TRANSFERABLE BOLTZMANN GENERATOR FOR DIRECT AIR CAPTURE

Anonymous authors
Paper under double-blind review

ABSTRACT

Screening metal–organic frameworks (MOFs) for direct air capture requires efficient estimation of dilute-limit CO_2 adsorption. Conventional Widom insertion spends many energy evaluations on unfavorable configurations, limiting the scale of screening with machine-learned interatomic potentials (MLIPs). We introduce a transferable MOF-conditioned Boltzmann generator for estimating the Henry coefficient without per-material retraining. The model uses equivariant MOF features from a pretrained MLIP to construct a normalized proposal over CO_2 position, orientation, and internal deformation, enabling importance-weighted estimation with flexible molecules. Periodically wrapped Gaussian position mixtures and conditional orientation and deformation mixtures support direct sampling and density evaluation, while MOF features are computed once and reused across draws. Across unseen MOFs spanning weak to strong adsorption, our model maintains high sampling efficiency and closely matches reference configuration and energy distributions after reweighting. It estimates flexible-gas Henry coefficients at roughly 700-fold lower inference cost than rigid-gas Widom insertion while capturing systematic effects of molecular deformation that additional rigid-gas sampling cannot recover.

1 INTRODUCTION

Climate mitigation scenarios include gigatonne-scale CO_2 removal through direct air capture (DAC) and storage, but high costs remain a barrier to deployment (Realmonte et al., 2019; Young et al., 2023). A key challenge is capturing CO_2 at its low atmospheric concentration of about 400 parts per million, motivating the search for adsorbents that perform well under these dilute conditions (Bose et al., 2024). Metal–organic frameworks (MOFs) are promising DAC adsorbents due to their high porosity and large, tunable design space (Bose et al., 2024). However, their vast design space is both an opportunity and a challenge: it may contain strong DAC adsorbents, but finding them requires evaluating adsorption across many candidate structures. Since experimental measurements at this scale are impractical, high-throughput virtual screening with molecular simulation has been widely used (Wilmer et al., 2012; Boyd et al., 2019; Lim et al., 2025).

Virtual screening for adsorption commonly uses Widom insertion to estimate the Henry coefficient K_H, which measures gas uptake per unit pressure in the dilute limit (Widom, 1963; Wilmer et al., 2012). Widom insertion averages the Boltzmann weights of randomly sampled molecular configurations, corresponding to importance sampling with a uniform proposal over positions and orientations. Although convergence may require millions of energy evaluations per material (Vandenbrande et al., 2018), Widom insertion has been applied to large-scale screening (Wilmer et al., 2012; Boyd et al., 2019) due to (a) the insertions being readily parallelized and (b) classical force fields being inexpensive to evaluate.

However, this screening strategy becomes less practical when classical force fields are replaced by machine-learned interatomic potentials (MLIPs). MLIPs have attracted growing interest for their ability to approximate density functional theory calculations and more accurately describe CO_2 binding at sites such as open metal centers (Goeminne et al., 2023; Lim et al., 2025; Edwards et al., 2026). However, their higher cost per energy evaluation (Zhang et al., 2021) limits the number of materials that can be screened with uniform insertion within a fixed computational budget. Fixing

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

the internal geometry of CO_2 is a common way to reduce sampling cost (Wilmer et al., 2012; Boyd et al., 2019), but it excludes internal molecular deformation and can bias estimates of K_H.

In this work, we consider Boltzmann generators (BGs) to accelerate MLIP-based adsorption screening. BGs learn proposal distributions that concentrate samples on thermodynamically favorable configurations, and reweight the samples to estimate free-energy differences (Noé et al., 2019). Motivated by the success of BGs and related models for peptides (Tan et al., 2025b; Rehman et al., 2026), protein ensembles (Lewis et al., 2025), and crystal configurations (Schebek et al., 2026), we study transferable BGs for sampling CO_2 across diverse MOFs, aiming to accelerate high-throughput screening with transferability to unseen materials.

Contribution. We introduce a transferable MOF-conditioned molecular Boltzmann generator for amortized K_H estimation. The model predicts a normalized joint proposal over CO_2 position, orientation, and internal deformation, allowing complete molecular configurations to be sampled and reweighted. To condition this proposal on the MOF, a trainable equivariant encoder refines frozen MLIP features, and a shared pose score couples the position and orientation distributions.

Our key design choice is the direct parameterization of proposal density as a mixture of Gaussians, which allows rapid sampling of CO_2 configurations. Flow-based BGs are commonly used, but they require repeated network evaluations for sampling and likelihood computation, which limits their applicability to adsorption (Tan et al., 2025b; Klein & Noé, 2024). The low-dimensional configuration space of CO_2 permits this more direct parameterization: localized, periodically wrapped Gaussians for position, and conditional mixtures for orientation and internal deformation. A shared pose score couples position and orientation, weighting each position by how well CO_2 can be oriented there. Each mixture is sampled and its exact density evaluated in closed form, and MOF features are computed once per material and reused, so additional draws do not rerun the MOF encoder.

To evaluate our method, we conduct experiments on 6,724 MOFs spanning a wide range of adsorption strengths and compare with conventional Widom insertion with optional annealed importance sampling. Our model estimates K_H for rigid CO_2 with roughly half the mean error of the Widom insertion baseline at approximately $700\times$ lower inference cost. The same model accurately estimates K_H for flexible CO_2, whose reference values differ from the rigid-gas values by 23% on average, capturing effects that additional rigid-gas sampling cannot recover. It also maintains efficient CO_2 sampling on flexible MOF structures without retraining, providing a step toward adsorption calculations with flexible frameworks. We plan to release the MOF–CO_2 configuration dataset used to train the model to support further research on transferable molecular sampling.

2 Background and Related Work

2.1 Boltzmann distributions and free-energy estimation

Equilibrium configurations follow a Boltzmann distribution defined by the system's energy function. For a configuration $x \in \mathcal{X}$ with energy $U(x)$, the target Boltzmann density is

$$p(x) = \frac{\widetilde{p}(x)}{Z}, \qquad \widetilde{p}(x) = e^{-\beta U(x)}, \qquad Z = \int_{\mathcal{X}} \widetilde{p}(x)\,\mathrm{d}x, \tag{1}$$

where $\mathcal{X}$ is the configuration space, $\widetilde{p}(x)$ is the unnormalized density, and Z is its normalizing constant. The inverse thermal energy is $\beta = (k_B T)^{-1}$, with Boltzmann constant k_B and temperature T.

For two states on $\mathcal{X}$ at the same temperature, the configurational free-energy difference is

$$\Delta F_{0\to 1} = F_1 - F_0 = -\beta^{-1} \log \frac{Z_1}{Z_0}, \tag{2}$$

where Z_0 and Z_1 are the normalizing constants of the two states.

Given a normalized proposal distribution with a tractable density $q_\theta(x)$, the normalizing constant can be expressed as an importance-weighted expectation,

$$Z = \mathbb{E}_{x\sim q_\theta}[w_\theta(x)], \qquad w_\theta(x) = \frac{\widetilde{p}(x)}{q_\theta(x)}, \tag{3}$$

where $w_\theta(x)$ is the unnormalized importance weight.

2.2 HENRY COEFFICIENT FOR GAS ADSORPTION

The Henry coefficient relates gas adsorption in a material to gas partial pressure at infinite dilution,

$$K_H = \lim_{P \to 0} \frac{n(P)}{P}, \qquad n(P) \simeq K_H P \quad \text{in the dilute regime.} \tag{4}$$

Here, $n(P)$ is the amount of gas adsorbed per unit mass of material and P is the gas partial pressure. The linear approximation applies to ppm-level gas mixtures when their partial pressure lies in this regime.

Estimating K_H amounts to estimating a bound–unbound free-energy difference. Let $\mathcal{M} = (A, R, L)$ specify a material structure by its atom types, atomic coordinates, and lattice matrix, and let X denote the gas molecule's Cartesian coordinates. At fixed temperature and lattice, the Henry coefficient is

$$K_H = C_{\mathcal{M}} e^{-\beta \Delta F} = C_{\mathcal{M}} \frac{Z_{\text{mat+gas}}}{Z_{\text{mat}} Z_{\text{gas}}}, \tag{5}$$

where $\Delta F = F_{\text{mat+gas}} - F_{\text{mat}} - F_{\text{gas}}$. The normalizing constants correspond to the interacting material–gas system, the material without gas, and the isolated gas molecule, respectively. The unbound reference factorizes as $Z_{\text{mat}} Z_{\text{gas}}$. The known prefactor $C_{\mathcal{M}}$ depends on temperature, unit-cell volume, and material mass per cell, but not on R. Appendix A provides the derivation.

This ratio is an equilibrium average of the conditional gas normalizing constant over material configurations. Define $p_{\text{mat}}(R) = e^{-\beta U_{\text{mat}}(R)}/Z_{\text{mat}}$ and $\Delta U(R, X) = U_{\text{mat+gas}}(R, X) - U_{\text{mat}}(R)$, where the two energies correspond to the material with and without gas. Then

$$K_H = \frac{C_{\mathcal{M}}}{Z_{\text{gas}}} \mathbb{E}_{R \sim p_{\text{mat}}}[I(R)], \qquad I(R) = \int e^{-\beta \Delta U(R,X)} \, \mathrm{d}X. \tag{6}$$

Here $I(R)$ is the normalizing constant of the conditional gas Boltzmann distribution given the material configuration R. The average varies with R at fixed A and L. Dependence on A and L is implicit in $I(R)$, the energies, p_{mat}, and Z_{gas}.

Rigidity assumptions. Rigidity assumptions reduce the cost of adsorption calculations by restricting the configurational degrees of freedom. We use K_H^{FF}, $K_H^{\text{RF}}(\mathcal{M})$, and $K_H^{\text{RR}}(\mathcal{M})$, with the first letter indicating the material and the second the gas. F denotes flexible, and R denotes rigid. The unrestricted coefficient defined above corresponds to K_H^{FF}. The rigid-material coefficients depend on a single structure $\mathcal{M}$. For a rigid material, R is held fixed, replacing the material ensemble average by evaluation at that configuration. A rigid gas fixes its internal geometry, leaving only position r and orientation Ω. The resulting normalizing constant is $I^{\text{rigid}}(R)$, and the corresponding gas reference is $Z_{\text{gas}}^{\text{rigid}}$. Table 1 summarizes the three settings considered in this work. Combining both assumptions with uniform position and orientation sampling gives the widely used rigid-material, rigid-gas form of Widom insertion ($K_H^{\text{RR}}(\mathcal{M})$). Appendix A gives the derivation.

Table 1: **Henry coefficients under rigidity assumptions.** Superscripts denote material and gas rigidity, respectively.

Material	Gas	Notation	Expression
Flexible	Flexible	K_H^{FF}	$\frac{C_{\mathcal{M}}}{Z_{\text{gas}}} \mathbb{E}_{p_{\text{mat}}}[I(R)]$
Rigid	Flexible	$K_H^{\text{RF}}(\mathcal{M})$	$\frac{C_{\mathcal{M}}}{Z_{\text{gas}}} I(R)$
Rigid	Rigid	$K_H^{\text{RR}}(\mathcal{M})$	$\frac{C_{\mathcal{M}}}{Z_{\text{gas}}^{\text{rigid}}} I^{\text{rigid}}(R)$

Writing the rigid gas coordinates as $X(r, \Omega)$ gives

$$I^{\text{rigid}}(R) = \int\int e^{-\beta \Delta U(R, X(r,\Omega))} \, \mathrm{d}r \, \mathrm{d}\Omega, \tag{7}$$

where r ranges over one unit cell and $\mathrm{d}\Omega$ denotes the normalized uniform orientation measure.

2.3 RELATED WORK

Boltzmann generators. Boltzmann generators learn proposals for equilibrium sampling and free-energy estimation, with reweighting correcting differences from the target distribution (Noé et al.,

2019). Flow-based BGs evaluate proposal densities through invertible transformations and their Jacobians. Sequential BGs combine flow proposals with sequential Monte Carlo refinement (Tan et al., 2025a), while autoregressive BGs construct densities directly from normalized conditional distributions without invertible flows (Rehman et al., 2026). PROSE and autoregressive BGs demonstrate transfer across peptide systems (Tan et al., 2025b; Rehman et al., 2026). In materials, scalable BGs sample bulk crystal configurations and transfer to larger simulation cells (Schebek et al., 2026). Our task instead requires a flexible gas proposal conditioned on different porous materials.

Density prediction. Structure-conditioned density models predict scalar fields over three-dimensional space. Electron-density models, including DeepDFT, ChargE3Net, InfGCN, and GP-WNO, map atomic structures to electronic density fields through spatial readouts or learned basis representations (Jørgensen & Bhowmik, 2022; Koker et al., 2024; Cheng & Peng, 2023; Kim & Ahn, 2024). Adsorption models predict spatial gas distributions. DeepAPD represents guests such as methane and xenon as single spherical sites and predicts their position distributions (Burner et al., 2026). MADField predicts molecular-center number densities whose integrals give the amount adsorbed (Kim et al., 2026). SorbIIT predicts spatially resolved adsorption from interaction-energy grids (Sun & Siepmann, 2024). These adsorption outputs describe spatial densities rather than joint distributions over molecular position, orientation, and internal deformation. Our estimator requires a normalized probability density over complete gas configurations so that each generated configuration has a tractable importance weight.

3 METHOD

We amortize Henry coefficient estimation by learning a molecular proposal conditioned on the material structure. The same model generates flexible CO_2 configurations and evaluates their proposal densities for importance sampling across materials, without per-material retraining.

3.1 AMORTIZED HENRY COEFFICIENT ESTIMATION

We learn a shared proposal $q_\theta(X \mid \mathcal{M})$ to amortize gas sampling across materials. Given $\mathcal{M} = (A, R, L)$, the model generates complete gas configurations and evaluates their densities. The target is $e^{-\beta \Delta U(R,X)}/I(R)$, with energies supplied by a fixed MLIP.

Gas importance sampling. The proposal estimates $I(R)$ through importance sampling,

$$I(R) = \mathbb{E}_{X \sim q_\theta(\cdot|\mathcal{M})} \left[\frac{e^{-\beta \Delta U(R,X)}}{q_\theta(X \mid \mathcal{M})} \right]. \tag{8}$$

Here q_θ is normalized in Cartesian gas coordinates and covers the target support. Learning q_θ changes sampling efficiency, not the target energy function. The Henry coefficients follow as

$$K_H^{\mathrm{RF}}(\mathcal{M}) = \frac{C_{\mathcal{M}}}{Z_{\mathrm{gas}}} I(R), \qquad K_H^{\mathrm{FF}} = \mathbb{E}_{R \sim p_{\mathrm{mat}}}[K_H^{\mathrm{RF}}(\mathcal{M})]. \tag{9}$$

Material reweighting. Estimating K_H^{FF} further requires averaging over material configurations at fixed A and L. For samples from a material proposal $q_{\mathrm{mat}}(R)$, importance reweighting gives

$$K_H^{\mathrm{FF}} = \frac{\mathbb{E}_{R \sim q_{\mathrm{mat}}}[w_{\mathrm{mat}}(R) K_H^{\mathrm{RF}}(\mathcal{M})]}{\mathbb{E}_{R \sim q_{\mathrm{mat}}}[w_{\mathrm{mat}}(R)]}, \qquad w_{\mathrm{mat}}(R) = \frac{e^{-\beta U_{\mathrm{mat}}(R)}}{q_{\mathrm{mat}}(R)}. \tag{10}$$

where w_{mat} is the material importance weight and q_{mat} covers the material target support. The ratio eliminates Z_{mat}. Its sample average is self-normalized importance sampling, which is consistent but biased at finite sample sizes. The gas reference Z_{gas} is computed separately to within 0.001 in $\ln K_H$ (Appendix A).

3.2 SYSTEM PARAMETRIZATION AND SYMMETRY

We construct a tractable proposal that retains molecular flexibility and respects the periodic material geometry. A CO_2 configuration X is reconstructed from molecular coordinates

$$X = \Phi_{\mathcal{M}}(c, u, v), \qquad c \in \mathbb{T}^3, \quad u \in S^2, \quad v \in \mathbb{R}^4, \tag{11}$$

where c is the fractional arithmetic center, u is the directed oxygen–oxygen axis, and v contains two bond-stretch and two bending coordinates. The reconstruction map is specified in Appendix D.

Periodic proposal density. We factorize the proposal into normalized distributions for position, orientation, and deformation,

$$q_\theta(c, u, v \mid \mathcal{M}) = q_\theta(c \mid \mathcal{M})\, q_\theta(u \mid c, \mathcal{M})\, q_\theta(v \mid c, u, \mathcal{M}). \tag{12}$$

The left-hand side is a density in molecular coordinates. Wrapped-Gaussian position components handle cell boundaries. Periodic atom images in the probe readout and periodic spatial queries give the conditional heads consistent material features across those boundaries. The orientation and deformation factors use Watson (Watson, 1965) and Gaussian mixtures, respectively. All three factors support direct sampling and density evaluation, as detailed in Appendix B.

Oxygen-exchange symmetry. We enforce oxygen-label invariance on the full molecular density, not only on its orientation factor. Let $q_{\theta,\mathrm{ord}}$ be the Cartesian density corresponding to Equation 12 (Appendix D). For the oxygen exchange operator P, we use

$$q_{\theta,\mathrm{sym}}(X \mid \mathcal{M}) = \frac{q_{\theta,\mathrm{ord}}(X \mid \mathcal{M}) + q_{\theta,\mathrm{ord}}(\mathsf{P}X \mid \mathcal{M})}{2}. \tag{13}$$

This accounts for the simultaneous changes in axis and internal coordinates under oxygen exchange. Sampling applies a random oxygen swap with probability $1/2$. For importance weighting, $q_\theta(X \mid \mathcal{M})$ denotes $q_{\theta,\mathrm{sym}}(X \mid \mathcal{M})$. We use the same symmetrization for training. Without symmetrization the importance-sampling estimate would remain unbiased; symmetrization reduces its variance and raises sampling efficiency (Table 3, Appendix D).

3.3 Model Architecture and Training

The network computes the three factors in Equation 12 from the material structure. We extract frozen atomic features from the pretrained MLIP MACE-MH-1 (Batatia et al., 2022; 2025) with its OMAT-PBE head. For material atom i, $f_i^{(l)}$ denotes an order-l irreducible feature, comprising scalar ($l = 0$), vector ($l = 1$), or symmetric traceless tensor ($l = 2$) channels. A learnable equivariant encoder updates these features. At each voxel center c_j, a probe aggregates messages from nearby material atoms and their periodic images, producing scalar features $h_j^{(0)} \in \mathbb{R}^{128}$ and vector features $h_j^{(1)} \in \mathbb{R}^{128\times 3}$.

$$\{f_i^{(0)}, f_i^{(1)}, f_i^{(2)}\}_i \xrightarrow{\text{atom encoder}} \{\bar{f}_i^{(0)}, \bar{f}_i^{(1)}, \bar{f}_i^{(2)}\}_i \xrightarrow[\text{at } c_j]{\text{atom-to-probe}} (h_j^{(0)}, h_j^{(1)}). \tag{14}$$

Small heads applied to the probe features of Equation 14 give three fields on the grid: a position logit b_j, scalar carbon and oxygen site fields, and a directional oxygen field from the probe vectors. For a position c and axis a, the pose score $S_\theta(c, a)$ adds the carbon field at the molecular center $r(c)$ and the oxygen fields at $r(c) \pm \ell a$, where ℓ is a reference C–O distance used only for scoring; it is high when the carbon and both oxygens fall on favorable sites. For voxel centers c_j and candidate axes a_k, the same score sets both position and orientation weights,

$$\pi_j \propto e^{b_j} \sum_k e^{S_\theta(c_j, a_k)}, \qquad \alpha_k(c) \propto e^{S_\theta(c, a_k)}, \tag{15}$$

normalized over j and k, respectively. A voxel is thus weighted by how many orientations fit there, and a sampled position chooses its axis among those orientations. These weights multiply the wrapped-Gaussian components of $q_\theta(c \mid \mathcal{M})$ and the Watson components of $q_\theta(u \mid c, \mathcal{M})$, which lie near the candidate axes. A small head predicts the Watson concentrations and the deformation mixture $q_\theta(v \mid c, u, \mathcal{M})$ from features read at the sampled pose. All fields are computed once per material, so scoring a new draw requires only interpolation on the grid.

The atom encoder and probe readout are equivariant, so probe vectors rotate with the material while probe scalars remain invariant. Wrapped position components and periodic image handling make the spatial conditioning consistent across unit-cell boundaries. The pose score is invariant to axis reversal, and the full proposal enforces oxygen exchange through Equation 13.

We train on material–CO_2 configurations by minimizing the negative log-likelihood $\mathcal{L}(\theta) = -\mathbb{E}_{(\mathcal{M},X)\sim\mathcal{D}}[\log q_{\theta,\mathrm{sym}}(X \mid \mathcal{M})]$, with the MLIP frozen. Mixture and score details are in Appendices B and C.

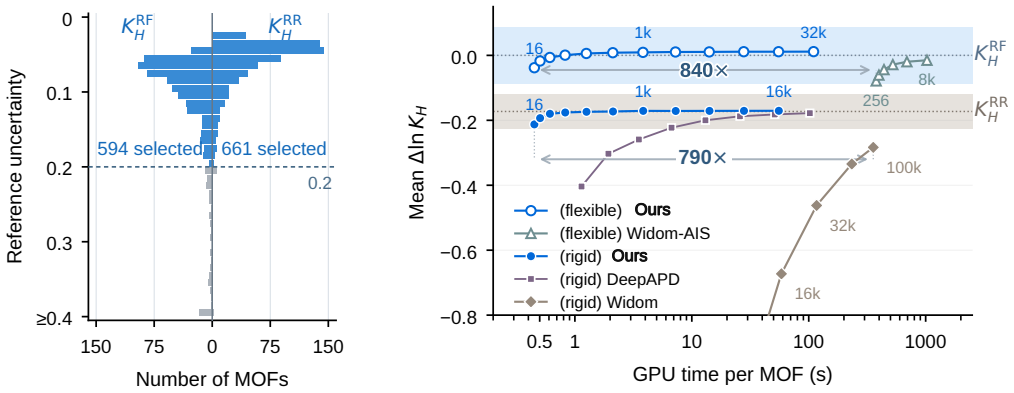


Figure 1: **Reference uncertainty and cost–accuracy comparison for K_H estimation.** (left) Uncertainty in reference K_H estimates computed from AIS-generated CO_2 samples for 672 unseen MOFs. For each MOF, we resample AIS particles with replacement 2,000 times and show the larger distance from the reference $\ln K_H$ estimate to the 2.5th or 97.5th percentile of the resulting estimates. Blue bars indicate MOFs below the 0.2 threshold: 594 for K_H^{RF} and 661 for K_H^{RR}. (right) Mean signed $\ln K_H$ difference from the K_H^{RF} reference versus GPU time for the common 588 MOFs passing the reference uncertainty threshold for both targets. Shaded bands average the per-MOF reference intervals for K_H^{RF} and K_H^{RR}. GPU times include one-time per-MOF encoding for Ours and DeepAPD. Widom-AIS denotes AIS initialized from 100,000 Widom insertions. Arrows show the GPU-time ratios at the marked budgets, where Ours is at least as close to the corresponding reference as the baseline.

4 Experiments

We evaluate whether a single material-conditioned Boltzmann generator estimates K_H across unseen MOFs. We first test K_H^{RR} and K_H^{RF} for given structures (Section 4.1), then test the proposal on flexible MOF structures as a prerequisite for K_H^{FF} (Section 4.2). We use 6,724 CO_2-accessible MOFs from LitMOF-DB, a curated version of the Cambridge Structural Database MOF subset that addresses known structural errors (Kim et al., 2025; White et al., 2025). We train on 6,052 MOFs with 512 flexible CO_2 configurations per MOF generated by sequential Monte Carlo (Del Moral et al., 2006) and evaluate on 672 unseen MOFs (Appendix E).

All energies are evaluated at 300 K with MACE-MH-1 (OMAT-PBE) (Batatia et al., 2022; 2025) and D3(BJ) (Grimme et al., 2010; 2011). We compare K_H^{RR} with rigid-CO_2 Widom insertion and DeepAPD, and K_H^{RF} with flexible-CO_2 annealed importance sampling (AIS) (Neal, 2001) initialized from 100,000 Widom insertions. We use separate AIS runs initialized from both Widom- and model-discovered sites for reference estimates because Widom-initialized AIS misses some CO_2 adsorption sites (Appendix F). We evaluate K_H accuracy with $|\Delta \ln K_H|$, sample quality with the normalized effective sample size (nESS) (Kong et al., 1994) and $\mathcal{E}$-$\mathcal{W}_2$, and cost with GPU time per MOF.

4.1 Henry Coefficient Estimation for Given MOF Structures

Reference precision and estimation cost. One model estimates both K_H^{RR} and K_H^{RF} without retraining and at substantially lower GPU cost than the baselines (Figure 1). Rigid CO_2 has no internal coordinates, so its proposal omits $q_\theta(v \mid c, u, \mathcal{M})$. We restrict the comparison to MOFs whose AIS reference has a 95% particle-resampling interval within ± 0.2 of its $\ln K_H$ estimate, leaving 661 MOFs for K_H^{RR} and 594 for K_H^{RF}, with 588 in both sets. On the common 588 MOFs, Ours approaches each reference while Widom and Widom-AIS remain farther from their respective targets at their largest evaluated budgets. The rigid- and flexible-gas estimates approach different reference values, so more rigid-gas samples cannot remove the difference. Appendix F gives the reference procedure and additional diagnostics.

Table 2: **Henry-coefficient results by quartile.** We report mean nESS (%), absolute $\ln K_H$ error, and $\mathcal{E}$-$\mathcal{W}_2$ normalized by the reference self-distance. MOFs are divided into quartiles by reference K_H, separately for the 661 K_H^{RR}-qualified and 594 K_H^{RF}-qualified structures. GPU-s denotes GPU seconds per MOF on a single NVIDIA H200 GPU, including one-time encoding for Ours and DeepAPD. Widom uses 100,000 insertions; AIS + Widom uses 8,192 particles initialized from those insertions; DeepAPD uses 4,096 draws; Ours uses 16 or 1,024 draws. † We benchmarked DeepAPD using its original atom-and-probe architecture, retrained to predict only the CO_2 center density $q(c \mid \mathcal{M})$. The CO_2 axis was sampled uniformly on S^2 for K_H^{RR}.

		Q1, weakest			Q2			Q3			Q4, strongest		
Model	GPU-s	nESS↑	$\lvert\Delta\rvert$ ↓	$\mathcal{E}$-$\mathcal{W}_2$ ↓	nESS↑	$\lvert\Delta\rvert$ ↓	$\mathcal{E}$-$\mathcal{W}_2$ ↓	nESS↑	$\lvert\Delta\rvert$ ↓	$\mathcal{E}$-$\mathcal{W}_2$ ↓	nESS↑	$\lvert\Delta\rvert$ ↓	$\mathcal{E}$-$\mathcal{W}_2$ ↓
K_H^{RR}													
Widom	356.3	0.614	0.915	8.744	0.371	0.191	3.087	0.143	0.182	3.107	0.041	0.375	5.053
DeepAPD†	14.2	3.32	0.243	2.813	5.03	0.107	1.909	3.53	0.109	2.250	1.69	0.218	2.883
Ours ($n = 16$)	0.51	**50.90**	0.271	5.001	**70.42**	0.195	5.907	**68.96**	0.165	6.085	**56.56**	0.256	6.013
Ours ($n = 1{,}024$)	3.9	43.77	**0.051**	**0.969**	65.18	**0.036**	**0.997**	63.18	**0.029**	**1.073**	49.32	**0.053**	**1.064**
K_H^{RF}													
AIS + Widom	1101.9	4.83	0.121	1.026	8.51	0.054	0.869	8.58	0.060	**0.907**	6.17	0.082	1.017
Ours ($n = 16$)	0.51	**57.45**	0.223	3.583	**68.57**	0.178	3.721	**65.12**	0.191	4.171	**54.36**	0.264	4.294
Ours ($n = 1{,}024$)	4.1	49.87	**0.051**	**0.803**	62.68	**0.047**	**0.789**	57.39	**0.042**	0.909	44.33	**0.057**	**0.895**

Accuracy across adsorption strengths. Even with 16 draws, Ours achieves 51%–70% mean nESS across the K_H quartiles at 0.51 GPU seconds per MOF (Table 2). At 1,024 draws, it has lower mean $|\Delta \ln K_H|$ and higher mean nESS than the corresponding baseline in every quartile, versus 4,096 draws for DeepAPD, 8,192 for Widom-AIS, and 100,000 for Widom.

Sample distributions. Across representative quartiles, unweighted Ours samples cover the main CO_2-center regions occupied by the AIS reference endpoints, although the proposal is broader in some cases (Figure 2).

Adsorption-site coverage. For 10 of 672 unseen MOFs, at least 5% of Ours' importance weight lies at CO_2 positions with no Widom-AIS sample within 1.5 Å (Figure 3). Only one MOF meets the same criterion with the methods reversed. These counts refer to MOFs with method-exclusive positions, not the number of adsorption sites. We compare directly against Widom-AIS without model-assisted reference particles and assess CO_2 positions only.

4.2 CO_2 Sampling on Flexible MOFs

The CO_2 proposal retains high sampling efficiency on flexible MOFs without retraining (Figure 4). Across 75,642 Hessian-Gaussian frames from 591 unseen MOFs, 128 CO_2 draws per frame give a median guest nESS of 53.2% at about 1.0 GPU second per frame. The estimated K_H^{RF} varies across frames, with a median absolute difference of 0.67 in $\ln K_H^{\mathrm{RF}}$ from the same MOF's rigid-structure estimate. For 39% of frames, this difference exceeds 1. These framewise estimates provide the inner calculation needed for K_H^{FF} (Equation 10), but the Hessian-Gaussian material proposal has a median effective sample size of only 1.5 over 128 frames. We therefore do not report a converged K_H^{FF} estimate.

4.3 Ablation Study

We retrain the model with one component removed, using the same data, schedule and seed, and report mean nESS for K_H^{RF} with 4,096 draws per MOF (Table 3). Every variant has lower mean nESS than the full model in every quartile. Replacing the pretrained MLIP features with a PaiNN encoder (Schütt et al., 2021) trained from scratch causes the largest decrease, 6.0–8.5 points, followed by removing the equivariant encoder (3.4–5.1) or the $l = 2$ features (3.4–4.6). Removing oxygen-exchange symmetrization at inference lowers nESS by 0.9–1.6 points. Mean $|\Delta \ln K_H|$ stays between 0.042 and 0.045 for all variants, because every proposal is exactly normalized; the components improve efficiency, not accuracy.

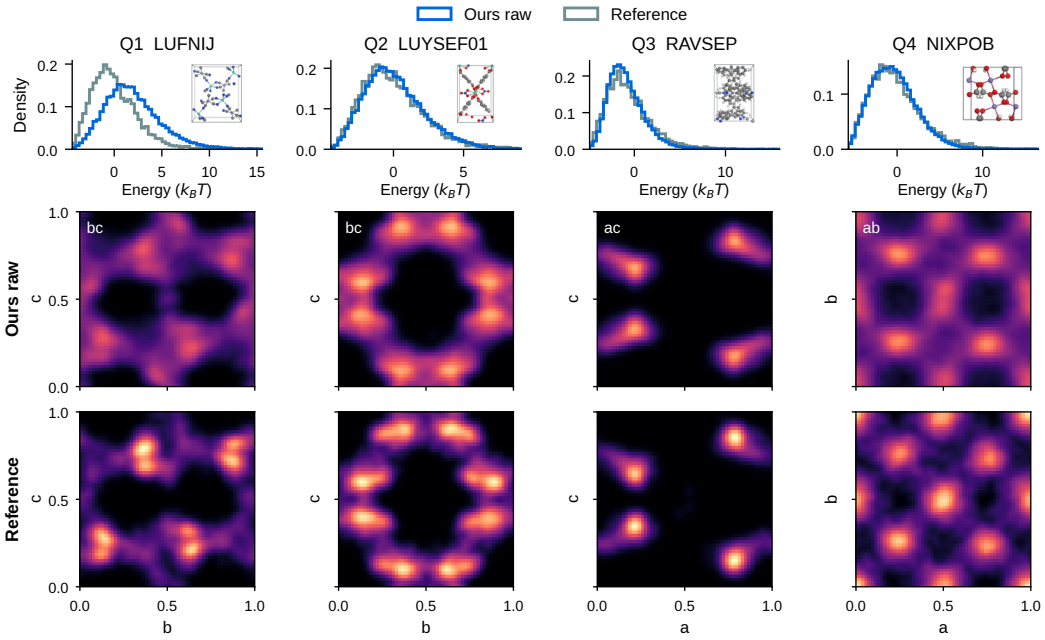


Figure 2: **CO_2 sample distributions across K_H^{RF} quartiles.** One MOF from each quartile of the 594 reference-qualified MOFs. Energy histograms and CO_2-center maps show unweighted samples from Ours and the frozen-schedule AIS reference endpoints. Energies are in k_BT relative to the mean reference energy for each MOF. Maps show fractional coordinates with a square-root color scale shared between the two rows for each MOF. Insets show the unit cells of the selected MOFs.

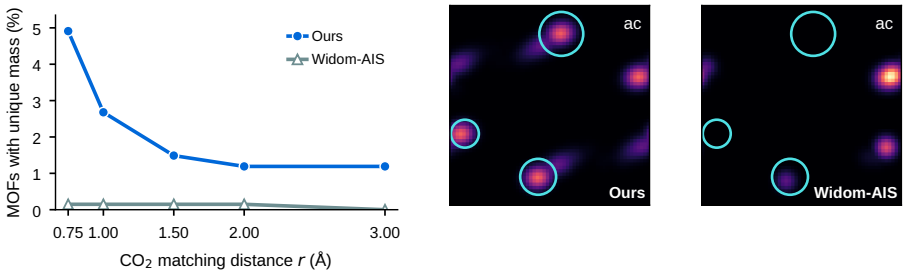
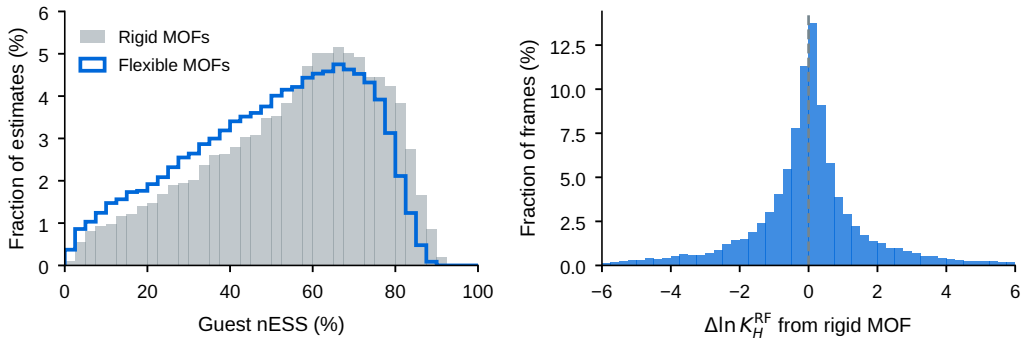


Figure 3: **Adsorption-site coverage.** (left) Fraction of 672 unseen MOFs for which at least 5% of one method's importance weight lies at CO_2 centers more than r from every sample of the other method. A higher Ours curve means that Ours finds high-weight positions missed by Widom-AIS. (right) Importance-weighted CO_2-center maps for the HUTRAQ MOF projected onto the ac plane. Both maps use the same color scale. Cyan circles indicate adsorption sites found by Ours but missed by Widom-AIS. HUTRAQ is one of the ten MOFs identified at $r = 1.5$ Å. Appendix G defines the metric and shows the other nine.

5 Conclusion

We introduced a material-conditioned sampler for adsorption free-energy estimation under an MLIP target. A shared pose score couples molecular position and orientation, while normalized conditional distributions cover all CO_2 degrees of freedom and support importance weighting. Training across MOFs enables transfer to unseen and flexible MOFs without retraining; converged flexible-MOF coefficients now hinge on a better material proposal. We evaluate sampling efficiency under the chosen potential, not agreement with experiment, and leave interacting adsorbates beyond infinite dilution to future work.



(a) Guest nESS (b) Framewise K_H^{RF} variation

Figure 4: **CO_2 sampling and K_H^{RF} variation on flexible MOFs.** Results for 75,642 Hessian-Gaussian frames from 591 unseen MOFs, with 128 flexible-CO_2 draws per frame. (a) Guest nESS on flexible MOF frames and on 128-draw subsets from the same MOFs held rigid at their relaxed structures. (b) Distribution of the framewise $\ln K_H^{\mathrm{RF}}$ difference from the corresponding rigid-MOF structure. The horizontal axis is limited to $[-6, 6]$, which contains 95.3% of frames. Host frames are not equilibrium-weighted, so panel (b) does not estimate K_H^{FF}.

Table 3: **Model ablations.** Mean nESS (%) for K_H^{RF} with 4,096 draws per MOF on the 594 reference-qualified unseen MOFs, in quartiles of reference K_H^{RF} (Q1 weakest). Each variant is retrained with the same data, schedule and seed. †Omitted symmetrization at inference.

Model	Q1	Q2	Q3	Q4
Full model	**49.58**	**61.97**	**56.50**	**43.29**
w/o equivariant encoder	45.99	56.96	51.40	39.89
PaiNN from scratch (no pretrained MLIP)	43.55	53.72	47.98	37.13
w/o pretrained MLIP $l = 2$ features	46.18	58.26	51.91	39.36
w/o oxygen-exchange symmetrization†	47.99	61.04	55.49	42.29

Reproducibility Statement

The proposal densities, network architecture, and training protocol are specified in Section 3 and Appendices B, C, and E. The data sources, the split into 6,052 training and 672 unseen MOFs, and the generation of training configurations are described in Section 4 and Appendix E; the reference procedure and gas normalization are given in Appendices F and A. The anonymized supplementary material contains the code, the trained model, inference and training scripts, the 672 unseen MOF structures, an installation script with pinned package versions, and tests that reproduce the reported proposal densities. The training configurations will be released upon publication.

AI Use Statement

We used generative AI tools to assist with writing, including drafting and editing the text of this paper. We have reviewed all AI-assisted text and take responsibility for the final content of this work.

REFERENCES

Ilyes Batatia, Dávid Péter Kovács, Gregor N. C. Simm, Christoph Ortner, and Gábor Csányi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. In *Advances in Neural Information Processing Systems*, volume 35, pp. 11423–11436, 2022.

Ilyes Batatia, Chen Lin, Joseph Hart, Elliott Kasoar, Alin M. Elena, Sam Walton Norwood, Thomas Wolf, and Gábor Csányi. Cross learning between electronic structure theories for unifying molecular, surface, and inorganic crystal foundation force fields. *arXiv preprint arXiv:2510.25380*, 2025.

Saptasree Bose, Debabrata Sengupta, Thomas M. Rayder, Xiaoliang Wang, Kent O. Kirlikovali, Ali K. Sekizkardes, Timur Islamoglu, and Omar K. Farha. Challenges and opportunities: Metal–organic frameworks for direct air capture. *Advanced Functional Materials*, 34(43):2307478, 2024.

Peter G Boyd, Arunraj Chidambaram, Enrique García-Díez, Christopher P Ireland, Thomas D Daff, Richard Bounds, Andrzej Gładysiak, Pascal Schouwink, Seyed Mohamad Moosavi, M Mercedes Maroto-Valer, et al. Data-driven design of metal–organic frameworks for wet flue gas co2 capture. *Nature*, 576(7786):253–256, 2019.

Jake Burner, Olivier Marchand, Rosa Cicciarella, Marco Gibaldi, and Tom K. Woo. Rapid prediction of single-site adsorbate probability distributions in metal–organic frameworks using graph neural networks. *Digital Discovery*, 5:2650–2668, 2026. doi: 10.1039/D6DD00027D.

Chaoran Cheng and Jian Peng. Equivariant neural operator learning with graphon convolution. In *Advances in Neural Information Processing Systems*, 2023.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006. doi: 10.1111/j.1467-9868.2006.00553.x.

Connor W. Edwards, Fengxu Yang, Konstantin Stracke, and Jack D. Evans. MLIP-MC: A framework for adsorption simulations using machine-learned interatomic potentials. *arXiv preprint arXiv:2602.13725*, 2026.

Ruben Goeminne, Louis Vanduyfhuys, Veronique Van Speybroeck, and Toon Verstraelen. DFT-quality adsorption simulations in metal-organic frameworks enabled by machine learning potentials. *Journal of Chemical Theory and Computation*, 19(18):6313–6325, 2023. doi: 10.1021/acs.jctc.3c00495.

Stefan Grimme, Jens Antony, Stephan Ehrlich, and Helge Krieg. A consistent and accurate *ab initio* parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu. *The Journal of Chemical Physics*, 132(15):154104, 2010. doi: 10.1063/1.3382344.

Stefan Grimme, Stephan Ehrlich, and Lars Goerigk. Effect of the damping function in dispersion corrected density functional theory. *Journal of Computational Chemistry*, 32(7):1456–1465, 2011. doi: 10.1002/jcc.21759.

Peter Bjørn Jørgensen and Arghya Bhowmik. Equivariant graph neural networks for fast electron density estimation of molecules, liquids, and solids. *npj Computational Materials*, 8:183, 2022. doi: 10.1038/s41524-022-00863-y.

Honghui Kim, Dohoon Kim, and Jihan Kim. LitMOF: An LLM multi-agent for literature-validated metal-organic frameworks database correction and expansion, 2025. URL https://arxiv.org/abs/2512.01693.

Seongsu Kim and Sungsoo Ahn. Gaussian plane-wave neural operator for electron density estimation. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 23805–23824, 2024.

Yoonho Kim, Seongsu Kim, Sungsoo Ahn, and Honghui Kim. MADField: Multi-fidelity amortized density field for adsorption in nanoporous materials. *arXiv preprint arXiv:2606.21284*, 2026.

Leon Klein and Frank Noé. Transferable Boltzmann generators. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Teddy Koker, Keegan Quigley, Eric Taw, Kevin Tibbetts, and Lin Li. Higher-order equivariant neural networks for charge density prediction in materials. *npj Computational Materials*, 10:161, 2024. doi: 10.1038/s41524-024-01343-1.

Augustine Kong, Jun S. Liu, and Wing Hung Wong. Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association*, 89(425):278–288, 1994. doi: 10.1080/01621459.1994.10476469.

Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew YK Foong, Victor García Satorras, Osama Abdin, Bastiaan S Veeling, Iryna Zaporozhets, et al. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science*, 389(6761): eadv9817, 2025.

Yunsung Lim, Hyunsoo Park, Aron Walsh, and Jihan Kim. Accelerating CO_2 direct air capture screening for metal-organic frameworks with a transferable machine learning force field. *Matter*, 8(7):102203, 2025. doi: 10.1016/j.matt.2025.102203.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Radford M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001. doi: 10.1023/A:1008923215028.

Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019. doi: 10.1126/science.aaw1147.

Giulia Realmonte, Laurent Drouet, Ajay Gambhir, James Glynn, Adam Hawkes, Alexandre C Köberle, and Massimo Tavoni. An inter-model assessment of the role of direct air capture in deep mitigation pathways. *Nature communications*, 10(1):3277, 2019.

Danyal Rehman, Charlie B. Tan, Yoshua Bengio, Avishek Joey Bose, and Alexander Tong. Autoregressive Boltzmann generators. In *International Conference on Machine Learning*, 2026.

Maximilian Schebek, Frank Noé, and Jutta Rogal. Scalable Boltzmann generators for equilibrium sampling of large-scale materials. *Nature Communications*, 17:5010, 2026. doi: 10.1038/s41467-026-73900-9.

Kristof T. Schütt, Oliver T. Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9377–9388, 2021.

Yangzesheng Sun and J. Ilja Siepmann. Understanding and predicting the spatially resolved adsorption properties of nanoporous materials. *Journal of Chemical Theory and Computation*, 20(12): 5259–5275, 2024. doi: 10.1021/acs.jctc.4c00149.

Charlie B. Tan, Avishek Joey Bose, Chen Lin, Leon Klein, Michael M. Bronstein, and Alexander Tong. Scalable equilibrium sampling with sequential Boltzmann generators. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 58467–58498, 2025a.

Charlie B. Tan, Majdi Hassan, Leon Klein, Saifuddin Syed, Dominique Beaini, Michael M. Bronstein, Alexander Tong, and Kirill Neklyudov. Amortized sampling with transferable normalizing flows. In *Advances in Neural Information Processing Systems*, 2025b. URL https://arxiv.org/abs/2508.18175.

Steven Vandenbrande, Michel Waroquier, Veronique Van Speybroeck, and Toon Verstraelen. Ab initio evaluation of henry coefficients using importance sampling. *Journal of Chemical Theory and Computation*, 14(12):6359–6369, 2018. doi: 10.1021/acs.jctc.8b00892.

G. S. Watson. Equatorial distributions on a sphere. *Biometrika*, 52(1–2):193–202, 1965. doi: 10.1093/biomet/52.1-2.193.

Andrew J. White, Marco Gibaldi, Jake Burner, R. Alex Mayo, and Tom K. Woo. High structural error rates in "computation-ready" MOF databases discovered by checking metal oxidation states. *Journal of the American Chemical Society*, 147(21):17579–17583, 2025. doi: 10.1021/jacs.5c04914.

Benjamin Widom. Some topics in the theory of fluids. *The Journal of Chemical Physics*, 39(11): 2808–2812, 1963. doi: 10.1063/1.1734110.

Christopher E. Wilmer, Michael Leaf, Chang Yeon Lee, Omar K. Farha, Brad G. Hauser, Joseph T. Hupp, and Randall Q. Snurr. Large-scale screening of hypothetical metal–organic frameworks. *Nature Chemistry*, 4(2):83–89, 2012.

John Young, Noah McQueen, Charithea Charalambous, Spyros Foteinis, Olivia Hawrot, Manuel Ojeda, Hélène Pilorgé, John Andresen, Peter Psarras, Phil Renforth, et al. The cost of direct air capture and storage can be reduced via strategic deployment but is unlikely to fall below stated cost targets. *One Earth*, 6(7):899–917, 2023.

Yaolong Zhang, Ce Hu, and Bin Jiang. Accelerating atomistic simulations with piecewise machine-learned ab initio potentials at a classical force field-like cost. *Physical Chemistry Chemical Physics*, 23:1815–1821, 2021. doi: 10.1039/D0CP05089J.

A ADSORPTION NORMALIZING CONSTANTS AND IMPORTANCE ESTIMATORS

States and reference measures. Use $\mathcal{M} = (\mathbf{A}, R, L)$ for the current material structure as in Section 2. The material is the host and the gas molecule is the guest. Fix the temperature, lattice, atom identities, and guest species. The host coordinates R have measure $\mathrm{d}R$. A guest configuration $X = (r_{\mathrm{C}}, r_{\mathrm{O}_1}, r_{\mathrm{O}_2})$ uses labeled Cartesian measure $\mathrm{d}X$. Its molecular center lies in one periodic cell and its intramolecular coordinates are unwrapped. Let $\mathcal{G}_{\mathcal{M}}$ be the intact-guest domain in that volume and $\mathcal{D}(R) \subseteq \mathcal{G}_{\mathcal{M}}$ the allowed domain in the host. Dependence on atom types and lattice is implicit in $\mathcal{D}(R)$, $I(R)$, the energies, and material distributions. The volume $V_{\mathcal{M}}$, mass $m_{\mathcal{M}}$, and prefactor $C_{\mathcal{M}}$ are constant as R varies at fixed atom types and lattice. Define

$$Z_{\mathrm{mat}} = \int e^{-\beta U_{\mathrm{mat}}(R)}\,\mathrm{d}R, \tag{16}$$

$$Z_{\mathrm{mat+gas}} = \int \int_{\mathcal{D}(R)} e^{-\beta U_{\mathrm{mat+gas}}(R,X)}\,\mathrm{d}X\,\mathrm{d}R, \tag{17}$$

$$Z_{\mathrm{gas}} = \int_{\mathcal{G}_{\mathcal{M}}} e^{-\beta U_{\mathrm{gas}}(X)}\,\mathrm{d}X. \tag{18}$$

All integrals are assumed finite. The gas has the same internal model and coordinate convention as the adsorbed guest, without host or periodic-image interactions. Its position is integrated over a reference volume equal to the host cell volume. This volume is an integration convention, not a periodic array of interacting gas molecules. Translational invariance gives $Z_{\mathrm{gas}} = V_{\mathcal{M}} z_{\mathrm{gas}}^{\mathrm{int}}$, where $z_{\mathrm{gas}}^{\mathrm{int}}$ includes the remaining orientation and internal integrals, so this reference-volume factor cancels the one in $C_{\mathcal{M}}$. Host and guest kinetic factors cancel in $Z_{\mathrm{mat+gas}}/(Z_{\mathrm{mat}} Z_{\mathrm{gas}})$. Using labeled oxygens in both $Z_{\mathrm{mat+gas}}$ and Z_{gas} also cancels the identical-particle symmetry factor. Inserting a factor of two on only one side would change K_H incorrectly.

Henry coefficient and units. We express the dilute-adsorption normalization underlying Widom insertion and Henry-coefficient estimation in our coordinate convention (Widom, 1963; Vandenbrande et al., 2018). Let η be the reservoir activity after absorbing guest kinetic factors, so a state with one guest has statistical weight $\eta Z_{\mathrm{mat+gas}}$. For an ideal gas in the reference cell, $\eta Z_{\mathrm{gas}} = \rho_{\mathrm{gas}} V_{\mathcal{M}}$ with $\rho_{\mathrm{gas}} = \beta P$. Let Ξ be the normalizing constant for the ensemble allowing a variable number of gas molecules. Expanding it in this activity gives

$$\Xi = Z_{\mathrm{mat}} + \eta Z_{\mathrm{mat+gas}} + O(\eta^2), \qquad \langle N \rangle = \eta \frac{\partial \log \Xi}{\partial \eta} = \eta \frac{Z_{\mathrm{mat+gas}}}{Z_{\mathrm{mat}}} + O(\eta^2).$$

States with two or more guests contribute only beyond first order. Substituting the gas-reference activity therefore gives

$$\langle N \rangle = \rho_{\mathrm{gas}} V_{\mathcal{M}} \frac{Z_{\mathrm{mat+gas}}}{Z_{\mathrm{mat}} Z_{\mathrm{gas}}} + o(P).$$

The amount of gas adsorbed per unit mass of material is $n(P) = \langle N \rangle / (N_A m_{\mathcal{M}})$. Here $V_{\mathcal{M}}$ and $m_{\mathcal{M}}$ are the host cell volume and mass, and N_A is Avogadro's constant. The flexible-host Henry coefficient is

$$K_H^{\mathrm{FF}} = \lim_{P \to 0} \frac{n(P)}{P} = C_{\mathcal{M}} \frac{Z_{\mathrm{mat+gas}}}{Z_{\mathrm{mat}} Z_{\mathrm{gas}}}, \qquad C_{\mathcal{M}} = \frac{\beta V_{\mathcal{M}}}{N_A m_{\mathcal{M}}}. \tag{19}$$

Using SI volume and mass gives $\mathrm{mol\,kg^{-1}\,Pa^{-1}}$. Multiplying its numerical value by 10^5 gives $\mathrm{mol\,kg^{-1}\,bar^{-1}}$. The prefactor contains temperature as well as unit conversion. Evaluating this SI prefactor requires k_B in $\mathrm{J\,K^{-1}}$. Boltzmann exponents instead use k_B in the same energy units as the potential, such as $\mathrm{eV\,K^{-1}}$ for the evaluator. The dimensionless quantity $\mathcal{B} = Z_{\mathrm{mat+gas}}/(Z_{\mathrm{mat}} Z_{\mathrm{gas}})$ satisfies $\mu_{\mathrm{ex}} = -\beta^{-1} \log \mathcal{B}$. This is an excess insertion free energy per molecule relative to an ideal guest in the same volume, not the logarithm of a dimensional Henry coefficient. Multiplication by N_A converts it to molar energy. A different adsorption standard state adds its specified reference term.

Energy offsets must also match across states. Under $U_{\mathrm{mat}} \mapsto U_{\mathrm{mat}} + a$, $U_{\mathrm{gas}} \mapsto U_{\mathrm{gas}} + b$, and $U_{\mathrm{mat+gas}} \mapsto U_{\mathrm{mat+gas}} + a + b$, the ratio is unchanged. Subtracting an isolated-guest energy from the loaded-state energy without making the corresponding change to Z_{gas} is not invariant.

Importance sampling. The identity in Equation 3 holds for an integrable target when $q_\theta > 0$ wherever $\widetilde{p} > 0$ and both densities use the same coordinate measure. Finite variance of the importance weight additionally requires $\int_{\mathcal{X}} \widetilde{p}(x)^2/q_\theta(x)\,\mathrm{d}x < \infty$. These statements concern exact density evaluation, with numerical approximations considered separately.

Joint normalizing-constant estimation. An alternative is to estimate Z_{mat} and $Z_{\mathrm{mat+gas}}$ separately using proposals $q_{\mathrm{mat}}(R)$ and $q_{\mathrm{mat+gas}}(R, X)$, then combine their estimates in the Henry ratio. For the loaded material, the weight is $w_{\mathrm{mat+gas}}(R, X) = e^{-\beta U_{\mathrm{mat+gas}}(R,X)}/q_{\mathrm{mat+gas}}(R, X)$ on the allowed domain. This estimator requires proposal coverage over the full material–gas configuration space. Agreement of generated configurations with typical energies alone does not ensure reliable normalizing-constant estimates.

Exploratory experiments with an earlier flow-based joint sampler on YOJMOC showed near-single-sample effective sample sizes when estimating $Z_{\mathrm{mat+gas}}$, including after increasing the number of divergence probes used to evaluate the density. These observations motivate conditional insertion in the present work. They concern an earlier model, not the current conditional architecture, and do not establish that every joint proposal has infinite variance. The mathematical issue is concentration of the importance weights, not the magnitude of the absolute energy zero. Conditional insertion removes the need to estimate Z_{mat} and $Z_{\mathrm{mat+gas}}$ separately, although its gas weights and any material reweighting still require adequate proposal coverage.

Conditional insertion and fixed hosts. Set $\Delta U(R, X) = U_{\mathrm{mat+gas}}(R, X) - U_{\mathrm{mat}}(R)$ and $p_{\mathrm{mat}}(R) = e^{-\beta U_{\mathrm{mat}}(R)}/Z_{\mathrm{mat}}$, with material dependence implicit. Then

$$I(R) = \int_{\mathcal{D}(R)} e^{-\beta \Delta U(R,X)}\,\mathrm{d}X, \tag{20}$$

$$\frac{Z_{\mathrm{mat+gas}}}{Z_{\mathrm{mat}}} = \mathbb{E}_{R\sim p_{\mathrm{mat}}}[I(R)], \tag{21}$$

$$K_H^{\mathrm{FF}} = \frac{C_{\mathcal{M}}}{Z_{\mathrm{gas}}}\mathbb{E}_{R\sim p_{\mathrm{mat}}}[I(R)]. \tag{22}$$

For any normalized conditional proposal $q(X \mid \mathcal{M})$ covering this domain,

$$I(R) = \mathbb{E}_{X\sim q(\cdot|\mathcal{M})}\left[\frac{\mathbf{1}_{\mathcal{D}(R)}(X)e^{-\beta\Delta U(R,X)}}{q(X \mid \mathcal{M})}\right]. \tag{23}$$

The averaging host distribution is that of the empty framework, not the loaded framework. The average is over I, not $\log I$. Holding R fixed defines a different target rather than an exact replacement of the flexible-host ensemble. The corresponding normalizing constant is

$$I(R) = \int_{\mathcal{D}(R)} e^{-\beta\Delta U(R,X)}\,\mathrm{d}X, \qquad K_H^{\mathrm{RF}}(\mathcal{M}) = \frac{C_{\mathcal{M}}}{Z_{\mathrm{gas}}}I(R). \tag{24}$$

Here $\mathcal{D}(R)$ is the allowed guest domain. Fixing the host leaves guest translation, orientation, and internal deformation free. A rigid-guest assumption separately constrains the internal geometry in both the adsorbed and gas-reference integrals. The rigid-material, flexible-gas excess insertion free energy is $\mu_{\mathrm{ex}}^{\mathrm{RF}} = -\beta^{-1}\log[I(R)/Z_{\mathrm{gas}}]$.

Define the conditional unnormalized gas density $\widetilde{p}(X \mid \mathcal{M}) = \mathbf{1}_{\mathcal{D}(R)}(X)e^{-\beta\Delta U(R,X)}$. Its normalizing constant is $I(R)$. For fixed material coordinates R, importance sampling gives

$$I(R) = \mathbb{E}_{X\sim q_\theta(\cdot|\mathcal{M})}\left[\frac{\widetilde{p}(X \mid \mathcal{M})}{q_\theta(X \mid \mathcal{M})}\right], \qquad K_H^{\mathrm{RF}}(\mathcal{M}) = \frac{C_{\mathcal{M}}}{Z_{\mathrm{gas}}}I(R). \tag{25}$$

Thus raw insertion weights estimate $I(R)$, whereas gas-normalized weights estimate $I(R)/Z_{\mathrm{gas}}$. Their normalized ESS is identical. Absolute Henry coefficients additionally depend on the accuracy of the gas reference normalization.

Rigid-gas measure and conventional Widom insertion. In this paragraph, the superscript or subscript rigid explicitly distinguishes the rigid-gas model from the flexible-gas model. For a fixed internal geometry, write $X = X(r, \Omega)$ and use measure $\mathrm{d}r\,\mathrm{d}\Omega$, where $\mathrm{d}\Omega$ is the normalized uniform orientation measure with $\int \mathrm{d}\Omega = 1$. For linear CO_2, orientation is an axis on S^2, with the same oxygen-label convention in both reference states. The isolated gas energy E^0_{gas} is independent of translation and orientation, so

$$Z^{\mathrm{rigid}}_{\mathrm{gas}} = V_{\mathcal{M}} e^{-\beta E^0_{\mathrm{gas}}}, \qquad q_{\mathrm{rigid}}(r, \Omega \mid \mathcal{M}) = V_{\mathcal{M}}^{-1}.$$

With both the material and gas rigid, the coefficient is

$$K^{\mathrm{RR}}_H(\mathcal{M}) = \frac{C_{\mathcal{M}}}{Z^{\mathrm{rigid}}_{\mathrm{gas}}} I^{\mathrm{rigid}}(R), \tag{26}$$

where I^{rigid} is the reduced-coordinate normalizing constant in Equation 7. Inserting the uniform proposal gives

$$K^{\mathrm{RR}}_H(\mathcal{M}) = C_{\mathcal{M}}\, \mathbb{E}_{r,\Omega\sim\mathrm{Unif}} \left[e^{-\beta[\Delta U(R, X(r,\Omega)) - E^0_{\mathrm{gas}}]} \right].$$

Thus the conventional expression averages the Boltzmann factor of the material–gas interaction energy. The subtraction of E^0_{gas} comes from dividing by the rigid-gas reference and does not redefine ΔU in the main text. Rigid-framework and rigid-molecule models are widely used in adsorption calculations, including the Henry-coefficient formulation of Vandenbrande et al. (2018). The more general Widom identity does not require either rigidity assumption.

Widom insertion and host proposals. Let $p_{\mathrm{gas}}(X) = e^{-\beta U_{\mathrm{gas}}(X)}/Z_{\mathrm{gas}}$ on $\mathcal{G}_{\mathcal{M}}$. Choosing $q(X \mid \mathcal{M}) = p_{\mathrm{gas}}(X)$ gives

$$\frac{I(R)}{Z_{\mathrm{gas}}} = \mathbb{E}_{X\sim p_{\mathrm{gas}}} \left[\mathbf{1}_{\mathcal{D}(R)}(X) e^{-\beta[U_{\mathrm{mat+gas}}(R,X) - U_{\mathrm{mat}}(R) - U_{\mathrm{gas}}(X)]} \right].$$

Uniform translation and rotation follow from the isolated gas reference. Its internal distribution is not uniform for a flexible molecule. A rigid-guest calculation changes both the adsorbed and gas configuration spaces. Host rigidity is a separate assumption, not a property of the insertion identity.

If material configurations are sampled from q_{mat} instead of p_{mat}, define $w_{\mathrm{mat}}(R) = e^{-\beta U_{\mathrm{mat}}(R)}/q_{\mathrm{mat}}(R)$. Provided the proposal covers the material target support, the coefficient is

$$K^{\mathrm{FF}}_H = \frac{\mathbb{E}_{R\sim q_{\mathrm{mat}}}[w_{\mathrm{mat}}(R) K^{\mathrm{RF}}_H(\mathcal{M})]}{\mathbb{E}_{R\sim q_{\mathrm{mat}}}[w_{\mathrm{mat}}(R)]}.$$

A Gaussian derived from a Hessian can serve as q_{mat}. Using it without host reweighting instead changes the target to a harmonic approximation. No flexible-host performance claim is made here.

Numerical conventions. The gas normalization must specify a matching intact-molecule reference domain and the same energy zero, temperature, coordinate measure, and volume convention.

For the evaluator's internal coordinate measure defined in Appendix D, the configured gas reference is $\log Z^{\mathrm{eval}}_{\mathrm{gas}} = \log V^{(\text{Å}^3)}_{\mathcal{M}} + 875.656344573973$ at 300 K. This value is a numerical reference input, not an analytic normalizer. Importance sampling of the isolated molecule with 1.6×10^7 samples converges to within 0.001 of this value, so the gas reference contributes negligible error. Under labeled Cartesian measure, $Z_{\mathrm{gas}} = 27 Z^{\mathrm{eval}}_{\mathrm{gas}}$. The corresponding factor in the proposal density cancels this conversion in every gas-normalized weight.

High sample ESS alone does not establish coverage or accuracy.

B NORMALIZED PROPOSAL DISTRIBUTIONS

All factors condition on the current material structure $\mathcal{M} = (\mathbf{A}, R, L)$, as in Section 2. Dependence on $\mathcal{M}$ is implicit in mixture parameters unless shown.

Position. On the fractional unit torus, the position density is

$$q_\theta(c \mid \mathcal{M}) = \sum_j \pi_j \mathcal{N}_{\mathrm{wrap}}(c; c_j + \delta_j, \Sigma_j), \tag{27}$$

where π_j is defined in Equation 38, δ_j is a learned offset within the voxel, and Σ_j is a positive diagonal covariance whose standard deviations are at least 0.08 Å and at most the smaller of the voxel width and 0.08 in fractional units. The wrapped density sums a Gaussian over integer translations:

$$\mathcal{N}_{\mathrm{wrap}}(c; m, \Sigma) = \sum_{n \in \mathbb{Z}^3} \mathcal{N}(c + n; m, \Sigma).$$

It integrates to one over $[0, 1)^3$. Sampling selects a component, draws its Gaussian, and wraps the result modulo one.

Orientation. The orientation factor is even in the directed axis u. Full oxygen exchange also transforms the internal coordinates. For density evaluation we use the sphere with normalized surface measure $\mathrm{d}\nu(u)$, with $\int_{S^2} \mathrm{d}\nu(u) = 1$. In this convention,

$$\mathcal{W}(u; \mu, \kappa) = \frac{\exp[\kappa(\mu^\top u)^2]}{{}_1F_1(1/2; 3/2; \kappa)}, \tag{28}$$

$$q_\theta(u \mid c, \mathcal{M}) = \sum_k \alpha_k(c) \mathcal{W}(u; \mu_k(c), \kappa_k(c)). \tag{29}$$

Here ${}_1F_1$ is the confluent hypergeometric function, μ_k is a unit axis within 4.5° of anchor a_k, and $\kappa_k \in [1, 80]$ controls concentration. Each component is even in u and normalized under ν. All mixture parameters depend on the position and material, not on the orientation being evaluated. The angular coordinates $u_z = 2\zeta - 1$, $\varphi = 2\pi\tau$ satisfy $\mathrm{d}\nu = \mathrm{d}\zeta\, \mathrm{d}\tau$ on $[0, 1] \times [0, 1)$.

Internal coordinates. For the standardized internal coordinates $v \in \mathbb{R}^4$ defined in Appendix D, we use

$$q_\theta(v \mid c, u, \mathcal{M}) = \sum_{m=1}^{2} g_m(c, u, \mathcal{M}) \mathcal{N}(v; \eta_m(c, u, \mathcal{M}), \operatorname{diag} s_m^2(c, u, \mathcal{M})). \tag{30}$$

The weights sum to one and the scales are positive. The parameter head uses local material features and an orientation-dependent summary weighted by the Watson-component responsibilities,

$$r_k(u) = \frac{\alpha_k(c) \mathcal{W}(u; \mu_k, \kappa_k)}{\sum_l \alpha_l(c) \mathcal{W}(u; \mu_l, \kappa_l)}.$$

This retains dependence between orientation and deformation without an additional normalization integral. Sequential integration of the three normalized factors proves normalization of the joint molecular-coordinate density.

Discretization and numerical evaluation. The reported configuration uses a 0.25 Å grid and 256 candidate axes. These discretizations determine mixture parameters, not the support of the continuous samples. Samples are drawn in single precision. The log density $\log q_\theta$ of every sample and the exactness checks are evaluated in double precision.

C MATERIAL REPRESENTATION AND POSE-SCORE CONSTRUCTION

All material features are evaluated from the current structure $\mathcal{M} = (\mathbf{A}, R, L)$. Dependence on $\mathcal{M}$ is implicit in the scores and mixture parameters below.

Equivariant refinement and spatial readout. The frozen material representation contains scalar, vector, and symmetric traceless tensor features $f_j^{(l)}$ for $l = 0, 1, 2$. Refinement and readout use messages of the form

$$m_i^{(l_o)} = \sum_{j \in \mathcal{N}(i)} \sum_{(l_i, l_e) \in \mathcal{P}(l_o)} w_{l_i l_e l_o}(r_{ij}) [f_j^{(l_i)} \otimes Y^{(l_e)}(\widehat{r}_{ij})]^{(l_o)}. \tag{31}$$

Here $\mathcal{N}(i)$ contains periodic neighbors of receiver i, r_{ij} is their separation, and $\widehat{r}_{ij}$ is its unit direction. The set $\mathcal{P}(l_o)$ contains the allowed input and edge angular orders coupling to output order l_o. The functions $w_{l_i l_e l_o}$ are learned radial filters. $Y^{(l_e)}$ denotes spherical harmonics and the bracket selects the output irreducible representation. In refinement, receivers are material atoms and a residual update retains features through order two. In spatial readout, receivers are grid points, and each grid point reads the 96 nearest atoms, including periodic images, within the cutoff. Probe updates couple scalar and vector states, and both are retained for the score heads.

The refinement uses learned radial filters. The readout uses learned coefficients over a fixed radial dictionary containing sinusoidal functions and Gaussian distance functions, including narrow functions near contact distances. The model uses one refinement layer, a 5 Å neighbor cutoff, and 128-dimensional invariant grid features. These features are computed at the full grid resolution.

Atom-type score fields. For voxel j, let $(h_j^{(0)}, h_j^{(1)})$ be the probe scalar and vector outputs from Equation 14. The position and site heads give

$$(b_j, \delta_j, \Sigma_j) = H_{\mathrm{pos}}(h_j^{(0)}, L), \tag{32}$$

$$(A_{\rho j}, B_{\rho j}, \gamma_{\rho j})_{\rho=1}^{D_s} = H_{\mathrm{site}}(h_j^{(0)}), \tag{33}$$

$$P_{\rho j} = H_{\mathrm{vec},\rho}(h_j^{(1)}, \gamma_{\rho j}). \tag{34}$$

Here $D_s = 6$. The position head yields a base logit, a within-voxel offset, and a positive diagonal covariance. The site head yields carbon and oxygen scalar fields and a gate. Channel projections of the probe vectors followed by scalar gating give the oxygen vector field.

Let $r(c)$ denote the Cartesian molecular center. We query carbon scalar fields A_ρ at $r(c)$ and oxygen scalar and vector fields B_ρ, P_ρ at $o_\pm = r(c) \pm \ell a$, with $\ell = 1.16$ Å. Defining $d_\pm = (r(c) - o_\pm)/\ell$, the score is

$$s_\rho(c, a) = A_\rho(r(c)) + \sum_{\sigma \in \{-,+\}} [B_\rho(o_\sigma) + P_\rho(o_\sigma)^\top d_\sigma], \tag{35}$$

$$S_\theta(c, a) = \log \left(\frac{1}{D_s} \sum_{\rho=1}^{D_s} e^{s_\rho(c,a)} \right), \tag{36}$$

with $D_s = 6$ learned channels. Field heads bound scalar outputs and vector magnitudes for numerical stability. The cached site-score grids use periodic trilinear interpolation. The pose-feature grids below use the same periodic interpolation. Swapping a for $-a$ exchanges the two oxygen summands and leaves the score unchanged. The reference bond length defines the scoring template. It does not constrain the generated internal coordinates.

Position and orientation weights. For voxel centers $\{c_j\}_{j=1}^J$ and candidate axes $\{a_k\}_{k=1}^K$, the shared score gives

$$\Lambda_j = \log \left(\frac{1}{K} \sum_{k=1}^{K} e^{S_\theta(c_j, a_k)} \right), \tag{37}$$

$$\pi_j = \mathrm{softmax}_j(b_j + \Lambda_j), \qquad \alpha_k(c) = \mathrm{softmax}_k S_\theta(c, a_k). \tag{38}$$

The orientation average Λ_j adjusts the position logit, and the same score sets the conditional orientation weights at a sampled position. This aggregation is learned and is not an exact target marginal. Material features and site fields are computed once per material and reused across gas samples. Field values at candidate poses use periodic interpolation from the cached grids.

Pose features. The cached score grids carry 32 additional carbon and oxygen channels, G_C and G_O, computed by two further heads from the same probe scalar features. For a position c and candidate axis a_k, the pose feature $[G_C(r(c)),\, G_O(o_+) + G_O(o_-)]$ is read with the same look-ups as $S_\theta(c, a_k)$. The sum over the two oxygen sites makes it invariant to axis reversal. A small network maps this feature to the Watson concentration $\kappa_k(c)$ and a bounded offset of $\mu_k(c)$ from the anchor a_k. The internal-coordinate mixture is conditioned on the probe scalar features at $r(c)$ and on the pose features pooled with the responsibilities $r_k(u)$. A channel projection of the probe vectors at $r(c)$, expressed in the tangent frame of u, supplies the bending direction. All inputs are look-ups in grids built once per material.

D COORDINATE MEASURES AND OXYGEN-EXCHANGE SYMMETRY

Molecular coordinates. For labeled atomic positions X, define the arithmetic center $r = (r_{\mathrm{C}} + r_{\mathrm{O}_1} + r_{\mathrm{O}_2})/3$, bonds $b_i = r_{\mathrm{O}_i} - r_{\mathrm{C}}$, lengths $\ell_i = \|b_i\|$, and directed axis $u = (b_1 - b_2)/\|b_1 - b_2\|$. The fractional center is $c \in \mathbb{T}^3$, with Cartesian map $r(c)$ given by the host lattice and a fixed origin. The implementation also uses a fixed orthogonal coordinate rotation, which does not change volume. Define dimensionless log lengths $t_i = \log(\ell_i/\ell_*)$, with reference length $\ell_* = 1$ Å, and

$$s_+ = (t_1 + t_2)/\sqrt{2}, \qquad s_- = (t_1 - t_2)/\sqrt{2}.$$

Let γ be the C–O bond angle and $e_\vartheta(u), e_\varphi(u)$ an orthonormal tangent frame. The bend disk vector $d \in \mathbb{R}^2$ has $\|d\|^2 = (1 + \cos\gamma)/2 < 1$. Its direction represents the projection of $b_1 + b_2$ onto that tangent frame. A linear molecule has $d = 0$ without a bending-plane azimuth. Set

$$y = \frac{d}{\sqrt{1 - \|d\|^2}}, \qquad a = (s_+, s_-, y_1, y_2), \qquad v = D_v^{-1}(a - a_0),$$

where a_0 is a fixed location vector and D_v is a positive diagonal scale matrix obtained from the training coordinates. Thus $v \in \mathbb{R}^4$ is standardized after the nonlinear disk transformation, not a vector of raw Cartesian displacements.

These coordinates reconstruct the bonds. Write $\ell_{1,2} = \ell_* \exp[(s_+ \pm s_-)/\sqrt{2}]$, $d = y/\sqrt{1 + \|y\|^2}$, $\chi = 2\|d\|^2 - 1$, and $L = (\ell_1^2 + \ell_2^2 - 2\ell_1\ell_2\chi)^{1/2}$. Then

$$b_1 - b_2 = Lu, \tag{39}$$

$$b_1 + b_2 = \frac{\ell_1^2 - \ell_2^2}{L}u + \frac{4\ell_1\ell_2\sqrt{1 - \|d\|^2}}{L}(d_1 e_\vartheta + d_2 e_\varphi). \tag{40}$$

Recover $r_{\mathrm{C}} = r - (b_1 + b_2)/3$ and $r_{\mathrm{O}_i} = r_{\mathrm{C}} + b_i$. This defines $X = \Phi_{\mathcal{M}}(z)$ for $z = (c, u, v)$ on regular coordinate patches. Spherical poles require a frame convention and form a set of zero surface measure.

Molecular-coordinate and Cartesian densities. Use normalized spherical surface measure $\mathrm{d}\nu(u)$, with $\int_{S^2} \mathrm{d}\nu(u) = 1$, and $\mathrm{d}\lambda_z = \mathrm{d}c\,\mathrm{d}\nu(u)\,\mathrm{d}v$. The Cartesian volume factor is defined by $\mathrm{d}X = J_{\mathcal{M}}(z)\,\mathrm{d}\lambda_z$. For these coordinates,

$$J_{\mathcal{M}}(z) = 16\pi V_{\mathcal{M}} \ell_1^3 \ell_2^3 \det(D_v)(1 + \|y\|^2)^{-2}. \tag{41}$$

The factors follow from $\mathrm{d}r = V_{\mathcal{M}}\mathrm{d}c$, $\mathrm{d}b_1\,\mathrm{d}b_2 = 16\pi\ell_1^3\ell_2^3\,\mathrm{d}\nu\,\mathrm{d}s_+\,\mathrm{d}s_-\,\mathrm{d}^2d$, and $\mathrm{d}^2d = (1 + \|y\|^2)^{-2}\mathrm{d}^2y$. All lengths and volumes must use the same coordinate units.

Using the conditional unnormalized density $\widetilde{p}(X \mid \mathcal{M})$ defined in Appendix A, the ordered Cartesian proposal and the target density in molecular coordinates on a regular patch are

$$q_{\theta,\mathrm{ord}}(\Phi_{\mathcal{M}}(z) \mid \mathcal{M}) = \frac{q_\theta(z \mid \mathcal{M})}{J_{\mathcal{M}}(z)}, \tag{42}$$

$$p_{\mathrm{coord}}(z \mid \mathcal{M}) = \frac{\widetilde{p}(\Phi_{\mathcal{M}}(z) \mid \mathcal{M}) J_{\mathcal{M}}(z)}{I(R)}. \tag{43}$$

Thus the target in molecular coordinates is not just the Cartesian Boltzmann factor. Its Jacobian must be included. Either converting the proposal to Cartesian measure or transforming the target to molecular coordinates gives the same importance weight.

The evaluator uses the linear coordinates $w = b_1 - b_2$ and $d_{\mathrm{ev}} = (b_1 + b_2)/6$ with center r. Their measure satisfies $\mathrm{d}X = 27\,\mathrm{d}r\,\mathrm{d}w\,\mathrm{d}d_{\mathrm{ev}}$. An exported density q_{ev} relative to that measure therefore corresponds to the Cartesian density $q_\theta = q_{\mathrm{ev}}/27$, with the same conditioning $\mathcal{M}$ on both sides, while $Z_{\mathrm{gas}} = 27 Z_{\mathrm{gas}}^{\mathrm{ev}}$. The gas-normalized importance weight is unchanged. This constant conversion is distinct from the configuration-dependent Jacobian in Equation 41.

Oxygen exchange. Let P exchange the two oxygen labels in Cartesian coordinates. With angular coordinates $u_z = 2\zeta - 1$, $\varphi = 2\pi\tau$, $\mathrm{d}\nu = \mathrm{d}\zeta\,\mathrm{d}\tau$. With $e_\vartheta(-u) = e_\vartheta(u)$ and $e_\varphi(-u) = -e_\varphi(u)$, the coordinate transformation is

$$(\zeta, \tau, s_+, s_-, d_1, d_2) \longmapsto (1 - \zeta, \tau + \tfrac{1}{2} \bmod 1, s_+, -s_-, d_1, -d_2).$$

The same signs act on the unbounded bend vector y. For standardized coordinates, the corresponding map is $v' = D_v^{-1}[T(a_0 + D_v v) - a_0]$, where $T = \mathrm{diag}(1, -1, 1, -1)$. A sign flip applied directly to v would be incorrect for nonzero centering constants. Orientation antipodality alone therefore does not enforce exchange symmetry of the full guest.

Normalized symmetric proposal. Since $\mathsf{P}^2 = I$ and the Cartesian permutation has unit absolute Jacobian,

$$q_{\theta,\mathrm{sym}}(X \mid \mathcal{M}) = \tfrac{1}{2}\big[q_{\theta,\mathrm{ord}}(X \mid \mathcal{M}) + q_{\theta,\mathrm{ord}}(\mathsf{P}X \mid \mathcal{M})\big]$$

is normalized and invariant. Sampling the ordered proposal and applying P with probability $1/2$ samples this mixture. Its stable log-density is

$$\log q_{\theta,\mathrm{sym}}(X \mid \mathcal{M}) = \mathrm{logaddexp}\big(\log q_{\theta,\mathrm{ord}}(X \mid \mathcal{M}), \log q_{\theta,\mathrm{ord}}(\mathsf{P}X \mid \mathcal{M})\big) - \log 2.$$

This is a log of a density mixture, not an average of log-densities. The target energy and domain are also invariant under P. All normalization and reweighting use the labeled space consistently.

Effect of symmetrization. Importance weighting with $q_{\theta,\mathrm{ord}}$ alone is also unbiased, because it is a normalized density on the same labeled space; symmetrization affects efficiency, not correctness. We compare the two proposals for the trained model, drawing from $q_{\theta,\mathrm{ord}}$ without the random exchange and weighting with $q_{\theta,\mathrm{ord}}$ (Table 3). The trained ordered density is close to, but not exactly, exchange invariant: $|\log q_{\theta,\mathrm{ord}}(X) - \log q_{\theta,\mathrm{ord}}(\mathsf{P}X)|$ averages 0.21 over its samples. Symmetrization raises mean nESS in every quartile and for 72% of the MOFs, while the mean $|\Delta \ln K_H|$ is unchanged (0.044 with and 0.042 without symmetrization).

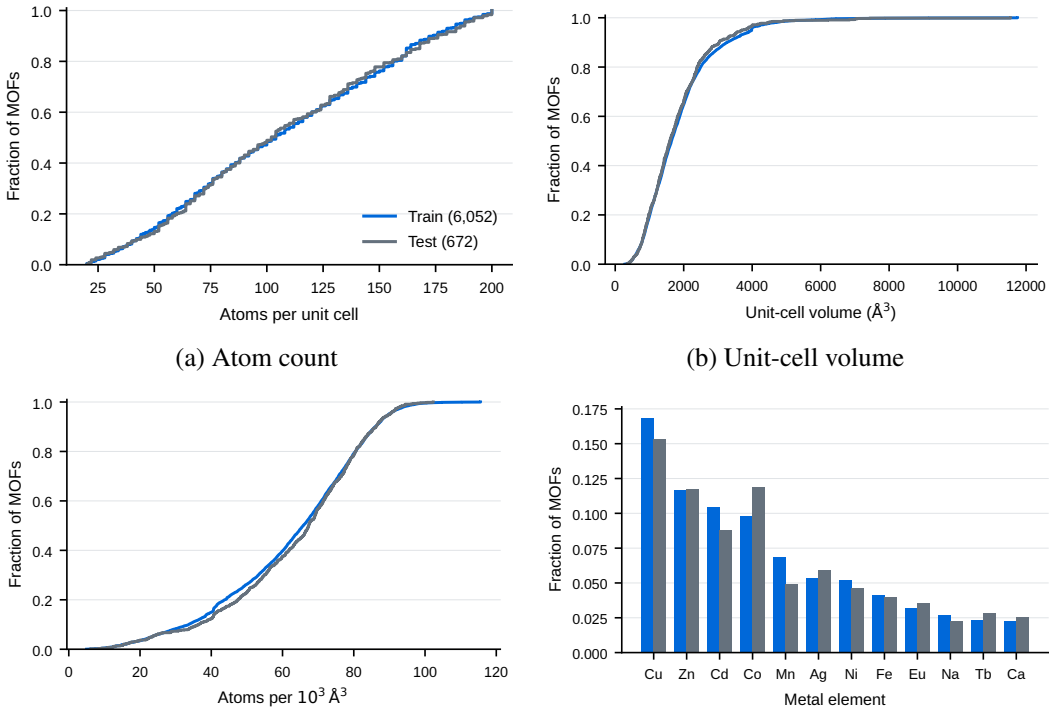
E DATASET CONSTRUCTION AND TRAINING PROTOCOL

Dataset construction. LitMOF-DB (Kim et al., 2025) is a curated subset of MOFs from the Cambridge Structural Database that addresses structural errors documented in existing MOF databases (White et al., 2025). We retain bound-solvent-removed and free-solvent-removed structures without remaining bound solvent. We require a pore-limiting diameter larger than the CO_2 diameter, at most 200 atoms per unit cell, and a pore-volume fraction above 0.1, yielding 11,523 MOFs. We geometry-optimize these structures using MACE-MH-1 with its OMAT-PBE head and D3(BJ) dispersion, with a maximum atomic-force convergence threshold of 0.05 eV $Å^{-1}$, and apply Niggli cell reduction. From 100,000 Widom insertions per MOF, 8,813 structures have at least one candidate adsorption site with negative interaction energy. We use a randomly selected 80% subset of these structures and exclude MOFs with fewer than 20 atoms per unit cell, leaving 6,724 MOFs. We split this set into 6,052 training and 672 held-out test MOFs.

Train/test split. The 6,052 training and 672 test MOFs are disjoint by MOF identifier. Figure E.1 compares their unit-cell size, volume, atomic density, and metal-element composition. The median atom counts are 104 and 102, and the median cell volumes are 1,642 and 1,596 $Å^3$, respectively. Every metal element present in the test set also appears in training, so this split tests transfer to new MOF structures rather than unseen elements.

Training configurations. For each training MOF, we generate 512 flexible CO_2 configurations at 300 K using adaptive sequential Monte Carlo (SMC) with the geometry-optimized MOF held fixed. The SMC target is the conditional Boltzmann density proportional to $e^{-\beta \Delta U(R,X)}$. We adapt each tempering step to a target conditional effective-sample-size fraction of 0.75, systematically resample when the particle effective sample size falls below 0.6 of the population, and apply two local Metropolis–Hastings moves per particle and stage, cycling through translation, rotation, and internal-coordinate perturbations. The final-stage configurations are resampled to uniform weights.

Training protocol. The model has 1,292,375 trainable parameters. We train from scratch for 2,400 steps with 16 MOFs per step and all 512 configurations of each MOF. The objective is the mean negative log-likelihood over configurations within each MOF, averaged over MOFs. We use AdamW (Loshchilov & Hutter, 2019) with weight decay 10^{-6}, a base learning rate of 2.5×10^{-4}, a 50-step linear warmup, and cosine decay to one tenth of the base rate. We clip the gradient norm at 10. During training, the position and orientation densities each include a 0.02-weight uniform component; this component is disabled at inference. We report the final checkpoint without early stopping or checkpoint selection.



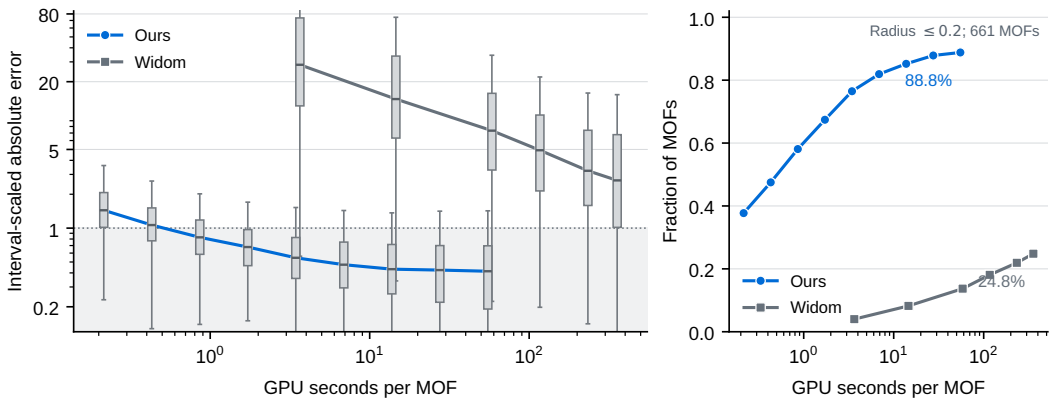
(a) Atom count

(b) Unit-cell volume

(c) Atomic density

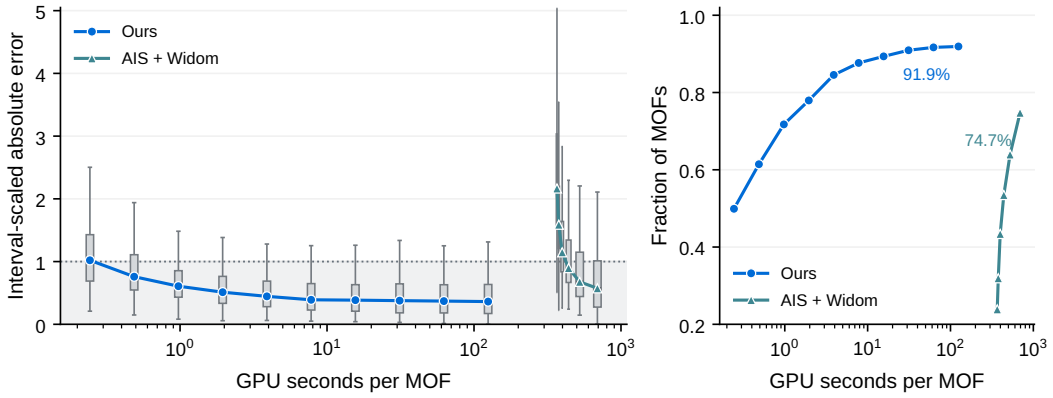
(d) Metal elements

Figure E.1: **Training and test MOF characteristics.** Cumulative distributions are normalized within each split. The 12 most frequent metal elements are shown. Metal bars report the fraction of MOFs containing each element; MOFs containing multiple metal elements contribute to multiple bars.



(a) Error distribution (b) Reference agreement

Figure F.1: **Rigid-gas reference agreement.** Results on 661 MOFs with 95% reference-interval radii at most 0.2 in $\ln K_H^{\mathrm{RR}}$. Left: median interval-scaled absolute error, interquartile boxes, and 1.5-IQR whiskers; the vertical axis is logarithmic. Right: fraction within each MOF's reference interval.



(a) Error distribution (b) Reference agreement

Figure F.2: **Flexible-gas reference agreement.** Results on 594 MOFs with 95% reference-interval radii at most 0.2 in $\ln K_H^{\mathrm{RF}}$. Left: median interval-scaled absolute error, interquartile boxes, and 1.5-IQR whiskers; one marks the reference interval boundary. Right: fraction within each MOF's reference interval. AIS + Widom time includes the 100,000-insertion Widom sweep.

F Reference precision

Baselines. Rigid-gas Widom insertion uses 100,000 uniformly sampled positions and orientations per MOF to estimate K_H^{RR}. For rigid-gas DeepAPD, we combine its predicted CO_2 position density with a uniform orientation density and use 32,768 draws per MOF. The flexible-gas AIS baseline uses 8,192 particles initialized from sites found in the 100,000-insertion Widom sweep to estimate K_H^{RF}.

Reference estimates. For each target, we obtain a separate AIS reference estimate initialized from the union of Widom- and model-discovered sites and evaluated in three frozen-schedule blocks. The 95% intervals are percentile bootstrap intervals from 2,000 resamples of the 6,144 particles in the three frozen-schedule AIS blocks per MOF and target. We select a MOF only when both sides of its interval are at most 0.2 in the corresponding $\ln K_H$ from the reference estimate. An interval-scaled error of one equals the upper or lower interval radius according to the sign of the estimation error, not a fixed $|\Delta \ln K_H|$ of 0.2.

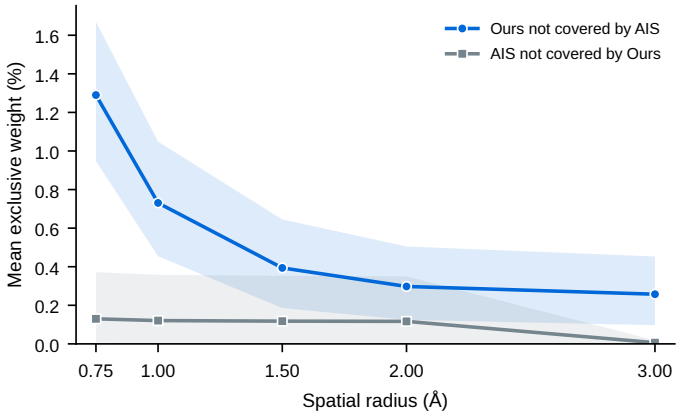


Figure G.1: **Mean exclusive importance weight.** Results over 672 unseen MOFs; shaded bands show 95% bootstrap intervals.

G Spatial Coverage of CO_2 Configurations

We compare Ours with the original 8,192-particle AIS baseline initialized from the 100,000-insertion Widom sweep, without using model-assisted reference particles. For method A and comparator B, the exclusive importance weight at radius r is

$$C_{A\setminus B}(r) = \sum_i \bar{w}_i^A \, \mathbf{1}\left[\min_j d_{\mathrm{PBC}}(c_i^A, c_j^B) > r\right], \tag{44}$$

where c_i is the CO_2 carbon position and $\bar{w}_i^A$ is its normalized importance weight. The distance is measured in Cartesian coordinates across periodic images of the unit cell. This is a directed comparison of sampled positions, not a measure of complete molecular mode coverage. Figure 3(a) reports the fraction of MOFs above 5% exclusive weight in both directions. Figure G.1 reports the mean exclusive weight without this threshold.

The ten MOFs with $C_{\mathrm{Ours}\setminus\mathrm{AIS}}(1.5\,\text{Å}) \geq 0.05$ are shown in Figure 3 and Figures G.2–G.4. For each MOF, we select the ab, ac, or bc plane with the largest Jensen–Shannon divergence between the smoothed, weighted center-of-mass maps. The power-law color scale (exponent 0.32) is shared between Ours and AIS within each MOF.

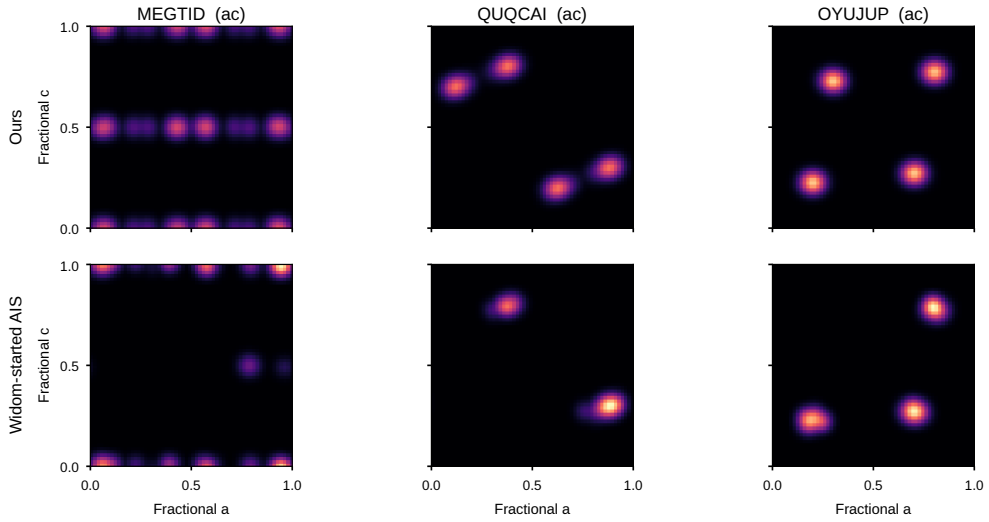


Figure G.2: **Weighted CO_2 center-of-mass distributions.** MEGTID, QUQCAI, and GEBTIU.

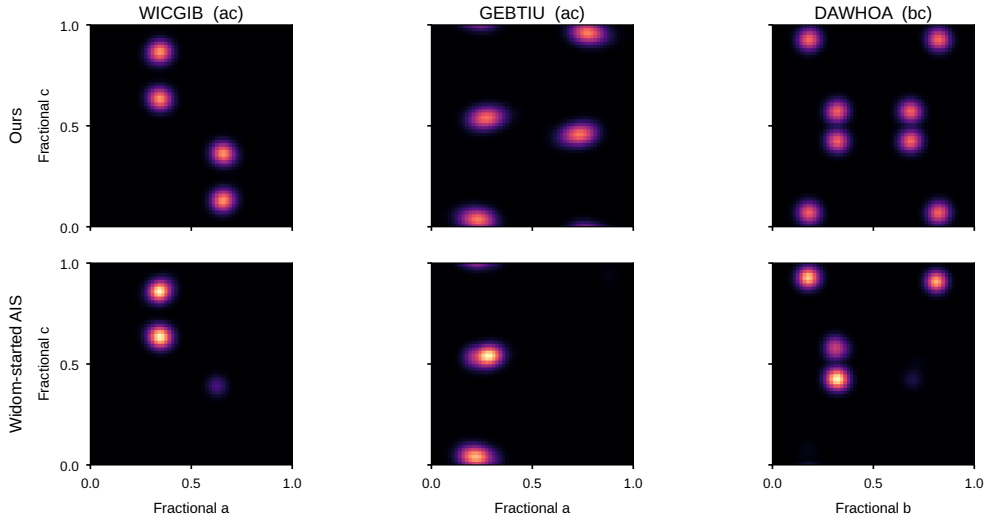


Figure G.3: **Weighted CO_2 center-of-mass distributions.** WICGIB, OYUJUP, and BUKLOJ.

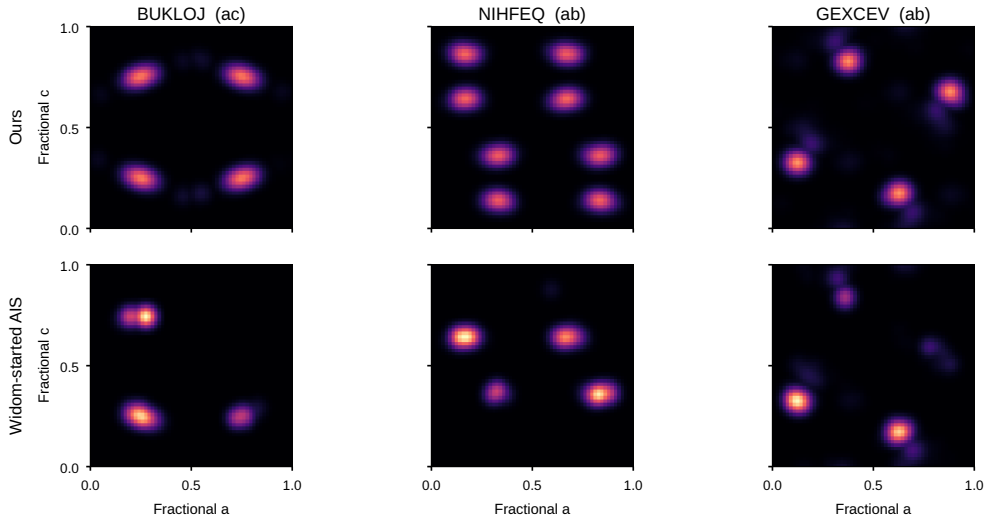


Figure G.4: **Weighted CO_2 center-of-mass distributions.** DAWHOA, NIHFEQ, and GEXCEV.