

Physics-Informed Analytical Models for Interpretable and Deployable Hydrogen Storage Prediction in MOFs

Alauddin Ahmed *

Mechanical Engineering Department, [University of Michigan](#), Ann Arbor, Michigan 48109, USA



(Received 10 April 2025; revised 15 September 2025; accepted 3 November 2025; published 3 December 2025)

Storing hydrogen safely and compactly is a central challenge for clean transportation, grid support, and industrial decarbonization. Metal-organic frameworks (MOFs) are compelling candidates because of their ultrahigh internal surface areas and large pore volumes, yet the design space is effectively infinite, making exhaustive synthesis and simulation impractical. We introduce a physics-informed analytical modeling approach that uses deterministic sparsity-promoting symbolic regression with explicit physical constraints to produce a family of closed-form equations for fast, interpretable, and deployable prediction of usable hydrogen storage in MOFs. On a curated dataset of 88 400 MOFs, the method discovers compact, physically consistent structure-property relationships using as few as one property (void fraction for volumetric capacity) or two properties (void fraction and pore volume for gravimetric capacity) selected from seven descriptors: single-crystal density, gravimetric surface area, volumetric surface area, void fraction, pore volume, largest cavity diameter, and pore-limiting diameter. Equations using one to three properties achieve design-grade accuracy ($R^2 > 0.99$ for gravimetric capacity and $R^2 > 0.95$ for volumetric capacity), while additional descriptors provide marginal gains for high-capacity edge cases, matching black-box machine learning performance and providing mechanistic insights into why some MOFs excel and others do not. Because evaluation is simulation-free and effectively costless, the models integrate directly into high-throughput screening, active learning workflows, and generative design to prioritize syntheses and avoid dead ends, establishing a generalizable path to interpretable and deployable surrogates for adsorption-based energy materials.

DOI: [10.1103/glnw-drm4](https://doi.org/10.1103/glnw-drm4)

I. INTRODUCTION

Hydrogen (H_2) represents our most promising pathway to deep decarbonization across transportation, industry, and grid storage, but its widespread adoption hinges critically on solving the storage challenge [1–4]. Its high energy density (120 MJ/kg), zero-carbon combustion, and abundance make it a promising alternative to fossil fuels [5–8]. Unlike batteries, which suffer from energy density limitations and long charging times, H_2 can provide on-demand high-power output, making it ideal for vehicular applications, aerospace, and grid energy storage [9]. However, one of the greatest challenges preventing the large-scale adoption of H_2 technologies is efficient storage [6,10].

For H_2 to be viable in transportation and stationary applications, storage systems must achieve high gravimetric and volumetric capacities while remaining lightweight, cost-effective, and safe [10–13]. Conventional methods, such as high-pressure compression (typically 350–700 bar) and cryogenic liquefaction (requiring temperatures as low as 20 K), introduce severe energy penalties and infrastructure challenges [11,12]. Solid-state materials [6], particularly nanoporous adsorbents [14], have emerged as an alternative storage medium, with metal-organic frameworks (MOFs) being the most promising [5,7,8,15–19]. Their exceptional surface area, tunable porosity, and reversible adsorption properties make them attractive for H_2 storage at moderate pressures and temperatures [20–23].

Commercial viability requires meeting performance benchmarks established by the U.S. Department of Energy (DOE), which has set system-level targets of 5.5 wt % gravimetric capacity and 40 g- H_2 L⁻¹ volumetric capacity for automotive hydrogen storage by 2025, with ultimate goals of 6.5 wt % and 50 g- H_2 L⁻¹ [24]. These targets represent complete storage systems—including tanks, valves, and all balance-of-plant components—not just the storage

*Contact author: alauddin@umich.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

material itself. Achieving these benchmarks would enable hydrogen vehicles to match gasoline vehicle performance in terms of driving range and refueling convenience, essential for widespread market adoption.

Since the first demonstration of H₂ adsorption in MOFs in 2003 [16], research has focused on identifying materials that can achieve high usable capacities—the fraction of stored H₂ that can be practically extracted under real-world conditions [5,6]. However, despite extensive computational and experimental screening efforts, no single MOF has yet met these performance benchmarks [5,8,15,25–27].

The complexity of this challenge is amplified by the vast chemical design space of MOFs [28–32]. With hundreds of thousands of possible organic linkers and metal clusters, over 100 trillion potential structures exist [33], making exhaustive experimental or computational screening impossible. To navigate this enormity, researchers have increasingly relied on high-throughput screening (HTS) [5, 8,15,34–39] and machine learning (ML) [7,37,38,40–55] approaches. While these tools have accelerated material discovery, they present important limitations that remain unresolved.

First, HTS campaigns still rely heavily on grand canonical Monte Carlo (GCMC) [56,57] simulations to calculate adsorption isotherms—an approach that, despite its accuracy, remains computationally demanding and limits scalability. While graphics processing unit (GPU)-accelerated GCMC software, such as GPU-optimized Monte Carlo (GOMC) and gRASPA [58], has reduced computational costs compared to CPU-based tools like RASPA [59], significant preprocessing challenges persist. These simulations require clean computation-ready atomic structure files (e.g., CIF) and accurate charge calculations, which often rely on density functional theory (DFT) or semiempirical methods. Furthermore, the assignment of interatomic potential parameters using force fields (e.g., UFF [60], DREIDING [61], TraPPE [62]) and the associated charge schemes introduces variability, as simulation results depend heavily on the specific combination of force field and charge assignment [6,8,33]. These steps become increasingly time-consuming when screening vast MOF libraries—for example, the over 100 trillion possible MOFs explored by Lee *et al.* [33]. Automating GCMC simulations at such scales, or integrating them into genetic algorithm-based and inverse design workflows, remains a significant challenge, even with GPU acceleration.

Second, ML models [63], although fast, often act as black boxes, obscuring the underlying physics of adsorption and offering limited interpretability. As a result, they provide little mechanistic insight to guide material design and optimization. Third, traditional empirical correlations, such as the Chahine rule [64–66] and its extensions [25, 32,39,67,68], are overly simplistic and derived from small datasets under fixed (P , T) conditions. These models tend to overestimate hydrogen storage, particularly in MOFs

with very high surface areas (e.g., $>4500 \text{ m}^2 \text{ g}^{-1}$), due to their reliance on single-descriptor fits and the assumption of linearity [8,32]. Moreover, they generally fail to capture the nuanced interplay of pore volume, density, and void fraction that governs volumetric capacities, and they break down when predicting performance across structurally diverse MOFs. Such inconsistencies highlight the inadequacy of surface area–based heuristics and reinforce the need for more rigorous physics-aware modeling frameworks.

Compounding these limitations, recent breakthroughs in AI-driven MOF discovery have introduced new opportunities but also revealed additional challenges. Generative systems like MOFGen [69], which use large language models (LLMs) and diffusion models for automated MOF design and have successfully synthesized five “AI-dreamt” MOFs [69], represent remarkable advances in computational MOF discovery. However, these powerful generative AI systems currently lack efficient guidance mechanisms to steer their design process toward structures with optimal hydrogen storage performance. Without property-informed constraints, they generate structures that must still undergo expensive *post-hoc* screening to identify high-performing candidates. This limitation underscores the critical need for compact deployable models that can be directly embedded into generative algorithms, inverse design workflows, and AI-driven discovery platforms to enable targeted synthesis of high-performing MOFs rather than generating random structures and subsequently screening them.

To overcome these challenges, we present a symbolic regression (SR) approach that derives closed-form analytical equations mapping MOF structural descriptors to usable hydrogen storage capacities. Recent advances have produced diverse SR methodologies, including physics-informed approaches [the sure independence screening and sparsifying operator (SISSO) [70–72]], genetic programming approaches [Eureqa [73], GP-GOMEA [74], Python symbolic regression (PySR) [75,76], gplearn [77]], sparse regression techniques (FFX [75], SINDy [78], LASSO [63,79]), neural-symbolic methods (AI Feynman [80], GraphDSR [81], PINN-symbolic hybrids [82]), LLM-integrated frameworks (LaSR [83], FormulaGPT [84]), and reinforcement learning approaches (deep symbolic regression [85]). Despite this methodological diversity, systematic comparative evaluation for hydrogen storage applications remains limited, leaving critical questions about method selection, fundamental prediction limits, and algorithmic trade-offs unresolved.

To address these knowledge gaps, we conduct comprehensive benchmarking across four symbolic regression frameworks: SISSO, AI Feynman, PySR, and gplearn. Our comparative analysis reveals fundamental trade-offs between statistical accuracy and scientific interpretability. Among these methods, we select SISSO [70–72] as

our primary approach for achieving an optimal balance of competitive performance with superior mechanistic interpretability, which is essential for material discovery applications. Key advantages include (1) competitive predictive accuracy while maintaining theoretical grounding; (2) systematic exploration of descriptor spaces exceeding 10^9 candidates through recursive algebraic operations and correlation-based sure independence screening while maintaining computational tractability; (3) the generation of compact equations that are interpretable and deployable without computational infrastructure; (4) deterministic open-source implementation, ensuring reproducible results; and (5) proven success in related material problems, including perovskite stability prediction [86] and methane storage in COFs [72]. SISO's grounding in compressed sensing theory [87,88] and ℓ_0 -regularized optimization provides theoretical guarantees, while its physics-constrained feature construction—building descriptors from physically meaningful structural properties through algebraic operations—increases the likelihood that derived equations reflect realistic structure-property relationships rather than statistical overfitting.

Here, we demonstrate that complex hydrogen storage behavior in MOFs can be captured by remarkably simple algebraic equations. Using SISO on 88 400 MOFs, compact 2–3 term expressions that provide physical insights guide rational design. Comprehensive benchmarking reveals that, while alternative methods achieve higher statistical performance—particularly PySR—all symbolic regression approaches exhibit universal challenges in volumetric prediction, indicating fundamental material physics constraints rather than algorithmic limitations. Unlike conventional ML [7,37,38,40–55], these theoretical equations are interpretable, theoretically grounded, and deployable in practice [72]. These equations can be implemented in spreadsheets, embedded in material databases, integrated into high-throughput screening pipelines, or even calculated manually—requiring no specialized software, computational clusters, or molecular simulation expertise. The cross-method validation confirms that physics-informed symbolic regression successfully captures underlying structure-property relationships while providing crucial advantages in interpretability and practical deployment over evolutionary optimization approaches. This democratizes hydrogen storage prediction and accelerates material discovery across research environments with varying computational resources.

II. METHODOLOGY

A. Dataset and data preparation

To develop predictive models for usable H_2 storage capacities in MOFs, we used a curated dataset of 88 400 MOFs derived from our previously published

work [5,7,72]. The complete dataset is publicly accessible through Zenodo [89]. This dataset consists of seven crystallographic properties that serve as primary structural descriptors: single-crystal density (ρ_c , g cm^{-3}), gravimetric surface area (S_g , $\text{m}^2 \text{g}^{-1}$), volumetric surface area (S_v , $\text{m}^2 \text{cm}^{-3}$), pore volume (V_p , $\text{cm}^3 \text{g}^{-1}$), void fraction (F_v , unitless), largest cavity diameter (D_i , Å), and pore limiting diameter (D_f , Å). Each MOF entry also includes its usable gravimetric capacity (C_g , wt%) and usable volumetric capacity (C_v , $\text{g-H}_2 \text{L}^{-1}$) under a pressure swing between 100 and 5 bar at 77 K. To maintain consistency in model training and validation, all MOFs with missing or zero values in any of these properties were excluded.

To extend the generalizability of the models, we leveraged a comprehensive superdatabase of approximately 900 000 MOFs, compiled from 19 open-source repositories and developed as part of our previously published work [7], also accessible through HyMARC [90]. After filtering out MOFs with missing or zero values in any of the seven crystallographic descriptors, approximately 600 000 MOFs remained. This filtered set enabled robust and generalizable model development across a broad design space, capturing meaningful structure-property trends relevant to hydrogen adsorption.

B. SISO implementation

For predictive modeling, the SISO [70–72] algorithm was employed to establish explicit algebraic relationships between MOF crystallographic descriptors and H_2 storage capacities. Given the computational demands associated with the SISO framework [70–72], the dataset was systematically analyzed to identify the smallest subset that preserved the statistical properties of the full dataset [72]. To validate the representativeness of the 400-MOF subset beyond individual descriptor distributions, we employed a two-tiered approach. First, we conducted two-sample t tests comparing the seven base descriptors between the 400-MOF subset and the full dataset. All p values exceeded 0.05, confirming no statistically significant differences in means and indicating that the subset preserved the central tendencies of the base descriptors. The continuous mapping theorem [91,92] ensures that these representative base descriptors yield statistically robust composite descriptors under SISO's continuous operations (e.g., exponential, logarithmic, arithmetic). Second, we evaluated higher-order moments—skewness and kurtosis—across 75 incrementally increasing subset sizes. In the final sampling quintile, the coefficients of variation for skewness and kurtosis were approximately 0.003 and 0.01, respectively, indicating sufficient convergence and stability. The 400-MOF subset's skewness and kurtosis values approach those of the full dataset, capturing asymmetric distributional behavior and rare-event configurations.

The SISSO algorithm was configured to apply four binary operations—addition, subtraction, multiplication, and division—along with seven unary operations, namely, exponentiation, logarithm, square root, cube root, inverse, square, and cube, to generate composite descriptors from the seven crystallographic descriptors. The algorithm iteratively searched an extensive space of composite descriptors, ranking them based on their correlation with usable H_2 capacity under the 100-to-5 bar pressure swing condition. The highest-ranked descriptors were then used to construct closed-form equations, with models consisting of one, two, or three terms, selected through linear and multilinear regression to balance predictive accuracy and interpretability.

To optimize model complexity and relevance, three feature selection strategies were explored. The first used the full seven-feature dataset, incorporating all available crystallographic descriptors. The second utilized a reduced five-feature subset, excluding pore geometric descriptors such as the largest cavity diameter and pore-limiting diameter. The third involved a three-feature subset that excluded void fraction and volumetric surface area. This systematic reduction of descriptors ensured that models retained high predictive performance while minimizing computational complexity, making the approach adaptable for both experimental validation and large-scale screening efforts. The workflow, dataset partitioning strategy, and high-throughput SISSO methodology employed in this study were adapted from prior work [72].

C. Comparative symbolic regression framework

To validate our physics-informed approach and establish performance benchmarks, we implemented three alternative symbolic regression methods: AI Feynman [80], PySR [75,76], and gplearn [77].

1. Dataset partitioning for benchmarking methods

The complete 88 400-MOF dataset was systematically partitioned into 18 data chunks, with 17 chunks containing 5000 samples each and the final chunk containing 3400 samples. Samples were randomly assigned to chunks using stratified sampling to ensure representative coverage of the crystallographic descriptor space within each partition. This partitioning strategy was necessitated by the computational demands of evolutionary and neural-symbolic methods, which require substantially more data than SISSO's deterministic approach to achieve statistical reliability in equation discovery.

2. AI Feynman implementation

The neural-symbolic hybrid approach [80] combines neural network pattern recognition with symbolic mathematics for equation discovery. Each of the 18 data chunks

was processed independently for both targets using 80-20 train-test splits. The algorithm configuration included a brute force search duration of 1 h, neural network training for 25 000 epochs, a polynomial fitting degree of 4, and a 14-operation symbolic operator set. The algorithm first applies neural networks to identify patterns, then constructs symbolic expressions using a brute force search over mathematical operations.

3. PySR implementation

Evolutionary symbolic regression [75,76] was implemented using the Julia backend. Each data chunk was processed with evolutionary parameters: 2500 iterations, 75 populations with 150 individuals per population, maximum equation complexity of 150 terms, maximum tree depth of 30, and 2000 evaluation cycles per iteration. Binary operators $\{+, -, \times, \div, \wedge\}$ and unary operators $\{\sin, \cos, \exp, \log, \sqrt{\cdot}, |\cdot|\}$ were employed with a parsimony coefficient of 0.001. Each chunk utilized a unique random seed to ensure independent evolutionary searches.

4. gplearn implementation

Classical genetic programming [77] was applied to each of the 18 data chunks independently. Evolutionary parameters included a population size of 75 000 individuals, 1500 generations, and a fitness convergence threshold of 0.001. Genetic operators were configured with crossover probability of 0.7, a subtree mutation probability of 0.1, a hoist mutation probability of 0.05, and a point mutation probability of 0.1. Tournament selection with a tournament size of 25 was employed. The function set included standard arithmetic operators $\{+, -, \times, \div\}$ augmented with protected functions: protected division (threshold 0.001), protected logarithm (offset 1×10^{-8}), protected square root, and bounded power operations $[-3, 3]$. A parsimony coefficient of 0.01 controlled equation complexity.

D. Physics-informed approach

This symbolic regression approach is physics-informed, as it derives explicit algebraic expressions using structural descriptors that govern hydrogen adsorption behavior. By selecting descriptors such as pore volume, void fraction, density, and surface area—each of which influences adsorption through well-established thermodynamic principles—the resulting equations reflect trends consistent with adsorption physics. For example, larger pore volumes accommodate greater gas storage, higher void fractions increase accessibility, and higher densities improve volumetric performance. Although the models do not explicitly solve thermodynamic equations, they capture the underlying trade-offs and saturation behaviors observed in classical adsorption theory. This ensures that the derived equations remain interpretable, physically meaningful, and broadly applicable across diverse MOF structures.

E. Validation and evaluation protocol

1. SISO equation validation

To validate the SISO-derived equations, a two-step screening process was implemented. The first step involved positivity enforcement, ensuring that all predicted H_2 capacities were physically meaningful and strictly non-negative. For this step, the unseen dataset of about 600 000 MOFs was used solely as a plausibility filter: only equations yielding non-negative capacities across this broader design space were retained. The second step involved performance benchmarking, where we retained only those equations capable of reproducing the actual usable volumetric hydrogen capacities with coefficient of determination (R^2), mean absolute error (MAE), and root mean square error (RMSE) values all simultaneously within 0.01, 1.0 g- H_2 L $^{-1}$, and 1.0 g- H_2 L $^{-1}$, respectively, of the highest-performing single-term equations.

2. Cross-method comparative evaluation

All benchmarking algorithms processed identical data chunks with consistent 80-20 train-test splits using fixed random seeds for reproducibility. Standardized evaluation metrics (R^2 , MAE, RMSE) enabled direct performance comparison. Feature standardization was applied where algorithmically required. Results were aggregated across all chunks per target, enabling systematic comparison of method performance, equation interpretability, and computational requirements. Final equations from all methods were benchmarked against GCMC-calculated data to assess generalizability across diverse MOF subclasses and pressure conditions.

III. RESULTS AND DISCUSSION

A. Predictive performance of SISO-derived equations for usable gravimetric H_2 capacity

From the 660 SISO-generated equations, only a fraction provided physically meaningful predictions for usable gravimetric H_2 capacity (C_g , wt %). Predictive performance was evaluated on the test dataset of 84 400 MOFs, while the unseen dataset of about 600 000 MOFs was employed only to enforce positivity checks. This distinction ensured that retained equations not only achieved high accuracy on the test set but also produced physically meaningful (non-negative) predictions across a much broader design space. Among them, 19 single-term, 31 two-term, and 8 three-term equations successfully predicted positive C_g values, while the remaining models either yielded negative predictions or redundant expressions. The fact that most equations failed to generalize across diverse MOFs highlights the importance of dataset curation in developing broadly applicable predictive equations.

A consistent trend emerged across all equation categories: pore-limiting diameter descriptors (D_i) were absent

from nearly all high-performing models, confirming that pore size played a limited role in governing usable H_2 storage under practical operating conditions. Additionally, the top-ranked models did not rely primarily on surface area descriptors (S_g, S_v), aligning with our previous ML studies that identified pore volume (V_p) and void fraction (F_v) as the dominant predictors of C_g [7].

To quantify descriptor interdependence, we computed Spearman (ρ) and Pearson (r) correlation matrices (Figs. S1 and S2 within the Supplemental Material [93]). We summarize monotonic association with $|\rho|$ (the absolute value of Spearman's ρ ; $|\rho| \in [0, 1]$), and retain the sign for Pearson's r to capture linear trends. Porosity descriptors exhibit very strong monotonic collinearity— $|\rho|(\rho_c, V_p) = 0.997$; $|\rho|(\rho_c, F_v) = 0.942$; $|\rho|(S_g, V_p) = 0.936$; $|\rho|(V_p, F_v) = 0.963$ —whereas the corresponding signed linear correlations are more moderate [$r(\rho_c, V_p) = -0.611$; $r(S_g, V_p) = 0.678$; $r(V_p, F_v) = 0.621$], consistent with nonlinear or saturating relationships. The pairwise scatter matrix (Fig. S3 within the Supplemental Material [93]) confirms these curvilinear patterns. Because V_p , F_v , ρ_c , and S_g encode overlapping information about accessible pore space or free volume, adding additional correlated terms yields little incremental explanatory variance, which accounts for the similar performance of the five-parameter gravimetric model and the compact three-parameter form.

Among the 19 single-term equations, the most accurate model relied exclusively on V_p and F_v , achieving an R^2 of 0.991, an MAE of 0.24 wt %, and an RMSE of 0.35 wt %. The simplicity of this model makes it particularly attractive for experimental and computational applications, as both V_p and F_v can be determined with lower cost and greater accuracy than surface area or pore diameter measurements. Unlike S_g , which depends on the diameter of the probe molecule used for measurement, V_p and F_v can be computed using Zeo++ in a probe-independent manner, reducing measurement variability [94]. Extending the model complexity to two-term equations improved accuracy while maintaining interpretability. Out of 660 equations, only 31 generated positive C_g values across both datasets. The top 12 models showed consistently strong performance, with R^2 values near 0.99, MAEs of 0.20–0.24 wt %, and RMSEs of 0.28–0.40 wt %. The highest-ranked two-term equation incorporated ρ_c, F_v, V_p , and D_i , achieving an R^2 of 0.994, MAE of 0.20 wt %, and RMSE of 0.28 wt %. Consistent with our prior studies [7,72], we find that substituting descriptors such as D_i with S_v yields similarly accurate predictions of C_g . This reinforces the notion that multiple crystallographic descriptors can act as interchangeable proxies for capturing porosity effects in MOFs, offering flexibility in model formulation without compromising predictive performance. A simplified model containing only ρ_c and V_p exhibited slightly lower accuracy ($R^2 = 0.980$, RMSE = 0.52 wt %),

TABLE I. Equations for predicting usable gravimetric capacities of MOFs under pressure swing conditions between 100 and 5 bar at 77 K.

No. of terms	Variables	Equation	R^2	MAE (wt %)	RMSE (wt %)
1	V_p, F_v	$C_g = A_{11}^g F_v^{1/3} V_p \exp(-1/V_p) + A_{12}^g$ (1)	0.991	0.24	0.35
2	ρ_c, V_p, F_v, D_i	$C_g = B_{11}^g \frac{F_v^3 \exp(-\rho_c/V_p)}{\sqrt{\rho_c}} + B_{12}^g \frac{D_i^4 F_v}{\sqrt[3]{\exp(D_i)}} + B_{13}^g$ (2)	0.994	0.20	0.28
2	ρ_c, V_p, F_v, S_v	$C_g = B_{21}^g \sqrt[3]{V_p} F_v^2 \exp(-1/V_p) + B_{22}^g \frac{V_p^3 \exp(-V_p)}{\rho_c S_v \ln(F_v)} + B_{23}^g$ (2a)	0.993	0.22	0.30
2	ρ_c, V_p	$C_g = B_{31}^g \rho_c^4 V_p^{13/3} + B_{32}^g \frac{\exp(-1/V_p) \ln(\rho_c)}{\rho_c V_p} + B_{33}^g$ (2b)	0.980	0.25	0.52
3	$\rho_c, V_p, F_v, S_g, D_i$	$C_g = C_{11}^g \rho_c V_p^{3/2} F_v^2 \ln(S_g) + C_{12}^g \frac{D_i^5}{\sqrt[3]{\exp(D_i)}} + C_{13}^g \frac{\exp(-V_p^2) \sqrt{S_g}}{\ln(F_v)} + C_{14}^g$ (3)	0.994	0.19	0.28

TABLE II. Coefficients and intercepts of equations presented in Table I.

Equation	Coefficient of first term	Coefficient of second term	Coefficient of third term	Intercept
(1)	$A_{11}^g = 12.6$	—	—	$A_{12}^g = 0.247$
(2)	$B_{11}^g = 8.15$	$B_{12}^g = 3.67 \times 10^{-3}$	—	$B_{13}^g = 5.54 \times 10^{-2}$
(2a)	$B_{21}^g = 13.4$	$B_{22}^g = 26.7$	—	$B_{23}^g = 0.106$
(2b)	$B_{31}^g = 6.92$	$B_{32}^g = -4.08$	—	$B_{33}^g = 0.147$
(3)	$C_{11}^g = 0.931$	$C_{12}^g = 1.15 \times 10^{-4}$	$C_{13}^g = 1.15 \times 10^{-4}$	$C_{14}^g = 0.2$

confirming that additional structural descriptors refine predictions without overfitting.

Several descriptors in Tables I and II are algebraically related—most notably $V_p = F_v/\rho_c$ and $S_v = \rho_c S_g = F_v S_g/V_p$. Consequently, equations that use V_p explicitly or, equivalently, F_v/ρ_c , and those that use S_v or $\rho_c S_g$, encode the same dimensional groups controlling accessible pore space per unit mass. Together with the strong monotonic associations among V_p , F_v , ρ_c , and S_g (see Figs. S1–S3 within the Supplemental Material [93]), this explains why alternative term arrangements yield similar R^2 : the models capture the same underlying adsorption physics despite differing algebraic form.

Expanding to three-term equations led to further refinements in predictive accuracy. Only 8 of the 660 three-term equations provided valid predictions across both datasets.

The top five models all produced R^2 values near 0.99, with MAEs between 0.19 and 0.20 wt%, and RMSEs between 0.28 and 0.33 wt%. The best-performing model incorporated ρ_c , V_p , F_v , S_g , and D_i , achieving the lowest RMSE among all models. Notably, an alternative three-term equation containing only ρ_c , V_p , and S_g delivered nearly identical predictive accuracy, suggesting that feature redundancy plays a role in stabilizing SISSO-generated equations [72].

The predictive performance of the SISSO-derived symbolic equations is illustrated in Fig. 1, which compares predicted and GCMC-calculated C_g for a dataset of 84 400 MOFs. Figures 1(a)–1(c) present scatter plots for three representative equations: Eq. (1) (a single-term model with two descriptors: V_p and F_v), Eq. (2) (a two-term model incorporating four descriptors: ρ_c , V_p , F_v , and D_i),




FIG. 1. Correlation between GCMC-calculated and Equation-predicted usable gravimetric hydrogen capacity (C_g , wt %) and corresponding error distributions for three representative predictive equations. (a)–(c) Parity plots comparing predicted versus actual C_g for Eqs. (1)–(3), showing strong linear agreement and high predictive accuracy. Dashed lines indicate the ideal 1:1 correlation. (d)–(f) Residual distributions with Gaussian fits, annotated with the number of samples (N), mean error (μ), standard deviation (σ), and median error. Equations (1) and (2) yield narrow unbiased residuals centered near zero, while Eq. (3) shows systematic underprediction.

and Eq. (3) (a more complex three-term model involving five descriptors, including S_g). All three equations exhibit strong linear correlation with minimal dispersion around the 1:1 reference line, demonstrating high predictive accuracy across a structurally diverse set of MOFs.

Systematic error analysis for high-capacity MOFs ($C_g > 25.0$ wt %, $N = 279$ MOFs, 0.32% of the dataset) reveals model-dependent prediction behavior. Single-term equations [Eq. (1)] systematically overpredict [mean error (2.31 ± 1.10) wt %]. Two-term models [Eq. (2)]—a multidescrptor form that integrates more than one crystallographic descriptor into the same algebraic expression, thereby capturing coupled physical effects that single-descriptor scaling cannot—achieve improved accuracy [mean error (0.38 ± 0.98) wt %] but still underpredict 33% of cases by (0.73 ± 0.70) wt % on average. Three-term equations [Eq. (3)] predominantly underpredict, affecting 55% of high-capacity MOFs by (1.57 ± 1.25) wt % on average. These deviations correlate with extreme structural parameters ($V_p > 2.85$ cm³ g⁻¹, $F_v > 0.88$), indicating that physics beyond geometric scaling becomes critical at performance limits. Accordingly, multidescrptor equations [e.g., Eq. (2)] are essential for mitigating errors in

edge cases corresponding to the highest-capacity MOFs. In this regime, single-descriptor scaling breaks down, and accurate prediction requires simultaneous accounting for porosity and density constraints.

Notably, Eq. (1) achieves excellent agreement with GCMC predictions across the majority of the dataset, highlighting the predictive power of a simple interpretable model. The corresponding error distributions [Figs. 1(d)–1(f)] further affirm the robustness of the models. The residuals are tightly centered around zero and follow approximately Gaussian distributions, with the vast majority of errors confined within ± 1.0 wt %. This behavior indicates both low bias and high precision across models. A comparison of Eqs. (1)–(3) reveals that reducing the number of descriptors through feature substitution or simplification does not significantly compromise predictive accuracy. This is consistent with our previous findings [7,72] and underscores the ability of symbolic regression to identify compact and physically meaningful expressions that retain generalizability.

Interestingly, the strong performance of Eq. (1), despite its parsimony, suggests that pore volume and void fraction are dominant drivers of hydrogen uptake, while other

descriptors such as pore diameter contribute marginally. The inclusion of additional descriptors in Eqs. (2) and (3) yields marginal improvements in performance for high-capacity MOFs, but for the majority of the dataset, simpler models suffice. These results reinforce the utility of physics-informed symbolic modeling in balancing interpretability, accuracy, and computational efficiency for rapid screening and design of MOFs for hydrogen storage.

B. Predictive performance of SISO-derived equations for usable volumetric H₂ capacity

The SISO-derived equations for usable volumetric H₂ capacity (C_v , g-H₂ L⁻¹) exhibit notable differences from those predicting usable gravimetric capacity (C_g , wt %). The fundamental contrast lies in the nature of the dominant descriptors: while C_g is mainly governed by porosity and void fraction, C_v also depends heavily on density-related parameters. This distinction reflects an inherent trade-off between maximizing hydrogen storage and maintaining sufficient framework density to ensure volumetric efficiency.

Of the 660 SISO-generated equations, 17 one-term, 94 two-term, and 9 three-term equations successfully predicted positive C_v values across both the test dataset (84 400 MOFs) and the unseen dataset (~600 000 MOFs). The greater number of valid two-term expressions for C_v compared to C_g suggests that volumetric capacity requires a broader set of descriptors to capture its structural variability—likely due to the additional dependence on ρ_c .

The best-performing one-term equation [Eq. (4) in Table III] relied solely on F_v , achieving an R^2 of 0.948, an MAE of 1.57 g-H₂ L⁻¹, and an RMSE of 2.20 g-H₂ L⁻¹. This is in contrast to C_g , where the one-term model used both V_p and F_v , highlighting that pore volume played a less dominant role in volumetric capacity predictions. The slightly reduced accuracy of this equation compared to its gravimetric counterpart ($R^2 = 0.948$ vs. 0.991) suggests

that a single feature is insufficient to fully capture volumetric trends, as C_v is inherently influenced by both porosity and single-crystal density.

Introducing a second term improved the predictive performance, with 94 two-term equations successfully predicting C_v values. The best-performing two-term equation [Eq. (5)] incorporates ρ_c , V_p , and F_v , yielding an R^2 of 0.955, with an MAE of 1.47 g-H₂ L⁻¹ and an RMSE of 2.05 g-H₂ L⁻¹. Unlike C_g , where pore diameter descriptors were mostly absent from high-performing models, D_i appears in the top volumetric equations (not shown in Tables III and IV), suggesting that pore size may contribute to space efficiency when assessing H₂ adsorption in compact environments.

Expanding to three-term models led to modest improvements in prediction accuracy. Of the 660 SISO-generated equations, 9 three-term expressions successfully predicted physically meaningful C_v values (not shown in Tables III and IV). This indicates that volumetric capacity can benefit from increased model complexity, though only select combinations of descriptors yield high-performing results. Unlike C_g , where S_v did not appear in top-ranked models, S_v contributes meaningfully to C_v predictions. This underscores the importance of available adsorption space per unit volume rather than just mass.

Comparing the predictive models for C_v and C_g reveals clear structural distinctions. Gravimetric models tend to prioritize V_p and F_v , emphasizing the role of maximizing porosity without accounting for material density. In contrast, C_v models incorporate ρ_c , highlighting the need to balance porosity with crystallographic density to improve volumetric performance. This interpretation is quantitatively confirmed by the SISO-derived equations in Table III. Equation (5) demonstrates this through its second term, $157.83 \times (-F_v^3 + F_v^2) \ln(\rho_c V_p)$, where density appears within the logarithmic function coupled with pore volume, weighted by a void fraction

TABLE III. Equations for predicting usable volumetric capacities of MOFs under pressure swing conditions between 100 and 5 bar at 77 K.

No. of terms	Variables	Equation	R^2	MAE (g-H ₂ L ⁻¹)	RMSE (g-H ₂ L ⁻¹)
1	F_v	$C_v = A_{11}^v F_v^{17/2} \ln(F_v) + A_{12}^v$ (4)	0.948	1.57	2.20
2	ρ_c, V_p, F_v	$C_v = B_{21}^v F_v^8 \ln(F_v) + B_{22}^v (-F_v^3 + F_v^2) \ln(\rho_c V_p) + B_{23}^v$ (5)	0.955	1.47	2.05

TABLE IV. Coefficients and intercepts of equations presented in Table III.

Equation	Coefficient of first term	Coefficient of second term	Coefficient of third term	Intercept
4	$A_{11}^v = -769.619$			$A_{12}^v = 4.520$
5	$B_{21}^v = -487.333$	$B_{22}^v = 157.83$		$B_{23}^v = 16.615$

term $F_v^2(1 - F_v)$. The positive coefficient (157.83) on this density-containing term validates that higher framework density contributes positively to volumetric capacity. The logarithmic dependence $\ln(\rho_c V_p)$ reveals that volumetric capacity follows a complex, nonlinear relationship rather than simple power laws, capturing the saturation effects that occur as materials approach physical limits of both density and porosity. This mathematical form elegantly encodes the fundamental challenge of volumetric storage: maximizing both the amount of material per unit volume (ρ_c) and the accessible storage space (V_p), which can be made more explicit through the relationship $\rho_c = M_{\text{framework}}/(V_p/F_v)$, where $M_{\text{framework}}$ is the MOF framework mass, yielding $\rho_c V_p = M_{\text{framework}}/F_v$, and therefore, $\ln(\rho_c V_p) = \ln(M_{\text{framework}}/F_v)$. This formulation directly quantifies the trade-off: heavier frameworks provide more adsorption sites but reduce void space, while highly porous MOFs offer greater accessible volume but fewer adsorption sites per unit volume. The logarithmic form indicates this balance follows a multiplicative relationship ($M_{\text{framework}}/F_v$) with diminishing returns, with the logarithmic scaling reflecting the diminishing returns as either parameter reaches extreme values. As this mathematical analysis confirms, this distinction suggests that

MOFs optimized for high gravimetric capacity—often highly porous—may not necessarily perform well volumetrically, making the simultaneous optimization of C_g and C_v a complex and nontrivial design challenge.

Figures 2(a) and 2(b) compare GCMC-calculated and SISSO-predicted C_v for two equations of increasing complexity [Eq. (4) vs. Eq. (5)]. Both track the 1:1 line closely ($R^2 \approx 0.95$), confirming that compact interpretable forms capture the dominant structure-property relationships. Higher-complexity Eq. (5) (three descriptors + bias) yields a modest but consistent gain, especially at larger capacities, with slightly lower errors (RMSE ≈ 2.04 vs. 2.20 g-H₂ L⁻¹; MAE ≈ 1.45 vs. 1.53 g-H₂ L⁻¹). Residual histograms in Figs. 2(c) and 2(d) (GCMC calculated – Equation-predicted) are narrowly centered near zero with approximately Gaussian shapes, indicating minimal systematic bias; the spread is somewhat broader than that for C_g , consistent with the added sensitivity of C_v to density and porosity alongside volumetric surface area. Overall, the strong performance of parsimonious Eq. (4) and the incremental benefit of Eq. (5) highlight a favorable simplicity-accuracy trade-off for C_v prediction.

Despite the slightly lower R^2 values for volumetric predictions compared to gravimetric models, the




FIG. 2. Correlation between GCMC-calculated and Equation-predicted usable volumetric hydrogen capacity (C_v , g-H₂ L⁻¹) and corresponding error distributions for two representative predictive equations. (a),(b) Parity plots comparing predicted versus actual C_v for Eqs. (4) and (5), showing strong linear agreement and good predictive accuracy. Dashed lines indicate the ideal 1:1 correlation. (c),(d) Residual distributions with Gaussian fits, annotated with the number of samples (N), mean error (μ), standard deviation (σ), and median error. Both equations yield residuals centered near zero with similar spreads, reflecting consistent predictive robustness across diverse MOF structures.

derived equations provide an efficient framework for C_v estimation. Their ability to replace computationally intensive GCMC simulations makes them valuable tools for large-scale screening of MOFs targeting vehicular H_2 storage applications, where volumetric efficiency is a key constraint. The ability to estimate C_v directly from crystallographic descriptors eliminates the need for high-throughput molecular simulations, enabling rapid preselection of promising MOFs before experimental synthesis and validation.

C. Real-world validation

To assess generalization beyond our computational training dataset, we compiled an independent set of 15 MOFs reported across two publications under identical operating conditions for usable hydrogen capacity (77 K, 5–100 bar). For each material, the sources provide experimentally measured usable capacity together with three crystallographic descriptors: ρ_c , V_p , and S_g . For 14 of the 15 MOFs, the same studies additionally report

computed V_p and S_g , enabling parallel evaluation with both experimental and computational descriptors. Missing descriptors required by our equations were calculated from available data ($S_v = \rho_c \times S_g$, $F_v = \rho_c \times V_p$). Because pore diameters were not reported, validation was restricted to equations independent of these parameters [C_g Eqs. (1) and (2b); C_v Eqs. (4) and (5)]. Coefficients were held fixed (no refitting).

Figure 3 presents predicted-versus-measured performance on this external set. For C_g , both equations showed robust performance with $R^2 = 0.844$ [Eq. (1)] and 0.905 [Eq. (2b)] on experimental descriptors, and $R^2 = 0.941$ – 0.950 on computational descriptors, reflecting systematic differences between measurement approaches. C_v was more challenging on this dataset [e.g., Eq. (4): $R^2 = 0.342$], consistent with the greater sensitivity of volumetric metrics to small uncertainties in ρ_c and V_p . Table SI within the Supplemental Material [93] reports R^2 , RMSE, and MAE for all equation-descriptor pairings. Overall, the independent experimental validation confirms that the symbolic equations retain reliable generalization capability




FIG. 3. Real-world validation of SISSO-derived equations on independent experimental MOF dataset. (a), (b) C_g predictions using Eqs. (1) and (2b), showing strong correlation across computational (blue circles) and experimental (red triangles) descriptor sets. (c), (d) C_v predictions using Eqs. (4) and (5) demonstrating reasonable trend capture despite increased scatter. Dashed lines represent perfect prediction.

to experimental MOFs not used in model development under realistic hydrogen storage conditions.

D. Comparative symbolic regression methodologies

To establish methodological benchmarks and validate our approach within the broader symbolic regression landscape, we systematically evaluated three alternative frameworks: AI Feynman (neural-symbolic hybrid) [80], PySR (evolutionary optimization) [75,76], and gplearn (genetic programming) [77]. All methods utilized identical physics-based crystallographic descriptors to ensure fair comparison across methodological approaches.

Table V presents a comparison of the performance across all methodologies for both C_g and C_v predictions. For C_g prediction, all methodologies achieved excellent statistical performance within a narrow range ($R^2 = 0.991$ – 0.998), demonstrating that the physics-based descriptor selection effectively captured the underlying structure-property relationships. The evolutionary optimization of PySR achieved the highest R^2 (0.998), followed by AI Feynman’s neural-symbolic approach ($R^2 = 0.996$) and gplearn’s genetic programming ($R^2 = 0.995$). SISSO demonstrated competitive performance ($R^2 = 0.991$) while maintaining critical advantages in interpretability and deployability.

C_v predictions revealed systematic challenges across all methodologies. This degradation manifested universally: PySR decreased from $R^2 = 0.998$ (C_g) to 0.985 (C_v), AI Feynman from 0.996 to 0.979, gplearn from 0.995 to 0.984, and SISSO from 0.991 to 0.960. Unlike C_g , which primarily depends on accessible pore volume and void fraction, C_v requires simultaneous optimization of conflicting properties: maximizing porosity while maintaining sufficient framework density.

While evolutionary approaches (PySR, gplearn) achieved marginally higher statistical metrics through extensive stochastic searches, they generated complex equations with limited mechanistic insights and variable reproducibility. In contrast, SISSO’s compressed sensing

framework produced compact, theoretically grounded equations that explicitly revealed structure-property relationships. The modest statistical differences ($\Delta R^2 < 0.007$ for C_g) become negligible compared to experimental uncertainties, where SISSO’s interpretability and deployability advantages outweigh marginal statistical improvements.

E. Dimensional analysis and physical interpretation of key equations

To understand the physical mechanisms that correlate with our symbolic regression results, we performed dimensional analysis of the highest-performing equations (Table VI). This analysis reveals that SISSO has identified mathematical relationships that reflect fundamental organizing principles governing hydrogen storage that extend beyond conventional descriptor-property relationships. Unless otherwise stated, all exponential and logarithmic functions act on nondimensionalized variables (e.g., V_p/V_0 , ρ_c/ρ_0 , D_i/D_0), so that the reported coefficients carry the necessary physical units to yield C_g (wt %) and C_v (g-H₂ L⁻¹). We list descriptor dimensions to aid physical interpretation. For example, $\exp(-\rho_c/V_p)$ is shorthand for $\exp[-(\rho_c/\rho_0)/(V_p/V_0)]$ with $\rho_0 = 1 \text{ g cm}^{-3}$ and $V_0 = 1 \text{ cm}^3 \text{ g}^{-1}$, which preserves the numerical values in our units while ensuring a dimensionless argument.

1. Hierarchical optimization architecture

Table VI demonstrates that multiterm equations operate through sequential physical constraints rather than additive contributions. In Eq. (2), the first term, $[F_v^3 \exp(-\rho_c/V_p)]/\sqrt{\rho_c}$, functions as a structural selectivity prefilter based on porosity-density, creating an exponential selectivity landscape that eliminates dense frameworks before the second term, $[D_i^4 F_v]/\sqrt[3]{\exp(D_i)}$, can optimize pore accessibility. This cascade architecture explains why linear descriptor combinations can underperform relative to multiterm cascades for C_g predictions—optimization requires satisfying density-porosity constraints before geometric effects become relevant.

Equation (5) reveals a two-stage optimization cascade for C_v . The dimensional analysis (Table VI) shows that the first term, with eighth-power porosity scaling modulated by a logarithm, correlates with enhanced performance in highly porous frameworks while imposing accessibility limits that prevent unrealistic divergence. The second term, combining quadratic and cubic porosity effects with a logarithmic density-volume factor, introduces empirical corrections for excessive void space and rewards balanced frameworks that couple pore capacity with structural density. Sign selectivity: since $(-F_v^3 + F_v^2 = F_v^3(1 - F_v) > 0)$ for $0 < F_v < 1$, the overall sign of this contribution is set by $\ln(\rho_c V_p)$; it subtracts

TABLE V. Performance comparison across symbolic regression methodologies under 100–5 bar pressure swing at 77 K.

Method	R^2	RMSE	MAE
C_g (wt %)			
SISSO	0.991	0.350	0.240
PySR	0.998	0.159	0.125
AI Feynman	0.996	0.179	0.141
gplearn	0.995	0.223	0.174
C_v (g-H ₂ L ⁻¹)			
SISSO	0.960	2.200	1.570
PySR	0.985	1.007	0.791
gplearn	0.984	1.128	0.819
AI Feynman	0.979	1.228	0.938

TABLE VI. Dimensional analysis and physical interpretation of representative equations.

Eq.	Term	Physical interpretation and expected behavior
Equation (1) (C_g)	$F_v^2 V_p^{1/3} \exp(-1/V_p)$	Physical mechanism: Correlates with fundamental accessibility-porosity coupling with exponential gate function that eliminates materials below critical pore volume thresholds. Moderate void fraction scaling prevents oversensitivity while cube-root pore volume scaling captures diminishing returns.
	Dimensional units, $\text{cm g}^{-1/3}$	Expected behavior: Sharp pore volume threshold, below which hydrogen storage becomes negligible, followed by steady growth with optimal balance between void accessibility and structural integrity.
Equation (2) (C_g)	$\frac{F_v^3 \exp(-\rho_c/V_p)}{\sqrt{\rho_c}}$	Physical mechanism: Captures the gravimetric optimization trade-off—cubic void fraction scaling drives performance but must be balanced against framework density. The exponential density-porosity ratio acts as a structural selectivity prefilter, strongly favoring ultralight highly porous materials while correlating negatively with dense frameworks.
	Dimensional units, $\text{cm}^{3/2} \text{g}^{-1/2}$	Expected behavior: Exponential preference for materials with high porosity-to-density ratios; a performance drop occurs when density overwhelms optimal porosity scaling.
	$\frac{D_i^4 F_v}{\sqrt[3]{\exp(D_i)}}$	Physical mechanism: Emphasizes pore-throat accessibility—quartic pore-diameter scaling yields large gains with increasing openings, while exponential attenuation prevents unphysical growth for macropores where surface area limits dominate.
	Dimensional units, $\text{\AA}^4$	Expected behavior: Sharp gains with increasing D_i until diffusion advantages saturate as surface area constraints dominate.
Equation (4) (C_v)	$F_v^{17/2} \ln(F_v)$	Physical mechanism: Storage dominated by porosity with high sensitivity; the logarithm tempers growth.
	Dimensionless	Expected behavior: Bounded contribution that vanishes as $F_v \rightarrow 0^+$ and $F_v \rightarrow 1^-$; maximal magnitude at $F_v = e^{-1/(17/2)} \approx 0.889$. With the negative fitted coefficient, this term provides a positive peak contribution at high—but not extreme—porosity.
Equation (5) (C_v)	$F_v^8 \ln(F_v)$	Physical mechanism: Strong porosity scaling moderated by a logarithm.
	Dimensionless	Expected behavior: Bounded contribution with a peak at $F_v = e^{-1/8} \approx 0.882$; negligible near the endpoints $F_v \rightarrow 0^+, 1^-$.
	$(-F_v^3 + F_v^2) \ln(\rho_c V_p)$	Physical mechanism: Balances quadratic and cubic porosity effects; the negative cubic penalizes excessive void space. The logarithmic dependence on the density-volume product incorporates mass and capacity trade-offs.
	Dimensionless	Expected behavior: Positive at moderate porosity but declines at higher porosity as penalties grow. Sign selectivity, the term subtracts when $\rho_c V_p < 1$ and adds when $\rho_c V_p > 1$.

when $\rho_c V_p < 1$ and adds when $\rho_c V_p > 1$. The mathematical structure indicates that volumetric optimization requires simultaneous success across both porosity amplification and density-volume trade-offs—explaining the persistent difficulty of achieving high C_v across symbolic regression methodologies.

2. Critical scaling and performance thresholds

The exponential and logarithmic terms identified in Table VI encode sharp performance thresholds rather than continuous scaling relationships. The $\exp(-1/V_p)$ gate

function in Eq. (1) creates a critical pore volume, below which hydrogen storage becomes negligible. In Eq. (4), the combination $F_v^{17/2} \ln(F_v)$ is bounded and vanishes as $F_v \rightarrow 0^+$ and $F_v \rightarrow 1^-$; its magnitude peaks at $F_v = e^{-1/(17/2)} \approx 0.889$, combining of high-power porosity scaling with a logarithmic factor approaching mathematical divergence as porosity nears its upper limit, while simultaneously suppressing performance in low-porosity regimes. In Eq. (5), the $F_v^8 \ln(F_v)$ term has analogous bounded behavior with a peak at $F_v = e^{-1/8} \approx 0.882$, and the quadratic-cubic balance introduces turning points where excessive porosity reduces capacity, unless compensated

for by density-volume contributions. These mathematical structures suggest that hydrogen storage materials operate in threshold-like regimes of descriptor space, where small structural variations in porosity or density can yield dramatic changes in performance.

3. Emergent design principles

Three universal principles emerge from the dimensional analysis:

(1) Threshold-dominated behavior: All high-performing equations contain exponential or logarithmic structures that correlate with sharp selectivity, suppressing performance at low porosity or pore volume while triggering rapid amplification only beyond critical thresholds. This indicates that hydrogen storage is governed by thresholdlike behavior rather than gradual scaling.

(2) Scale-separated optimization: Single-term equations capture optimization within narrow descriptor regimes, while multiterm equations, such as Eqs. (2) and (5), impose sequential filters that couple porosity, density, and geometry. This cascade structure demonstrates that performance requires success across multiple constraint layers, not simply maximizing a single descriptor.

(3) Physical limit enforcement: The high-power porosity scaling in Eq. (4) is naturally bounded and peaks at high—but not extreme—porosity, while the quadratic-cubic balance in Eq. (5) imposes penalties on excessive porosity. These mathematical descriptors correlate with physical constraints and suggest the equations capture realistic material limitations, indicating that optimal materials approach structural feasibility limits.

These findings reveal that successful hydrogen storage optimization requires understanding hierarchical constraint sequences rather than individual descriptor effects. The nonadditive cascade physics captured by SISSO helps explain when simple linear combinations underperform relative to multiterm cascades and suggests that materials near optimal performance will exhibit high sensitivity to synthesis conditions. More broadly, the mathematical structures—hierarchical optimization with critical scaling—represent empirical correlations that may reflect structure-property relationships in materials science, suggesting that symbolic regression may be uncovering statistical patterns that correspond to underlying physical principles.

F. Applicability domain and model limitations

The physics-informed analytical models developed in this work provide interpretable and computationally efficient predictions of usable hydrogen storage capacities in MOFs. However, their reliability is constrained by the characteristics and scope of the training dataset. To define the domain within which these models are valid, Fig. 4




FIG. 4. Distributions of usable H_2 capacities and crystallographic descriptors for the dataset used to develop the analytical models. (a) Usable gravimetric capacity (C_g) and (b) usable volumetric capacity (C_v) are shown alongside the structural descriptors: (c) crystallographic density (ρ_c), (d) gravimetric surface area (S_g), (e) volumetric surface area (S_v), (f) void fraction (F_v), (g) pore volume (V_p), and (h) largest pore diameter (D_i). These distributions define the applicability domain of the derived equations. Models are expected to yield accurate and physically meaningful predictions within these empirical ranges, and caution is advised when extrapolating beyond them.

presents the distributions of crystallographic descriptors and hydrogen uptake values used during training. These distributions reveal the empirical boundaries over which the equations were derived—e.g., C_v values ranging from approximately 0 to 40 $\text{g-H}_2 \text{ L}^{-1}$, and ρ_c between 0.2 and 2.0 g cm^{-3} . Similar bounds apply to surface area, pore volume, void fraction, and pore diameter.

Because the equations were obtained using symbolic regression over this defined domain, they should not be extrapolated beyond these feature ranges. Predictions for

materials outside the training space may result in physically implausible outputs, reduced accuracy, or numerical instability—especially due to nonlinear terms such as exponentials or logarithms present in the equations.

The consistent performance degradation for C_v predictions across all symbolic regression methodologies—from SISO’s deterministic approach to PySR’s extensive evolutionary search—confirms that these challenges reflect fundamental materials physics constraints rather than algorithmic deficiencies. Cross-methodology analysis revealed systematic underprediction for high- C_v MOFs (>25 g-H₂ L⁻¹) across all approaches, suggesting that exceptional C_v performance arose from subtle structural features not captured by standard crystallographic descriptors. Accurate modeling of these edge cases is critical, as they correspond to candidate MOFs approaching U.S. DOE performance targets; multidescrptor forms provide the necessary resolution to guide design in this high-capacity regime.

This limitation affects all methodologies equally, reinforcing that enhanced algorithmic sophistication cannot overcome fundamental descriptor limitations in materials science applications. Unlike C_g , which exhibits near-perfect predictability ($R^2 > 0.99$) across methodologies, C_v predictions consistently plateau below $R^2 = 0.99$, indicating fundamental physical constraints in the density-porosity optimization trade-off.

While the models successfully capture key adsorption trends through interpretable physics-based descriptors, they do not explicitly account for molecule-specific interactions, such as polarization or open metal site effects. These limitations may affect prediction accuracy for certain materials. Future improvements to volumetric predictions could involve several promising approaches. A particularly attractive avenue involves hybrid methods combining symbolic regression with machine learning. For instance, quantum-informed descriptors capturing specific chemical interactions (e.g., metal-H₂ binding energies, polarization effects, open metal site characteristics) could be incorporated as additional descriptors in the SISO framework, potentially improving the accuracy of volumetric capacity predictions where the current models show slightly lower R^2 values. Alternatively, symbolic equations could serve as physically motivated basis functions in neural network architectures, preserving interpretability while capturing complex nonlinearities that may be particularly important for volumetric storage optimization. Such hybrid approaches could also better account for framework flexibility effects and pore shape dependencies that become critical at high pressures, addressing some of the limitations observed in the current symbolic models. Thus, Fig. 4 serves as a practical reference, delineating the valid operational range for each input variable and providing a clear boundary map for safe and meaningful model deployment in MOF screening workflows.

G. Adaptability and design-guided optimization

A key advantage of the equations developed lies in their adaptability to evolving datasets. Grounded in symbolic regression and physically meaningful descriptors, these models can seamlessly incorporate new MOFs and hydrogen storage data, enabling ongoing refinement of structure-property relationships and broader generalization across chemical and structural spaces. Unlike black-box machine learning models or purely empirical fits, the derived equations encode the underlying physics of adsorption, explicitly linking MOF structure to hydrogen uptake. This not only improves interpretability but also enhances predictive confidence.

Beyond prediction, these equations serve as a valuable tool for design-guided optimization. Their explicit algebraic form relates performance metrics to specific structural parameters, supporting rational tuning of MOF descriptors to achieve target storage capacities. For example, maximizing void fraction and pore volume tends to favor high C_g , while C_v requires a more delicate balance between porosity and crystallographic density. Such insights empower more informed material design decisions than traditional heuristic-based screening.

Moreover, the interpretable and compact nature of these expressions makes them ideally suited for integration into generative and inverse design workflows. When embedded in high-throughput computational pipelines, they can rapidly guide the discovery of MOFs with tailored hydrogen storage capabilities—while maintaining physical realism, which is often lost in purely data-driven approaches. In summary, these symbolic models bridge the gap between empirical screening and computationally intensive simulations. They offer a scalable, interpretable, and physics-informed approach for accelerating the design and discovery of high-performance MOFs for hydrogen storage applications.

IV. CONCLUSIONS

This study introduces a symbolic regression-based framework for predicting usable hydrogen storage capacities in MOFs, addressing key limitations of empirical heuristics and black-box machine learning models. Leveraging the SISO algorithm, we derive compact and interpretable equations from a dataset of 88 400 MOFs, enabling the identification of key structural descriptors that govern hydrogen uptake under 100-to-5 bar pressure swing conditions at 77 K. These equations were rigorously validated against GCMC-calculated data, demonstrating high predictive accuracy and broad applicability across diverse MOF chemistries and architectures.

Cross-methodology benchmarking against AI Feynman, PySR, and gplearn confirmed that SISO achieved competitive statistical performance ($R^2 = 0.991$ for C_g , 0.960 for C_v) while providing superior interpretability and

deployability advantages. The narrow performance range across all methods ($\Delta R^2 < 0.007$ for C_g) demonstrates that physics-based descriptor selection effectively captures underlying structure-property relationships, with universal C_v prediction challenges confirming fundamental material constraints rather than algorithmic limitations.

For C_g , most predictive models consistently relied on V_p and F_v , reinforcing that surface area alone was insufficient for capturing performance trends. In contrast, C_v predictions emphasized ρ_c and S_v , highlighting the trade-off between porosity and packing density that underpinned compact storage. The best-performing single-term equations achieved R^2 values exceeding 0.99 for C_g and 0.96 for C_v , with multiterm models offering further improvements in accuracy.

Beyond predictive performance, a major advantage of the SISSO-derived equations lies in their on-demand usability. Unlike traditional ML models or molecular simulations, these algebraic expressions require only basic crystallographic descriptors—available directly from structural databases—and can be evaluated using spreadsheets, scripts, or even by hand. This enables instantaneous simulation-free screening of large MOF libraries, making the approach especially valuable for experimentalists and design engineers seeking rapid feasibility assessments without computational overheads.

These equations also offer an interpretable design tool, directly linking structural parameters to performance and enabling rational optimization of MOF descriptors for target storage capacities. Their compact form makes them ideal for integration into inverse design algorithms, high-throughput pipelines, and experimental workflows aimed at next-generation hydrogen storage materials.

In summary, this work establishes a physics-informed data-driven paradigm for MOF screening that bridges the gap between empirical models and computationally intensive simulations. By embedding adsorption physics into symbolic regression, we provide a robust, scalable, and interpretable alternative to black-box ML. The resulting equations not only enhance predictive accuracy but also support rapid, accessible, and scientifically grounded MOF discovery—opening a path toward practical high-performance hydrogen storage technologies for clean energy and transportation applications.

ACKNOWLEDGMENTS

This work was partly supported by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, under Grant No. DE-EE0007046, and by the University of Michigan Graham Sustainability Institute. Computational resources and services were provided in part by Advanced Research Computing—Technology Services (ARC-TS), a division of Information and Technology Services (ITS) at the University of Michigan, Ann Arbor.

The authors gratefully acknowledge Professors Randall Snurr and Tom Woo for providing access to their MOF databases, and Dr. Maciej Haranczyk for the use of the Zeo++ code and the mail-order MOF database.

DATA AVAILABILITY

The data that support the findings of this article are openly available [89].

- [1] G. Glenk and S. Reichelstein, Economics of converting renewable power to hydrogen, *Nat. Energy* **4**, 216 (2019).
- [2] S. van Renssen, The hydrogen solution?, *Nat. Clim. Change* **10**, 799 (2020).
- [3] N. Mac Dowell, N. Sunny, N. Brandon, H. Herzog, A. Y. Ku, W. Maas, A. Ramirez, D. M. Reiner, G. N. Sant, and N. Shah, The hydrogen economy: A pragmatic path forward, *Joule* **5**, 2524 (2021).
- [4] A. Majumdar, J. M. Deutch, R. S. Prasher, and T. P. Griffin, A framework for a hydrogen economy, *Joule* **5**, 1905 (2021).
- [5] A. Ahmed, S. Seth, J. Purewal, A. G. Wong-Foy, M. Veenstra, A. J. Matzger, and D. J. Siegel, Exceptional hydrogen storage achieved by screening nearly half a million metal-organic frameworks, *Nat. Commun.* **10**, 1568 (2019).
- [6] M. D. Allendorf, *et al.*, An assessment of strategies for the development of solid-state adsorbents for vehicular hydrogen storage, *Energy Environ. Sci.* **11**, 2784 (2018).
- [7] A. Ahmed and D. J. Siegel, Predicting hydrogen storage in MOFs via machine learning, *Patterns* **2**, 100291 (2021).
- [8] A. Ahmed, Y. Liu, J. Purewal, L. D. Tran, M. Veenstra, A. Wong-Foy, A. Matzger, and D. Siegel, Balancing gravimetric and volumetric hydrogen density in MOFs, *Energy Environ. Sci.* **10**, 2459 (2017).
- [9] M. A. Pellow, C. J. M. Emmott, C. J. Barnhart, and S. M. Benson, Hydrogen or batteries for grid storage? A net energy analysis, *Energy Environ. Sci.* **8**, 1938 (2015).
- [10] M. D. Allendorf, V. Stavila, J. L. Snider, M. Witman, M. E. Bowden, K. Brooks, B. L. Tran, and T. Autrey, Challenges to developing materials for the transport and storage of hydrogen, *Nat. Chem.* **14**, 1214 (2022).
- [11] I. A. Hassan, H. S. Ramadan, M. A. Saleh, and D. Hissel, Hydrogen storage technologies for stationary and mobile applications: Review, analysis and perspectives, *Renewable Sustainable Energy Rev.* **149**, 111311 (2021).
- [12] C. Moran, P. Deane, S. Yousefian, and R. F. D. Monaghan, The hydrogen storage challenge: Does storage method and size affect the cost and operational flexibility of hydrogen supply chains?, *Int. J. Hydrogen Energy* **52**, 1090 (2024).
- [13] D. D. Papadias and R. K. Ahluwalia, Bulk storage of hydrogen, *Int. J. Hydrogen Energy* **46**, 34527 (2021).
- [14] Z. Chen, K. O. Kirlikovali, K. B. Idrees, M. C. Wasson, and O. K. Farha, Porous materials for hydrogen storage, *Chemistry* **8**, 693 (2022).
- [15] Z. Chen, Penghao Li, Ryther Anderson, Xingjie Wang, Xuan Zhang, Lee Robison, Louis R. Redfern, Shinya Moribe, Timur Islamoglu, Diego A. Gómez-Gualdrón,

- Taner Yildirim, J. Fraser Stoddart, and Omar K. Farha, Balancing volumetric and gravimetric uptake in highly porous materials for clean energy, *Science* **368**, 297 (2020).
- [16] N. L. Rosi, J. Eckert, M. Eddaoudi, D. T. Vodak, J. Kim, M. O’Keeffe, and O. M. Yaghi, Hydrogen storage in microporous metal-organic frameworks, *Science* **300**, 1127 (2003).
- [17] L. J. Murray, M. Dinc, and J. R. Long, Hydrogen storage in metal-organic frameworks, *Chem. Soc. Rev.* **38**, 1294 (2009).
- [18] M. P. Suh, H. J. Park, T. K. Prasad, and D. W. Lim, Hydrogen storage in metal-organic frameworks, *Chem. Rev.* **112**, 782 (2012).
- [19] M. Dincă and J. R. Long, Hydrogen storage in microporous metal-organic frameworks with exposed metal sites, *Angew. Chem., Int. Ed.* **47**, 6766 (2008).
- [20] J. L. C. Rowsell and O. M. Yaghi, Strategies for hydrogen storage in metal-organic frameworks, *Angew. Chem., Int. Ed.* **44**, 4670 (2005).
- [21] A. Schoedel, Z. Ji, and O. M. Yaghi, The role of metal-organic frameworks in a carbon-neutral energy cycle, *Nat. Energy* **1**, 16034 (2016).
- [22] A. G. Wong-Foy, A. J. Matzger, and O. M. Yaghi, Exceptional H₂ saturation uptake in microporous metal-organic frameworks, *J. Am. Chem. Soc.* **128**, 3494 (2006).
- [23] D. Chakraborty, A. Yurdusen, G. Mouchaham, F. Nouar, and C. Serre, Large-scale production of metal-organic frameworks, *Adv. Funct. Mater.* **34**, 2309089 (2024).
- [24] DOE Technical Targets for Onboard Hydrogen Storage for Light-Duty Vehicles, <https://energy.gov/eere/fuelcells/doe-technical-targets-onboard-hydrogen-storage-light-duty-vehicles>.
- [25] X. Zhang, R. Lin, J. Wang, B. Wang, B. Liang, T. Yildirim, J. Zhang, W. Zhou, and B. Chen, Optimization of the pore structures of MOFs for record high hydrogen volumetric working capacity, *Adv. Mater.* **32**, 1907995 (2020).
- [26] D. A. Gomez-Gualdrón, T. C. Wang, P. García-Holley, R. M. Sawelewa, E. Argueta, R. Q. Snurr, J. T. Hupp, T. Yildirim, and O. K. Farha, Understanding volumetric and gravimetric hydrogen adsorption trade-off in metal-organic frameworks, *ACS Appl. Mater. Interfaces* **9**, 33419 (2017).
- [27] D. Zhao, X. Wang, L. Yue, Y. He, and B. Chen, Porous metal-organic frameworks for hydrogen storage, *Chem. Commun.* **58**, 11059 (2022).
- [28] B. Mouriño, S. Majumdar, X. Jin, F. McIlwaine, J. van Herck, A. Ortega-Guerrero, S. Garcia, and B. Smit, Exploring the chemical design space of metal-organic frameworks for photocatalysis, *Chem. Sci.* **16**, 11434 (2025).
- [29] S. M. Moosavi, A. Nandy, K. M. Jablonka, D. Ongari, J. P. Janet, P. G. Boyd, Y. Lee, B. Smit, and H. J. Kulik, Understanding the diversity of the metal-organic framework ecosystem, *Nat. Commun.* **11**, 4068 (2020).
- [30] J. Kim, D. Nam, H. J. Cho, E. Cho, D. Sivanesan, C. Cho, J. Lee, J. Kim, and W. Choe, Up-down approach for expanding the chemical space of metal-organic frameworks, *Nat. Synth.* **3**, 1518 (2024).
- [31] T.-W. Liu, Q. Nguyen, A. Bouso Dieng, D. A. G. Omez-Gualdrón, and G. Gualdrón, Diversity-driven, efficient exploration of a MOF design space to optimize MOF properties, *Chem. Sci.* **15**, 18903 (2024).
- [32] K. Liu, Z. Chen, T. Islamoglu, S. J. Lee, H. Chen, T. Yildirim, O. K. Farha, and R. Q. Snurr, Exploring the chemical space of metal-organic frameworks with *rht* topology for high capacity hydrogen storage, *J. Phys. Chem. C* **128**, 7435 (2024).
- [33] S. Lee, B. Kim, H. Cho, H. Lee, S. Y. Lee, E. S. Cho, and J. Kim, Computational screening of trillions of metal-organic frameworks for high-performance methane storage, *ACS Appl. Mater. Interfaces* **13**, 23647 (2021).
- [34] S. Majumdar, S. M. Moosavi, K. M. Jablonka, D. Ongari, and B. Smit, Diversifying databases of metal organic frameworks for high-throughput computational screening, *ACS Appl. Mater. Interfaces* **13**, 61004 (2021).
- [35] R. B. Getman, Y.-S. Bae, C. E. Wilmer, and R. Q. Snurr, Review and analysis of molecular simulations of methane, hydrogen, and acetylene storage in metal-organic frameworks, *Chem. Rev.* **112**, 703 (2012).
- [36] Y. J. Colón, D. Fairen-Jimenez, C. E. Wilmer, and R. Q. Snurr, High-throughput screening of porous crystalline materials for hydrogen storage capacity near room temperature, *J. Phys. Chem. C* **118**, 5383 (2014).
- [37] A. W. Thornton, Cory M. Simon, Jihan Kim, Ohmin Kwon, Kathryn S. Deeg, Kristina Konstas, Steven J. Pas, Matthew R. Hill, David A. Winkler, Maciej Haranczyk, and Berend Smit, Materials genome in action: Identifying the performance limits of physical hydrogen storage, *Chem. Mater.* **29**, 2844 (2017).
- [38] G. Anderson, B. Schweitzer, R. Anderson, and D. A. Gómez-Gualdrón, Attainable volumetric targets for adsorption-based hydrogen storage in porous crystals: Molecular simulation and machine learning, *J. Phys. Chem. C* **123**, 120 (2019).
- [39] P. García-Holley, Benjamin Schweitzer, Timur Islamoglu, Yangyang Liu, Lu Lin, Stephanie Rodriguez, Mitchell H. Weston, Joseph T. Hupp, Diego A. Gómez-Gualdrón, Taner Yildirim, and Omar K. Farha, Benchmark study of hydrogen storage in metal-organic frameworks under temperature and pressure swing conditions, *ACS Energy Lett.* **3**, 748 (2018).
- [40] L. T. Glasby and P. Z. Moghadam, Hydrogen storage in MOFs: Machine learning for finding a needle in a haystack, *Patterns* **2**, 100305 (2021).
- [41] B. J. Bucior, N. S. Bobbitt, T. Islamoglu, S. Goswami, A. Gopalan, T. Yildirim, O. K. Farha, N. Bagheri, and R. Q. Snurr, Energy-based descriptors to rapidly predict hydrogen storage in metal-organic frameworks, *Mol. Syst. Des. Eng.* **4**, 162 (2019).
- [42] G. Borboudakis, T. Stergiannakos, M. Frysali, E. Klontzas, I. Tsamardinos, and G. E. Froudakis, Chemically intuited, large-scale screening of MOFs by machine learning techniques, *npj Comput. Mater.* **3**, 40 (2017).
- [43] Z. Wang, Y. Zhou, T. Zhou, and K. Sundmacher, Identification of optimal metal-organic frameworks by machine learning: Structure decomposition, feature integration, and predictive modeling, *Comput. Chem. Eng.* **160**, 107739 (2022).
- [44] A. A. Yamde, V. G. Lade, A. B. Bindwal, M. S. Tiwari, and R. P. Birmood, Machine learning approaches for the prediction of hydrogen uptake in metal-organic-frameworks: A comprehensive review, *Int. J. Hydrogen Energy* **98**, 1131 (2025).

- [45] H. Qiu, Y. Xia, C. Xiang, F. Xu, L. Sun, and Y. Zou, Prediction of hydrogen storage in metal-organic frameworks using CatBoost-based approach, *Int. J. Hydrogen Energy* **79**, 952 (2024).
- [46] S. Ghude and C. Chowdhury, Exploring hydrogen storage capacity in metal-organic frameworks: A Bayesian optimization approach, *Chem. - Eur. J.* **29**, e202301840 (2023).
- [47] C. Altintas, O. F. Altundal, S. Keskin, and R. Yildirim, Machine learning meets with metal organic frameworks for gas storage and separation, *J. Chem. Inf. Model.* **61**, 2131 (2021).
- [48] S. Meduri and J. Nandanavanam, Prediction of hydrogen uptake of metal organic frameworks using explainable machine learning, *Energy AI* **12**, 100230 (2023).
- [49] M. Seyyedattar, S. Zendehboudi, A. Ghamartale, and M. Afshar, Advancing hydrogen storage predictions in metal-organic frameworks: A comparative study of LightGBM and random forest models with data enhancement, *Int. J. Hydrogen Energy* **69**, 158 (2024).
- [50] C. Altintas and S. Keskin, On the shoulders of high-throughput computational screening and machine learning: Design and discovery of MOFs for H₂ storage and purification, *Mater. Today Energy* **38**, 101426 (2023).
- [51] Y. Cao, H. A. Dhahad, S. G. Zare, N. Farouk, A. E. Anqi, A. Issakhov, and A. Raise, Potential application of metal-organic frameworks (MOFs) for hydrogen storage: Simulation by artificial intelligent techniques, *Int. J. Hydrogen Energy* **46**, 36336 (2021).
- [52] X. Wang, H. M. Breunig, and P. Peng, Broad range material-to-system screening of metal-organic frameworks for hydrogen storage using machine learning, *Appl. Energy* **383**, 125346 (2025).
- [53] X. Lu, Z. Xie, X. Wu, M. Li, and W. Cai, Hydrogen storage metal-organic framework classification models based on crystal graph convolutional neural networks, *Chem. Eng. Sci.* **259**, 117813 (2022).
- [54] R. M. Giappa, E. Tylianakis, M. Di Gennaro, K. Gkagkas, and G. E. Froudakis, A combination of multi-scale calculations with machine learning for investigating hydrogen storage in metal organic frameworks, *Int. J. Hydrogen Energy* **46**, 27612 (2021).
- [55] K. Salehi, M. Rahmani, and S. Atashrouz, Machine learning assisted predictions for hydrogen storage in metal-organic frameworks, *Int. J. Hydrogen Energy* **48**, 33260 (2023).
- [56] R. J. Sadus, *Molecular Simulation of Fluids: Theory, Algorithms, and Object-Orientation* (Elsevier, Amsterdam, 1999).
- [57] T. L. Hill, *An Introduction to Statistical Thermodynamics* (Dover Publications, New York, USA, 1986).
- [58] Z. Li, K. Shi, D. Dubbeldam, M. Dewing, C. Knight, Á Vázquez-Mayagoitia, and R. Q. Snurr, Efficient implementation of Monte Carlo algorithms on graphical processing units for simulation of adsorption in porous materials, *J. Chem. Theory Comput.* **20**, 10649 (2024).
- [59] D. Dubbeldam, S. Calero, D. E. Ellis, and R. Q. Snurr, RASPA: Molecular simulation software for adsorption and diffusion in flexible nanoporous materials, *Mol. Simul.* **42**, 81 (2016).
- [60] A. K. Rappe, C. J. Casewit, K. S. Colwell, W. A. Goddard, and W. M. Skiff, UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations, *J. Am. Chem. Soc.* **114**, 10024 (1992).
- [61] S. L. Mayo, B. D. Olafson, and W. A. Goddard III, DREIDING: A generic force field for molecular simulations, *J. Phys. Chem.* **94**, 8897 (1990).
- [62] M. G. Martin and J. I. Siepmann, Transferable potentials for phase equilibria. 1. United-atom description of *n*-alkanes, *J. Phys. Chem. B* **102**, 2569 (1998).
- [63] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning* (Springer New York, New York, NY, 2009).
- [64] B. Panella, M. Hirscher, and S. Roth, Hydrogen adsorption in different carbon nanostructures, *Carbon N. Y.* **43**, 2209 (2005).
- [65] K. J. Gross, Recommended Best Practices for the Characterization of Storage Properties of Hydrogen Storage Materials (Dec. 31, 2015). <https://doi.org/10.23722/1718912>
- [66] R. Chahine and T. K. Bose, Low-pressure adsorption storage of hydrogen, *Int. J. Hydrogen Energy* **19**, 161 (1994).
- [67] M. G. Nijkamp, J. E. M. J. Raaymakers, A. J. Van Dillen, and K. P. De Jong, Hydrogen storage using physisorption-materials demands, *Appl. Phys. A* **72**, 619 (2001).
- [68] A. Züttel, P. Sudan, P. Mauron, and P. Wenger, Model for the hydrogen adsorption on carbon nanostructures, *Appl. Phys. A* **78**, 941 (2004).
- [69] Theo Jaffrelot Inizan, Sherry Yang, Aaron Kaplan, Yensu Lin, Jian Yin, Saber Mirzaei, Mona Abdelgaid, Ali H. Alawadhi, KwangHwan Cho, Zhiling Zheng, Ekin Dogus Cubuk, Christian Borgs, Jennifer T. Chayes, Kristin A. Persson, and Omar M. Yaghi, System of agentic AI for the discovery of metal-organic frameworks (2025). <https://doi.org/10.48550/arXiv.2504.14110>.
- [70] R. Ouyang, S. Curtarolo, E. Ahmetcik, M. Scheffler, and L. M. Ghiringhelli, SISSO: A compressed-sensing method for identifying the best low-dimensional descriptor in an immensity of offered candidates, *Phys. Rev. Mater.* **2**, 83802 (2018).
- [71] T. A. R. Purcell, M. Scheffler, and L. M. Ghiringhelli, Recent advances in the SISSO method and their implementation in the SISSO++ code, *J. Chem. Phys.* **159**, 114110 (2023).
- [72] A. Ahmed, Machine-learning-enhanced symbolic regression for methane storage prediction in covalent organic frameworks, *Phys. Rev. Mater.* **8**, 115408 (2024).
- [73] *Commercial Eureka Software Sold by DataRobot*: <https://www.datarobot.com/nutonian/>
- [74] M. Virgolin, T. Alderliesten, C. Witteveen, and P. A. N. Bosman, Improving model-based genetic programming for symbolic regression of small expressions, *Evol. Comput.* **29**, 211 (2021).
- [75] M. Cranmer, Interpretable machine learning for science with PySR and SymbolicRegression.jl (2023). <https://doi.org/10.48550/arXiv.2305.01582>.
- [76] A. Tonda, Review of PySR: high-performance symbolic regression in PYTHON and JULIA, Genet. Program. Evolvable Mach. **26**, 1 (2024).

- [77] gplearn: https://gplearn.readthedocs.io/en/stable/change_log.html.
- [78] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci. U. S. A.* **113**, 3932 (2016).
- [79] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. R. Stat. Soc. Ser. B* **58**, 267 (1996).
- [80] S. M. Udrescu and M. Tegmark, AI Feynman: A physics-inspired method for symbolic regression, *Sci. Adv.* **6**, eaay2631 (2020).
- [81] J. Liu, W. Li, L. Yu, M. Wu, W. Li, Y. Li, and M. Hao, Mathematical expression exploration with graph representation and generative graph neural network, *Neural Netw.* **187**, 107405 (2025).
- [82] L. Podina, B. Eastman, and M. Kohandel, in *Proceedings of the 40th International Conference on Machine Learning, PMLR* (2023), Vol. 202, p. 27948.
- [83] *LaSR: Symbolic Regression with a Learned Concept Library*. <https://doi.org/10.48550/arXiv.2409.09359>.
- [84] Y. Li, W. Li, L. Yu, M. Wu, J. Liu, W. Li, M. Hao, S. Wei, and Y. Deng, Generative pre-trained transformer for symbolic regression base in-context reinforcement learning. (2024). <https://arxiv.org/html/2404.06330v1>.
- [85] B. K. Petersen, M. L. Larma, T. N. Mundhenk, C. P. Santiago, S. K. Kim, and J. T. Kim, in *International Conference on Learning Representations (ICLR)*. (2021). <https://openreview.net/forum?id=m5Qsh0kBQG>.
- [86] C. J. Bartel, C. Sutton, B. R. Goldsmith, R. Ouyang, C. B. Musgrave, L. M. Ghiringhelli, and M. Scheffler, New tolerance factor to predict the stability of perovskite oxides and halides, *Sci. Adv.* **5**, 2375 (2019).
- [87] L. M. Ghiringhelli, J. Vybiral, E. Ahmetcik, R. Ouyang, S. V. Levchenko, C. Draxl, and M. Scheffler, Learning physical descriptors for materials science by compressed sensing, *New J. Phys.* **19**, 023017 (2017).
- [88] R. Ouyang, E. Ahmetcik, L. M. Ghiringhelli, and M. Scheffler, in *APS March Meeting 2017, Abstract Id. B1.004* (2017).
- [89] A. Ahmed, Input dataset used to develop the analytical equations in “physics-informed analytical models for interpretable and deployable hydrogen storage prediction in MOFs,” <https://zenodo.org/records/17108560>.
- [90] A. Ahmed and D. J. Siegel, *HyMARC Datahub*, <https://datahub.hymarc.org/dataset/computational-prediction-of-hydrogen-storage-capacities-in-mofs>.
- [91] A. W. van der Vaart, *Asymptotic Statistics* (Cambridge University Press, New York, USA, 1998).
- [92] P. Billingsley, *Probability and Measure*, Wiley Series in Probability and Statistics, (John Wiley & Sons, Inc., New York, USA, 1995) 3rd ed.
- [93] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/glnw-drn4> for detailed correlation analyses (Spearman and Pearson correlation matrices and pairwise scatter plots for all structural descriptors), real-world validation performance metrics on experimental MOF datasets, and comparative complexity analysis of symbolic regression methods including equation structure comparisons between SISSO and PySR approaches. No additional references are cited in the Supplemental Material.
- [94] T. F. Willems, C. H. Rycroft, M. Kazi, J. C. Meza, and M. Haranczyk, Algorithms and tools for high-throughput geometry-based analysis of crystalline porous materials, *Microporous Mesoporous Mater.* **149**, 134 (2012).