synthesis. *Trans. Tianjin Univ.* **2020**, *26*, 67–91.

(7) Liu, T.-W.; Gorky, F.; Carreon, M. L.; Gómez-Gualdrón, D. A. Energetics of Reaction Pathways Enabled by N and H Radicals during Catalytic, Plasma-Assisted NH$_3$ Synthesis. *ACS Sustain. Chem. Eng.* **2022**. 10, 6, 2034–2051

(8) Aziz, M.; Wijayanta, A. T.; Nandiyanto, A. B. D. Ammonia as effective hydrogen storage: A review on production, storage and utilization. *Energies* **2020**, *13*, 3062.

(9) Hong, J.; Prawer, S.; Murphy, A. B. Plasma catalysis as an alternative route for ammonia production: status, mechanisms, and prospects for progress. *ACS Sustain. Chem. Eng.* **2018**, *6*, 15–31.

(10) Miao, Y.; Yokochi, A.; Jovanovic, G.; Zhang, S.; von Jouanne, A. Application-oriented non-thermal plasma in chemical reaction engineering: A review. *Green Energy and Resources* **2023**, *1*, 100004.

(11) Liu, T.-W.; Gorky, F.; Carreon, M. L.; Gomez-Gualdron, D. A. Plasma Radicals as Kinetics-Controlling Species during Plasma-Assisted Catalytic NH$_3$ Formation: Support from Microkinetic Modeling. *ACS Sustain. Chem. Eng.* **2023**. 11, 47, 16749–16763

(12) van 't Veer, K.; Engelmann, Y.; Reniers, F.; Bogaerts, A. Plasma-Catalytic Ammonia Synthesis in a DBD Plasma: Role of Microdischarges and Their Afterglows. *J. Phys. Chem. C* **2020**, *124*, 22871–22883.

(13) Gorky, F.; Guthrie, S. R.; Smoljan, C. S.; Crawford, J. M.; Carreon, M. A.; Carreon, M. L. Erratum: Plasma ammonia synthesis over mesoporous silica SBA-15 (2021 J. Phys. D: Appl. Phys. 54 264003). *J. Phys. D Appl. Phys.* **2021**, *54*, 449601.

(14) Rouwenhorst, K. H. R.; Mani, S.; Lefferts, L. Improving the Energy Yield of Plasma-Based Ammonia Synthesis with In Situ Adsorption. *ACS Sustain. Chem. Eng.* **2022**, *10*, 1994–2000.

(15) Gorky, F.; Nguyen, H. M.; Lucero, J. M.; Guthrie, S.; Crawford, J. M.; Carreon, M. A.; Carreon, M. L. CC3 porous organic cage crystals and membranes for the non-thermal plasma catalytic ammonia synthesis. *Chemical Engineering Journal Advances* **2022**, *11*, 100340.

(16) Mizushima, T.; Matsumoto, K.; Sugoh, J.; Ohkita, H.; Kakuta, N. Tubular membrane-like catalyst for reactor with dielectric-barrier-discharge plasma and its performance in ammonia synthesis. *Applied Catalysis A: General* **2004**, *265*, 53–59.

(17) Kirchon, A.; Feng, L.; Drake, H. F.; Joseph, E. A.; Zhou, H.-C. From fundamentals to applications: a toolbox for robust and multifunctional MOF materials. *Chem. Soc. Rev.* **2018**, *47*, 8611–8638.

(18) Anderson, R.; Gómez-Gualdrón, D. A. Increasing topological diversity during computational "synthesis" of porous crystals: how and why. *CrystEngComm* **2019**, *21*, 1653–1665.

(19) 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.

(20) O'Keeffe, M.; Peskov, M. A.; Ramsden, S. J.; Yaghi, O. M. The Reticular Chemistry Structure Resource (RCSR) database of, and symbols for, crystal nets. *Acc. Chem. Res.* **2008**, *41*, 1782–1789.

(21) Shah, J.; Wu, T.; Lucero, J.; Carreon, M. A.; Carreon, M. L. Nonthermal Plasma Synthesis of Ammonia over Ni-MOF-74. *ACS Sustain. Chem. Eng.* **2019**, *7*, 377–383.

(22) Yang, L.; Carreon, M. A. Green deoxygenation of fatty acids to transport fuels over metal-organic frameworks as catalysts and catalytic supports. In *Metal-Organic Frameworks (MOFs) for Environmental Applications*; Elsevier, **2019**; pp. 285–318.

(23) Ma, Y.; Lu, W.; Han, X.; Chen, Y.; da Silva, I.; Lee, D.; Sheveleva, A. M.; Wang, Z.; Li, J.; Li, W.; et al. Direct Observation of Ammonia Storage in UiO-66 Incorporating Cu(II) Binding Sites. *J. Am. Chem. Soc.* **2022**, *144*, 8624–8632.

(24) Han, X.; Lu, W.; Chen, Y.; da Silva, I.; Li, J.; Lin, L.; Li, W.; Sheveleva, A. M.; Godfrey, H. G. W.; Lu, Z.; et al. High Ammonia Adsorption in MFM-300 Materials: Dynamics and Charge Transfer in Host-Guest Binding. *J. Am. Chem. Soc.* **2021**, *143*, 3153–3161.

(25) Guo, L.; Hurd, J.; He, M.; Lu, W.; Li, J.; Crawshaw, D.; Fan, M.; Sapchenko, S.; Chen, Y.; Zeng, X.; et al. Efficient capture and storage of ammonia in robust aluminium-based metal-organic frameworks. *Commun. Chem.* **2023**, *6*, 55.

(26) Kim, D. W.; Kang, D. W.; Kang, M.; Choi, D. S.; Yun, H.; Kim, S. Y.; Lee, S. M.; Lee, J.-H.; Hong, C. S. High Gravimetric and Volumetric Ammonia Capacities in Robust Metal-Organic Frameworks Prepared via Double Postsynthetic Modification. *J. Am. Chem. Soc.* **2022**, *144*, 9672–9683.

(27) Rieth, A. J.; Tulchinsky, Y.; Dincă, M. High and Reversible Ammonia Uptake in Mesoporous Azolate Metal-Organic Frameworks with Open Mn, Co, and Ni Sites. *J. Am. Chem. Soc.* **2016**, *138*, 9401–9404.

(28) Argueta, E.; Shaji, J.; Gopalan, A.; Liao, P.; Snurr, R. Q.; Gómez-Gualdrón, D. A. Molecular Building Block-Based Electronic Charges for High-Throughput Screening of Metal-Organic Frameworks for Adsorption Applications. *J. Chem. Theory Comput.* **2018**, *14*, 365–376.

(29) Islamov, M.; Babaei, H.; Anderson, R.; Sezginel, K. B.; Long, J. R.; McGaughey, A. J. H.; Gomez-Gualdron, D. A.; Wilmer, C. E. High-throughput screening of hypothetical metal-organic frameworks for thermal conductivity. *npj Comput. Mater.* **2023**, *9*, 11.

(30) Colón, Y. J.; Snurr, R. Q. High-throughput computational screening of metal-organic frameworks. *Chem. Soc. Rev.* **2014**, *43*, 5735–5749.

(31) Taw, E.; Neaton, J. B. Accelerated discovery of ch$_4$ uptake capacity metal–organic frameworks using bayesian optimization. *Adv. Theory Simul.* **2022**, *5*. 2100515

(32) Gantzler, N.; Deshwal, A.; Doppa, J. R.; Simon, C. Multi-fidelity Bayesian Optimization of Covalent Organic Frameworks for Xenon/Krypton Separations. *Digital Discovery*, **2023**,*2*, 1937-1956.

(33) Chung, Y. G.; Bai, P.; Haranczyk, M.; Leperi, K. T.; Li, P.; Zhang, H.; Wang, T. C.; Duerinck, T.; You, F.; Hupp, J. T.; et al. Computational screening of nanoporous materials for hexane and heptane isomer separation. *Chem. Mater.* **2017**, *29*, 6315–6328.

(34) Bai, P.; Jeon, M. Y.; Ren, L.; Knight, C.; Deem, M. W.; Tsapatsis, M.; Siepmann, J. I. Discovery of optimal zeolites for challenging separations and chemical transformations using predictive materials modeling. *Nat. Commun.* **2015**, *6*, 5912.

(35) Anderson, R.; Rodgers, J.; Argueta, E.; Biong, A.; Gómez-Gualdrón, D. A. Role of pore chemistry and topology in the co$_2$ capture capabilities of mofs: from molecular simulation to machine learning. *Chem. Mater.* **2018**, *30*, 6325–6337.

(36) Schweitzer, B.; Archuleta, C.; Seong, B.; Anderson, R.; Gómez-Gualdrón, D. A. Electronic effects due to organic linker-metal surface interactions: implications on screening of MOF-encapsulated catalysts. *Phys. Chem. Chem. Phys.* **2020**, *22*, 2475–2487.

(37) Graff, D. E.; Shakhnovich, E. I.; Coley, C. W. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.* **2021**, *12*, 7866–7881.

(38) Chung, Y. G.; Haldoupis, E.; Bucior, B. J.; Haranczyk, M.; Lee, S.; Zhang, H.; Vogiatzis, K. D.; Milisavljevic, M.; Ling, S.; Camp, J. S.; et al. Advances, Updates, and Analytics for the Computation-Ready, Experimental Metal–Organic Framework Database: CoRE MOF 2019. *J. Chem. Eng. Data* **2019**, *64*, 5985–5998.

(39) Moosavi, S. M.; Nandy, A.; Jablonka, K. M.; Ongari, D.; Janet, J. P.; Boyd, P. G.; Lee, Y.; Smit, B.; Kulik, H. J. Understanding the diversity of the metal-organic framework ecosystem. *Nat. Commun.* **2020**, *11*, 4068.

(40) 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. *Micropor. Mesopor. Mat.* **2012**, *149*, 134–141.

(41) Dubbeldam, D.; Calero, S.; Ellis, D. E.; Snurr, R. Q. RASPA: molecular simulation software for adsorption and diffusion in flexible nanoporous materials. *Mol. Simul.* **2016**, *42*, 81–101.

(42) Widom, B. Some topics in the theory of fluids. *J. Chem. Phys.* **1963**, *39*, 2808–2812.

(43) Voter, A. F. Introduction to the kinetic monte carlo method. In *Radiation effects in solids*; Sickafus, K. E.; Kotomin, E. A.; Uberuaga, B. P., Eds.; Springer Netherlands: Dordrecht, **2007**; pp. 1–23.

(44) Karasawa, N.; Goddard, W. A. Acceleration of convergence for lattice sums. *J. Phys. Chem.* **1989**, *93*, 7320–7327.

(45) Chen, Z.-M.; Çağin, T.; Goddard, W. A. Fast Ewald sums for general van der Waals potentials. *J. Comput. Chem.* **1997**, *18*, 1365–1370.

(46) Chen, B.; Siepmann, J. I. Transferable Potentials for Phase Equilibria. 3. Explicit-Hydrogen Description of Normal Alkanes. *J. Phys. Chem. B* **1999**, *103*, 5370–5379.

(47) Zhang, L.; Siepmann, J. I. Development of the trappe force field for ammonia. *Collect. Czechoslov. Chem. Commun.* **2010**, *75*, 577–591.

(48) Darkrim, F.; Levesque, D. Monte Carlo simulations of hydrogen adsorption in single-walled carbon nanotubes. *J. Chem. Phys.* **1998**, *109*, 4981–4984.

(49) Mayo, S. L.; Olafson, B. D.; Goddard, W. A. DREIDING: a generic force field for molecular simulations. *J. Phys. Chem.* **1990**, *94*, 8897–8909.

(50) Rappe, A. K.; Casewit, C. J.; Colwell, K. S.; Goddard, W. A.; Skiff, W. M. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *J. Am. Chem. Soc.* **1992**, *114*, 10024–10035.

(51) Lorentz, H. A. Ueber die Anwendung des Satzes vom Virial in der kinetischen Theorie der Gase. *Ann. Phys.* **1881**, *248*, 127–136.

(52) Kancharlapalli, S.; Gopalan, A.; Haranczyk, M.; Snurr, R. Q. Fast and Accurate Machine Learning Strategy for Calculating Partial Atomic Charges in Metal-Organic Frameworks. *J. Chem. Theory Comput.* **2021**, *17*, 3052–3064.

(53) Liu, S.; Luan, B. Benchmarking various types of partial atomic charges for classical all-atom simulations of metal-organic frameworks. *Nanoscale* **2022**, *14*, 9466–9473.

(54) Zhang, Y.-R.; Van Laer, K.; Neyts, E. C.; Bogaerts, A. Can plasma be formed in catalyst pores? A modeling investigation. *Appl. Catal. B* **2016**, *185*, 56–67.

(55) Castro-Dominguez, B.; Leelachaikul, P.; Messaoud, S. B.; Takagaki, A.; Sugawara, T.; Kikuchi, R.; Oyama, S. T. The optimal point within the Robeson upper boundary. *Chemical Engineering Research and Design* **2015**, *97*, 109–119.

(56) Rouwenhorst, K. H. R.; Van der Ham, A. G. J.; Lefferts, L. Beyond Haber-Bosch: The renaissance of the Claude process. *Int. J. Hydrogen Energy* **2021**, *46*, 21566–21579.

(57) Anderson, R.; Schweitzer, B.; Wu, T.; Carreon, M. A.; Gómez-Gualdrón, D. A. Molecular simulation insights on xe/kr separation in a set of nanoporous crystalline membranes. *ACS Appl. Mater. Interfaces* **2018**, *10*, 582–592.

(58) Shi, K.; Li, Z.; Anstine, D. M.; Tang, D.; Colina, C. M.; Sholl, D. S.; Siepmann, J. I.; Snurr, R. Q. Two-Dimensional Energy Histograms as Features for Machine Learning to Predict Adsorption in Diverse Nanoporous Materials. *J. Chem. Theory Comput.* **2023**, *19*, 4568–4583.

(59) Ram, P.; Sinha, K. Revisiting kd-tree for Nearest Neighbor Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining - KDD '19*; ACM Press: New York, New York, USA, **2019**; pp. 1378–1388.

(60) Chen, Q. P.; Xue, B.; Siepmann, J. I. Using the k-d Tree Data Structure to Accelerate Monte Carlo Simulations. *J. Chem. Theory Comput.* **2017**, *13*, 1556–1565.

(61) Anzola, J.; Pascual, J.; Tarazona, G.; González Crespo, R. A Clustering WSN Routing Protocol Based on k-d Tree Algorithm. *Sensors* **2018**, *18*. 580808

(62) Wei, Q.; Lucero, J. M.; Crawford, J. M.; Way, J. D.; Wolden, C. A.; Carreon, M. A. Ammonia separation from N2 and H2 over LTA zeolitic imidazolate framework membranes. *J. Memb. Sci.* **2021**, *623*, 119078.

(63) Lucero, J.; Osuna, C.; Crawford, J. M.; Carreon, M. A. Microwave-assisted synthesis of porous organic cages CC3 and CC2. *CrystEngComm* **2019**, *21*, 4534–4537.

(64) Nandy, A.; Duan, C.; Kulik, H. J. Using Machine Learning and Data Mining to Leverage Community Knowledge for the Engineering of Stable Metal-Organic Frameworks. *J. Am. Chem. Soc.* **2021**, *143*, 17535–17547.

(65) Nandy, A.; Terrones, G.; Arunachalam, N.; Duan, C.; Kastner, D. W.; Kulik, H. J. MOFSimplify, machine learning models with extracted stability data of three thousand metal-organic frameworks. *Sci. Data* **2022**, *9*, 74.

TOC