LEARNING HIERARCHICAL REPRESENTATIONS OF METAL–ORGANIC FRAMEWORKS WITH COUPLED EUCLIDEAN AND HYPERBOLIC GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning representations of structured materials requires combining local geometry with connectivity across chemical building units. Metal–organic frameworks (MOFs) make this challenge explicit: their properties depend on both atomic environments and the organization of metal-containing units and organic linkers. We introduce GeoDuet, a graph representation model that couples Euclidean and hyperbolic geometry. A Euclidean branch encodes periodic local interactions, while a Lorentz branch processes an atomic covalent graph enriched with building-block features. Graph-level attention combines the two representations for property prediction. Building blocks augment atom features rather than define a separate coarse-grained graph. Across the reported hMOF and CoRE MOF adsorption benchmarks, GeoDuet achieves the lowest MAE among the evaluated models; additional comparisons cover MOF-specific architectures, repeated training runs, and accessible surface area. Geometric diagnostics show lower normalized hyperbolicity than degree-matched rewired controls and better graph-distance preservation for the evaluated Lorentz encoder. However, hyperbolicity has little association with per-structure prediction gains. These findings support coupled geometry as a useful inductive bias for MOF pore-related properties, while separating geometric motivation from claims about the causes of predictive improvement.

1 INTRODUCTION

Graph representation learning must often reconcile two kinds of structure: local metric relationships and connectivity across groups of interacting entities. A single embedding geometry need not represent both equally well. Euclidean encoders provide a natural language for distances and orientations, whereas hyperbolic representations offer an inductive bias for hierarchical relationships (Nickel & Kiela, 2017; Ganea et al., 2018; Chami et al., 2019). The challenge is to combine these representations without assuming that every graph is intrinsically hierarchical or that a geometric fit guarantees better downstream prediction.

Metal–organic frameworks (MOFs) provide a concrete setting for this problem. They consist of metal-containing secondary building units (SBUs) connected by organic linkers, producing periodic structures with chemically diverse environments and pores. Gas adsorption depends on this structural organization, making accurate representations useful for screening candidate materials (Wilmer et al., 2012; Chung et al., 2019). Atomic crystal graphs encode local interactions, but their inputs do not necessarily identify the building units to which atoms belong. Conversely, building-block descriptions alone need not preserve the spatial information available in the crystal structure. Learning how these views complement each other is therefore a representation problem as well as an application in materials prediction.

We propose GeoDuet, which couples a Euclidean local encoder with a Lorentz graph encoder. The former uses periodic neighborhoods and geometric edge features. The latter operates on an *atomic covalent graph*, with SBU and linker information mapped back to constituent atoms. Thus, the model retains atom-level resolution while exposing chemical organization through its node features. The two graph representations are fused by attention after the Lorentz representation is mapped to a Euclidean

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.

tangent space. The coupling occurs at the graph-level readout; we do not construct a separate graph whose nodes are pores, cages, or building blocks.

Our geometric motivation is empirical and conditional. Hyperbolic spaces can represent tree-like relationships compactly, but periodic MOF graphs contain cycles and are not generally trees. On a diagnostic collection of 1,377 structures, real atomic covalent graphs have lower normalized Gromov hyperbolicity than degree-matched rewired controls, and the evaluated Lorentz encoder better preserves graph distances. Crucially, neither observation establishes why prediction improves: the distance-preservation encoders differ in capacity and edge features, and graph hyperbolicity is only weakly associated with per-structure prediction gains. We retain these negative findings because they delimit the interpretation of the model.

Our contributions are threefold. First, we formulate a coupled Euclidean–Lorentz representation of MOFs that combines periodic local geometry with building-block-aware atomic connectivity. Second, we evaluate its geometric motivation using matched graph controls, distance-preservation diagnostics, and the relationship between graph geometry and prediction gains. Third, we assess adsorption and accessible-surface-area prediction, including MOF-specific comparisons, a three-run adsorption comparison, and training-cost measurements. The contribution is this combination of representations and evidence; hierarchical MOF modeling and hyperbolic graph learning each have substantial prior work.

2 Related work

Crystal and MOF representations. CGCNN and subsequent crystal graph networks learn from periodic atomic neighborhoods (Xie & Grossman, 2018; Chen et al., 2019). ALIGNN includes bond-angle information, while Matformer, PotNet, and ComFormer develop geometric encodings for periodic crystals (Choudhary & DeCost, 2021; Yan et al., 2022; Lin et al., 2023; Yan et al., 2024). MOF-specific methods expose complementary information: MOFormer uses a structured string representation; MOFNet predicts adsorption isotherms with a graph transformer; MSAIGNN develops multiscale graph representations; and Uni-MOF combines three-dimensional pretraining with adsorption prediction (Cao et al., 2023; Chen et al., 2022a; Li et al., 2025; Wang et al., 2024). These approaches already establish the value of representations beyond a basic atomic graph. GeoDuet investigates the coupling of two embedding geometries while retaining atomic nodes and adding building-unit information as features.

Hyperbolic representation learning. Poincaré embeddings learn hierarchical representations in negatively curved space (Nickel & Kiela, 2017). Hyperbolic neural networks and graph convolutional networks extend this idea to trainable feature transformations and message passing (Ganea et al., 2018; Chami et al., 2019; Liu et al., 2019). Lorentz networks implement operations on the hyperboloid (Chen et al., 2022b; Bdeir et al., 2024), and Riemannian optimization provides updates for manifold-valued parameters (Bécigneul & Ganea, 2019; Kochurov et al., 2020). Combining Euclidean and hyperbolic graph representations also predates our work: GIL learns interactions between the two spaces for node classification and link prediction (Zhu et al., 2020). GeoDuet instead assigns periodic local geometry and annotated covalent connectivity to separate encoders and couples their graph-level summaries for MOF regression. This task-specific construction, rather than the general idea of combining geometries, is our focus. The Lorentz and Poincaré models describe the same hyperbolic geometry; differences between trained implementations do not establish different underlying expressivity.

3 Representations and geometric motivation

Periodic structures and atomic graphs. A crystal is represented by $\mathcal{S} = (\mathbf{L}, \mathbf{P}, \mathbf{A})$, where $\mathbf{L} = [\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3] \in \mathbb{R}^{3\times3}$ contains lattice vectors, $\mathbf{P} = [\mathbf{p}_1, \ldots, \mathbf{p}_n] \in \mathbb{R}^{3\times n}$ contains Cartesian atomic positions, and $\mathbf{A}$ contains atom features. Its periodic extension is

$$\{(\mathbf{a}_i, \mathbf{p}_i + \mathbf{L}\mathbf{t}) : i \in \{1, \ldots, n\},\ \mathbf{t} \in \mathbb{Z}^3\}. \tag{1}$$

We distinguish a periodic local-neighbor graph G_{loc} from the covalent-connectivity graph G_{atom}. Both use atomic indices; their edges encode different relationships. Building-block extraction identifies SBUs and linkers and associates constituent atoms with these units. It supplies node features, not a third, coarse-grained message-passing graph.

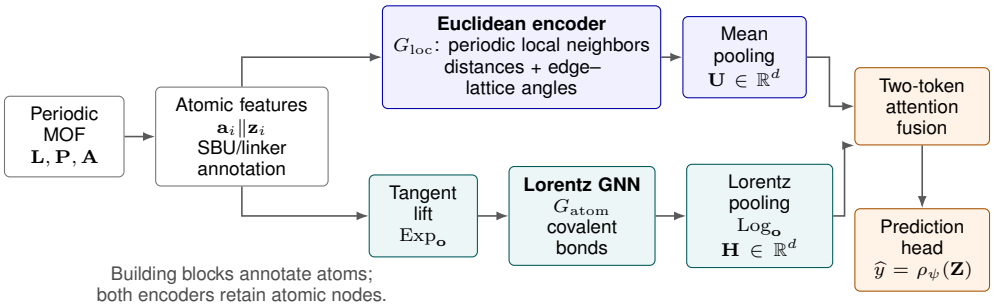


Figure 1: **GeoDuet overview.** Block descriptors are attached to constituent atoms. The Euclidean branch encodes periodic local geometry, and the Lorentz branch processes atomic covalent connectivity. After separate graph pooling, the Lorentz representation is mapped to the tangent space for attention fusion and scalar property prediction. Arrows show data flow, not physical gas-transport paths.

Lorentz geometry. For $K > 0$, the Lorentz model of curvature $-1/K$ is

$$\mathbb{H}^{d,K} = \{\mathbf{x} \in \mathbb{R}^{d+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = -K,\ x_0 > 0\}, \qquad \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = -x_0 y_0 + \sum_{r=1}^{d} x_r y_r. \tag{2}$$

The ambient bilinear form has matrix $\text{diag}(-1, 1, \ldots, 1)$; K sets the hyperboloid radius. Its restriction to a tangent space gives the Riemannian metric. The geodesic distance is

$$d_{\mathcal{L}}(\mathbf{x}, \mathbf{y}) = \sqrt{K}\, \text{arcosh}(-\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} / K)\,. \tag{3}$$

The origin is $\mathbf{o} = (\sqrt{K}, \mathbf{0})$. Exponential and logarithmic maps connect the manifold to its tangent space, enabling Euclidean inputs and readouts (Appendix B).

Why consider hyperbolic representations? Tree-embedding results motivate negatively curved spaces for hierarchical metrics (Sarkar, 2011; Nickel & Kiela, 2017). Applying this motivation to MOFs requires checking the actual input graph rather than treating a schematic pore hierarchy as evidence. We therefore measure normalized four-point hyperbolicity of G_{atom} and compare it with rewired, nearest-neighbor, and spanning-tree controls (Section 5.4). This diagnostic concerns finite graph distances. It neither assigns a physical curvature to a MOF nor supplies a universal threshold for stable materials, and it imposes no guarantee on supervised prediction error.

4 GeoDuet

4.1 Building-block-aware atom features

Let $\mathcal{B} = \{B_1, \ldots, B_M\}$ denote the extracted SBUs and organic linkers, and let $b(i)$ identify the unit assigned to atom i. We denote the feature encoder for a building unit by $F_{\mathcal{B}}$ and write

$$\mathbf{z}_i = F_{\mathcal{B}}(B_{b(i)}) \in \mathbb{R}^{d_z}, \qquad \mathbf{v}_i = \mathbf{a}_i \,\|\, \mathbf{z}_i \in \mathbb{R}^{92+d_z}. \tag{4}$$

Here $\mathbf{a}_i$ is the 92-dimensional CGCNN atomic descriptor and $\|$ denotes concatenation. Atoms assigned to the same building unit receive its block descriptor while retaining their individual atomic features. The notation $F_{\mathcal{B}}$ separates descriptor construction from graph message passing: building blocks annotate nodes, and all Lorentz messages remain indexed by atoms and covalent edges. MOFstructure and MOFfragmentor provide the structure-processing tools used for this extraction. Appendix C describes the atom-assignment interface, forward computation, and training configuration.

Complementary neighborhoods. The two graphs expose different relationships on the same atomic indices. A local-neighbor edge records spatial proximity, including nearby atoms that need not be covalently bonded. A path in G_{atom} instead follows chemical connectivity through the framework.

Block annotation adds the chemical-unit context to these atomic representations without collapsing all atoms of a linker or SBU into a single node. Consequently, the Lorentz branch does not receive a precomputed tree: its input retains cycles and the connectivity produced by the extraction pipeline. The design combines two graph views and two embedding geometries. Distinguishing their individual contributions requires controls that match both the graph inputs and encoder capacity, which motivates our qualified interpretation of the ablations.

4.2 Euclidean local encoder

For a periodic edge $(i, j, \mathbf{t}) \in E_{\text{loc}}$, define the displacement $\mathbf{e}_{ij\mathbf{t}} = \mathbf{p}_j + \mathbf{L}\mathbf{t} - \mathbf{p}_i$ and its length $r_{ij\mathbf{t}} = \|\mathbf{e}_{ij\mathbf{t}}\|_2$. The encoder combines an inverse-distance feature $\eta(r) = c/r$, with $c = -0.75$, and orientation features relative to each lattice vector:

$$q_{ij\mathbf{t},k} = \cos\theta_{ij\mathbf{t},k} = \frac{\mathbf{e}_{ij\mathbf{t}}^{\top}\mathbf{l}_k}{r_{ij\mathbf{t}}\|\mathbf{l}_k\|_2}, \qquad k = 1, 2, 3. \tag{5}$$

The scalar c is a coefficient, not an exponent. These are edge–lattice angles, distinct from three-atom bond angles. Radial basis functions (RBFs) and spherical Fourier–Bessel features (SBFs) encode the inputs before a learned projection:

$$\widehat{\mathbf{e}}_{ij\mathbf{t}} = \Phi\Big(\Psi_{\text{RBF}}(\eta(r_{ij\mathbf{t}})), [\Psi_{\text{RBF}}(\eta(\|\mathbf{l}_k\|_2)), \Psi_{\text{SBF}}(\eta(r_{ij\mathbf{t}}), q_{ij\mathbf{t},k})]_{k=1}^{3}\Big). \tag{6}$$

This follows geometric feature designs used in crystal and molecular networks (Lin et al., 2023; Yan et al., 2024; Gasteiger et al., 2020). The features are invariant to a joint rigid rotation of the coordinates and lattice. This statement does not imply invariance to an arbitrary change of lattice basis.

More explicitly, under $\mathbf{p}_i \mapsto \mathbf{R}\mathbf{p}_i + \mathbf{b}$ and $\mathbf{L} \mapsto \mathbf{R}\mathbf{L}$, with $\mathbf{R}^{\top}\mathbf{R} = \mathbf{I}$, each edge displacement becomes $\mathbf{R}\mathbf{e}_{ij\mathbf{t}}$. Its norm and its normalized inner products with the rotated lattice vectors are unchanged. Shared node transformations, neighborhood summation, and graph pooling also respect consistent atom reindexing, provided graph construction and block assignment respect that reindexing. These properties explain which coordinate choices the representation removes. They do not establish invariance to unit-cell replication, arbitrary lattice-basis changes, or alternative building-block decompositions; those depend on preprocessing choices beyond the rigid-motion calculation.

With initial embeddings $\mathbf{u}_i^{(0)} = E_E(\mathbf{v}_i)$, local message passing takes the form

$$\mathbf{u}_i^{(\ell)} = \sigma\left(g_E^{(\ell)}\left(\mathbf{u}_i^{(\ell-1)}, \sum_{(j,\mathbf{t})\in\mathcal{N}_{\text{loc}}(i)} f_E^{(\ell)}(\mathbf{u}_i^{(\ell-1)}, \mathbf{u}_j^{(\ell-1)}, \widehat{\mathbf{e}}_{ij\mathbf{t}})\right)\right). \tag{7}$$

The functions $f_E^{(\ell)}$ and $g_E^{(\ell)}$ are trainable transformations. Mean pooling after L layers gives $\mathbf{U} = n^{-1}\sum_i \mathbf{u}_i^{(L)}$.

4.3 Lorentz encoder on the atomic covalent graph

A learned input map E_H projects $\mathbf{v}_i$ to $\mathbf{w}_i \in \mathbb{R}^d$. We lift $(0, \mathbf{w}_i)$ from $T_{\mathbf{o}}\mathbb{H}^{d,K}$ onto the manifold:

$$\mathbf{h}_i^{(0)} = \text{Exp}_{\mathbf{o}}(0, \mathbf{w}_i) = \left(\sqrt{K}\cosh\frac{\|\mathbf{w}_i\|_2}{\sqrt{K}}, \sqrt{K}\frac{\sinh(\|\mathbf{w}_i\|_2/\sqrt{K})}{\|\mathbf{w}_i\|_2}\mathbf{w}_i\right), \tag{8}$$

with the continuous value $\mathbf{o}$ when $\mathbf{w}_i = \mathbf{0}$. Lorentz message passing uses covalently connected neighbors:

$$\mathbf{h}_i^{(\ell)} = \Pi_{\mathcal{L}}\left[g_{\mathcal{L}}^{(\ell)}\left(\mathbf{h}_i^{(\ell-1)}, \sum_{j\in\mathcal{N}_{\text{atom}}(i)} f_{\mathcal{L}}^{(\ell)}(\mathbf{h}_i^{(\ell-1)} \oplus_{\mathcal{L}} \mathbf{h}_j^{(\ell-1)})\right)\right]. \tag{9}$$

The maps $f_{\mathcal{L}}^{(\ell)}$ and $g_{\mathcal{L}}^{(\ell)}$ are manifold-aware neural transformations (Chen et al., 2022b; Bdeir et al., 2024). Neighborhood summation is over G_{atom}, not over all atoms or a block graph. To make

the manifold constraint explicit, let $\mathbf{x}_r = (t_r, \mathbf{s}_r)$ denote N Lorentz points with the same radius parameter K. Their concatenation is

$$\bigoplus_{\mathcal{L},r=1}^{N} \mathbf{x}_r = \left(\sqrt{\sum_{r=1}^{N} t_r^2 - (N-1)K},\ \mathbf{s}_1 \| \cdots \| \mathbf{s}_N \right). \tag{10}$$

The reconstructed time coordinate ensures Lorentz norm $-K$ in the enlarged space. For a future-directed time-like vector, normalization is

$$\Pi_{\mathcal{L}}(\mathbf{x}) = \sqrt{\frac{K}{-\langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}}}}\, \mathbf{x}, \qquad x_0 > 0, \quad \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} < 0. \tag{11}$$

The domain restriction matters: scalar normalization cannot turn a space-like vector into a time-like one. The neural transformations must therefore respect this restriction before normalization. Appendix B gives the algebraic justification and the limiting cases of the manifold maps.

We obtain a graph representation by normalizing the ambient mean,

$$\mathbf{H}^{\mathcal{L}} = \Pi_{\mathcal{L}} \left(\frac{1}{n} \sum_{i=1}^{n} \mathbf{h}_i^{(L)} \right), \qquad \mathbf{H} = \left[\mathrm{Log}_{\mathbf{o}}(\mathbf{H}^{\mathcal{L}}) \right]_{1:d}. \tag{12}$$

The last expression keeps the spatial coordinates of the tangent vector. For $\mathbf{H}^{\mathcal{L}} = (t, \mathbf{s})$, this readout can be written explicitly as

$$\mathbf{H} = \sqrt{K}\, \mathrm{arcosh}\left(\frac{t}{\sqrt{K}} \right) \frac{\mathbf{s}}{\|\mathbf{s}\|_2}, \tag{13}$$

with value $\mathbf{0}$ when $\mathbf{H}^{\mathcal{L}} = \mathbf{o}$. Its direction follows the spatial component, and its Euclidean norm equals the hyperbolic distance from the origin. Thus, the map preserves this radial information when constructing a common coordinate system for fusion; it does not preserve every pairwise hyperbolic distance as a Euclidean distance. The pooled representation summarizes learned atomic connectivity, rather than explicitly reconstructing pore channels or gas-transport paths.

4.4 Coupling the representations

Stacking the two graph embeddings gives $\mathbf{E} = [\mathbf{U}; \mathbf{H}] \in \mathbb{R}^{2 \times d}$. Multi-head attention learns interactions between these two representation tokens:

$$\mathbf{Q}_r = \mathbf{E}\mathbf{W}_r^Q, \quad \mathbf{K}_r = \mathbf{E}\mathbf{W}_r^K, \quad \mathbf{V}_r = \mathbf{E}\mathbf{W}_r^V, \tag{14}$$

$$\mathbf{Z} = \left[\Big\|_{r=1}^{h} \mathrm{softmax}\left(\frac{\mathbf{Q}_r \mathbf{K}_r^\top}{\sqrt{d_k}} \right) \mathbf{V}_r \right] \mathbf{W}^O, \qquad \widehat{y} = \rho_\psi(\mathbf{Z}). \tag{15}$$

The readout $\rho_\psi : \mathbb{R}^{2 \times d} \to \mathbb{R}$ produces a scalar property prediction. The encoders and readout are trained jointly with the task loss. RiemannianAdam handles manifold-constrained parameters; ordinary weight matrices are Euclidean parameters. Manifold-valued activations alone do not require every parameter to be optimized on a manifold.

Each attention head forms a 2×2 matrix: its rows correspond to the Euclidean and tangent-space summaries, and its columns determine how each summary draws information from the two branches. These weights depend on the input structure through $\mathbf{E}$, allowing the readout to adapt the combination of representations across MOFs. The attention step acts on graph summaries, so its token count is independent of the number of atoms. Atom-level information exchange remains within the respective message-passing graphs. This separates the computational role of fusion from that of the encoders: fusion combines their learned summaries without constructing all-pairs attention between atoms. Attention weights alone do not establish a physical contribution of either branch to adsorption.

Forward computation. Table 1 consolidates the complete path from a periodic structure to a property prediction. Both branches start from atom-indexed features, process their respective graph neighborhoods, and reduce their node representations before fusion. The exponential map supplies valid initial Lorentz points, while the logarithmic map places the pooled Lorentz representation in the coordinate space used by the prediction head. Building-block extraction precedes these learned operations; it supplies the chemical annotations used throughout the forward pass.

Table 1: GeoDuet forward pass. All message-passing nodes are atoms.

1	Read the periodic structure $(\mathbf{L}, \mathbf{P}, \mathbf{A})$ and construct G_{loc} and G_{atom}.
2	Extract SBUs/linkers, map their descriptors to constituent atoms, and form $\mathbf{v}_i = \mathbf{a}_i \Vert \mathbf{z}_i$.
3	Encode distances and edge–lattice angles on G_{loc} using Equations 5–6.
4	Apply Euclidean message passing on G_{loc} and mean-pool to obtain $\mathbf{U}$.
5	Lift the Lorentz-branch inputs with $\mathrm{Exp}_{\mathbf{o}}$ and apply Lorentz message passing on G_{atom}.
6	Normalize the mean Lorentz embedding, apply $\mathrm{Log}_{\mathbf{o}}$, and keep the spatial coordinates to obtain $\mathbf{H}$.
7	Fuse the two tokens $[\mathbf{U}; \mathbf{H}]$ with attention and apply the scalar readout ρ_ψ.

Table 2: CH_4 adsorption on hMOF: test MAE (mol kg^{-1}) at five pressures (bar). Entries are point estimates; lower is better.

Model	0.05	0.5	0.9	2.5	4.5
Stoichiometric-120	0.044	0.259	0.372	0.587	0.739
SOAP	0.036	0.161	0.172	0.330	0.431
RACs	0.031	0.198	0.293	0.518	0.673
CGCNN	0.025	0.119	0.195	0.391	0.483
MOFormer	0.031	0.156	0.216	0.353	0.431
Matformer	0.023	0.085	0.112	0.186	0.242
PotNet	0.018	0.079	0.108	0.187	0.239
ComFormer	0.019	0.067	0.091	0.146	0.180
DsLAPM	0.018	0.074	0.101	0.179	0.224
GeoDuet	**0.016**	**0.063**	**0.079**	**0.130**	**0.164**

5 Experiments

5.1 Tasks and evaluation settings

We evaluate CH_4 and CO_2 adsorption on hMOF (Wilmer et al., 2012), Ar adsorption on CoRE MOF 2019 (Chung et al., 2019), and accessible surface area as an additional pore-related target. The source hMOF collection contains 137,953 structures. The adsorption conditions are listed in Tables 2–4; MAE is lower-is-better. hMOF uptake is expressed in mol kg^{-1}. For CoRE, both training and MAE evaluation use $\log(1+y)$ targets (Table 4); its errors are therefore reported on the transformed scale, not in mol kg^{-1}. Appendix C.2 gives the dataset and metric definitions.

The primary benchmark includes Stoichiometric-120, SOAP, and RAC descriptors, alongside CGCNN, MOFormer, Matformer, PotNet, ComFormer, and DsLAPM (Meredig et al., 2014; Himanen et al., 2020; Moosavi et al., 2020; Xie & Grossman, 2018; Cao et al., 2023; Yan et al., 2022; Lin et al., 2023; Yan et al., 2024; Wang et al., 2025). Appendix A.2 reports supplementary from-scratch comparisons with MOF-specific architectures.

GeoDuet uses three message-passing layers, hidden dimension 256, four fusion attention heads, batch size 64, and 100 training epochs on NVIDIA H800 hardware, with a one-cycle learning-rate schedule. Most benchmark entries are point estimates. We report repeated runs only where available and do not infer significance for unreplicated conditions.

5.2 Adsorption prediction and repeated runs

GeoDuet has the lowest reported MAE in all ten hMOF conditions. Relative to the best competing entry in each column, reductions range from 6.0% to 13.2% for CH_4 and from 4.6% to 10.5% for CO_2. For example, CH_4 MAE at 4.5 bar decreases from 0.180 to 0.164, while CO_2 MAE at 2.5 bar decreases from 0.346 to 0.330. The relative gains do not increase monotonically with pressure, so these results alone do not support a pressure-dependent pore-filling mechanism for the improvement.

Table 3: CO_2 adsorption on hMOF: test MAE (mol kg^{-1}) at five pressures (bar). Entries are point estimates; lower is better.

Model	0.01	0.05	0.1	0.5	2.5
Stoichiometric-120	0.098	0.283	0.429	0.930	1.791
SOAP	0.059	0.126	0.222	0.477	0.901
RACs	0.079	0.232	0.355	0.853	1.735
CGCNN	0.055	0.135	0.196	0.473	0.984
MOFormer	0.062	0.167	0.238	0.505	0.889
Matformer	0.047	0.103	0.139	0.279	0.527
PotNet	0.041	0.093	0.130	0.268	0.499
ComFormer	0.038	0.083	0.113	0.218	0.346
DsLAPM	0.040	0.093	0.124	0.253	0.446
GeoDuet	**0.034**	**0.076**	**0.107**	**0.205**	**0.330**

Table 4: Ar adsorption on CoRE MOF 2019: test MAE on the $\log(1 + y)$ scale at three pressures (bar). Both training and evaluation use transformed targets. Entries are point estimates; lower is better.

Model	0.001	0.01	0.1
Stoichiometric-120	1.086	1.112	0.959
SOAP	0.854	0.998	0.926
RACs	0.881	1.022	0.943
CGCNN	0.807	0.994	0.898
Matformer	0.767	0.905	0.850
PotNet	0.719	0.837	0.783
ComFormer	0.687	0.807	0.758
DsLAPM	0.709	0.812	0.760
GeoDuet	**0.672**	**0.794**	**0.747**

Table 5: CH_4 adsorption at 4.5 bar: MAE across three paired training runs (mean $\pm$ standard deviation).

Model	Run 1	Run 2	Run 3	Mean $\pm$ SD
ComFormer	0.180	0.178	0.163	0.174 ± 0.009
GeoDuet	**0.164**	**0.160**	**0.145**	$\mathbf{0.156 \pm 0.010}$

On CoRE MOF, GeoDuet has the lowest reported error in each of the three Ar conditions. The reductions relative to ComFormer are 2.2%, 1.6%, and 1.5%, respectively. These are modest point-estimate improvements on experimentally derived structures, not evidence that every condition yields a statistically reliable gain.

For CH_4 at 4.5 bar, three paired runs give MAE 0.156 ± 0.010 for GeoDuet and 0.174 ± 0.009 for ComFormer, with lower error for GeoDuet in each run (Table 5). Using the displayed run-level values, the relative reduction of mean MAE is 9.98%; Appendix A.3 distinguishes this calculation from averaging per-run percentages. This supports repeatability for the tested condition; three runs cannot establish robustness across all targets or data partitions.

5.3 Component and manifold analysis

Table 6 examines building-block features and the coupled representation. Adding block features to the Euclidean configuration reduces CH_4 MAE from 0.226 to 0.178; using the Lorentz branch reduces it further to 0.164, a 7.9% reduction relative to the atom-plus-block configuration. The first comparison supports chemical block annotation, and the second supports the evaluated dual-geometry configuration. Neither comparison by itself rules out differences in capacity or optimization as contributing factors.

Table 6: Component comparison for CH_4 at 4.5 bar. BB denotes building-block features attached to atoms. Entries are single-run MAE in mol kg^{-1}. The task is evaluated on hMOF.

Configuration	MAE ↓
Base: Euclidean, atomic features only	0.226
Base + BB: Euclidean, block-aware features	0.178
GeoDuet: Euclidean + Lorentz, block-aware features	**0.164**

The Poincaré-ball variant has MAE 0.215 versus 0.164 for CH_4 at 4.5 bar and 0.421 versus 0.330 for CO_2 at 2.5 bar. These results compare implementations and training settings of isometric hyperbolic models. Without optimization diagnostics or a matched tuning study, they do not establish numerical instability as the cause of the difference (Appendix C.3).

5.4 Geometry: supporting evidence and its limits

Table 7: Normalized hyperbolicity on the geometric diagnostic subset.

Graph	Median δ_{norm}	Interquartile range
Real MOF covalent graph	0.2308	[0.1944, 0.2727]
Degree-matched rewiring	0.2500	[0.2500, 0.2727]
kNN-16 atomic graph	0.3750	[0.3333, 0.4000]
Spanning tree	0.0000	[0.0000, 0.0000]

Table 8: Graph-distance preservation of the compared learned representations. Higher Spearman correlation and lower stress/distortion are better.

Representation	Spearman ↑	Normalized stress ↓	Relative distortion ↓
Lorentz	**0.634**	**0.311**	**0.401**
Euclidean	0.438	0.402	0.517

Table 9: Association between normalized hyperbolicity and per-MOF Lorentz-branch prediction gain in the 1,377-structure diagnostic.

Adsorption condition	Spearman ρ	p
CH_4, 4.5 bar	0.015	0.58
CO_2, 2.5 bar	0.082	0.0022

We examine $\delta_{\text{norm}}(G) = \delta(G)/\operatorname{diam}(G)$, where $\delta(G)$ is four-point Gromov hyperbolicity. Smaller values indicate a more tree-like shortest-path metric under this diagnostic. On 1,377 MOFs, the median is 0.2308 for the real covalent graphs, 0.2500 for degree-matched rewired controls, and 0.3750 for kNN-16 graphs; spanning trees have value zero. The reported paired real-minus-rewired difference is -0.0278, with bootstrap 95% interval $[-0.0313, -0.0221]$, and 63.7% of real graphs have the lower value. These are relative comparisons with controls, not evidence that MOFs satisfy a fixed hyperbolicity threshold.

The evaluated Lorentz representations also preserve graph distances better: Spearman correlation is 0.634 versus 0.438, normalized stress is 0.311 versus 0.402, and relative distortion is 0.401 versus 0.517. Because the encoders differ in capacity and edge-feature design, this is a descriptive diagnostic rather than a controlled attribution to curvature. Tables 7–9 present the diagnostics; Appendix A.1 provides additional interpretation.

The available evidence does not support a monotonic relationship between tree-likeness and prediction gain. The Spearman correlation between δ_{norm} and per-MOF gain is $\rho = 0.015$ ($p = 0.58$) for CH_4 at 4.5 bar and $\rho = 0.082$ ($p = 0.0022$) for CO_2 at 2.5 bar. The latter association is weak despite its small p-value, and its positive sign does not support greater gains for lower hyperbolicity. Graph

geometry therefore provides a motivation for the representation, but is not a sufficient explanation of downstream benefit.

5.5 ACCESSIBLE SURFACE AREA AND COMPUTATIONAL COST

Table 10: Accessible-surface-area prediction on MOFs. MAE is in m^2/g; values are mean $\pm$ SD over the reported runs.

Model	MAE ↓	R^2 ↑
GeoDuet	$\mathbf{172.0 \pm 2.9}$	$\mathbf{0.9704 \pm 0.0010}$
Atom + block features	183.4 ± 2.1	0.9688 ± 0.0007
iComFormer	340.7 ± 9.1	0.9183 ± 0.0026

Table 11: Reported training cost in the matched CO_2 adsorption setup on H800 hardware. Building-block preprocessing is not included in these training times.

Model	Params (M)	ms/batch	min/epoch	Peak GB
Matformer	3.64	634	18.2	45.9
iComFormer	4.97	538	15.4	35.9
GeoDuet	4.42	562	16.1	23.2
GeoDuet (TF32)	4.42	370	10.6	23.2

Accessible surface area tests whether the representation is useful for a second MOF property. GeoDuet achieves MAE 172.0 ± 2.9 m^2/g, compared with 183.4 ± 2.1 for the atom-plus-block baseline and 340.7 ± 9.1 for iComFormer. The corresponding R^2 values are 0.9704 ± 0.0010, 0.9688 ± 0.0007, and 0.9183 ± 0.0026. The reported paired bootstrap MAE reductions are 11.3 m^2/g (95% interval [2.8, 19.7]) and 168.6 m^2/g ([151.0, 186.5]), respectively. These results broaden evaluation within pore-related MOF properties; they do not establish transfer to unrelated crystal targets or demonstrate that an adsorption-trained encoder transfers without retraining. Table 10 reports the complete comparison; Appendix A.4 gives the paired comparisons.

Table 11 reports the CO_2 training-cost comparison. In this setting, GeoDuet has 4.42M parameters and requires 562 ms per batch, versus 4.97M and 538 ms for iComFormer. Peak memory is 23.2 GB versus 35.9 GB. TF32 reduces the reported GeoDuet batch time to 370 ms, but this timing alone does not establish unchanged predictive accuracy. Building-block extraction is a separate preprocessing cost and is excluded from these training measurements (Appendix A.5).

6 DISCUSSION AND CONCLUSION

GeoDuet combines periodic local geometry with a Lorentz representation of building-block-aware atomic connectivity. The reported adsorption benchmarks, component comparisons, and surface-area results support this coupled representation for MOF pore-related prediction. Geometric diagnostics provide additional motivation, while the weak relationship between hyperbolicity and per-structure gains prevents a simple geometric explanation of accuracy.

The scope of these findings is defined by the evaluation. Repeated training covers one adsorption condition, and the distance-preservation encoders differ in capacity and edge features. Further controls with matched graph inputs, parameter budgets, and optimization settings would help isolate the effect of geometry. Generalization to unseen framework topologies or chemical building units also remains an open question. Within the evaluated settings, the results support coupling local geometric information with chemically annotated connectivity for MOF property prediction. The weak association between hyperbolicity and prediction gains emphasizes that graph geometry, chemistry, and the downstream target must be considered together.

AI USE STATEMENT

We used generative AI tools for manuscript organization, language editing, LaTeX formatting, a schematic figure, literature checks, and the presentation of existing experimental results. No reported experimental measurements were generated by the tool. The authors are responsible for verifying the manuscript, references, and experimental claims.

REFERENCES

Ahmad Bdeir, Kristian Schwethelm, and Niels Landwehr. Fully hyperbolic convolutional neural networks for computer vision. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ekz1hN5QNh.

Gary Bécigneul and Octavian-Eugen Ganea. Riemannian adaptive optimization methods. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=r1eiqi09K7.

Zhonglin Cao, Rishikesh Magar, Yuyang Wang, and Amir Barati Farimani. Moformer: self-supervised transformer model for metal–organic framework property prediction. *Journal of the American Chemical Society*, 2023.

Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. *Advances in Neural Information Processing Systems*, 2019.

Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 2019.

Pin Chen, Rui Jiao, Jinyu Liu, Yang Liu, and Yutong Lu. Interpretable graph transformer network for predicting adsorption isotherms of metal–organic frameworks. *Journal of Chemical Information and Modeling*, 62(22):5446–5456, 2022a. doi: 10.1021/acs.jcim.2c00876.

Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Fully hyperbolic neural networks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022b.

Kamal Choudhary and Brian DeCost. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 2021.

Yongchul G Chung, Emmanuel Haldoupis, Benjamin J Bucior, Maciej Haranczyk, Seulchan Lee, Hongda Zhang, Konstantinos D Vogiatzis, Marija Milisavljevic, Sanliang Ling, Jeffrey S Camp, et al. Advances, updates, and analytics for the computation-ready, experimental metal–organic framework database: CoRE MOF 2019. *Journal of Chemical & Engineering Data*, 2019.

Nathann Cohen, David Coudert, and Aurélien Lancin. On computing the Gromov hyperbolicity. *Journal of Experimental Algorithmics (JEA)*, 2015.

Hervé Fournier, Anas Ismail, and Antoine Vigneron. Computing the Gromov hyperbolicity of a discrete metric space. *Information Processing Letters*, 2015.

Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in Neural Information Processing Systems*, 2018.

Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*, 2020.

Lauri Himanen, Marc OJ Jäger, Eiaki V Morooka, Filippo Federici Canova, Yashasvi S Ranawat, David Z Gao, Patrick Rinke, and Adam S Foster. Dscribe: Library of descriptors for machine learning in materials science. *Computer Physics Communications*, 2020.

Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in PyTorch. *arXiv preprint arXiv:2005.02819*, 2020.

Lujun Li, Haibin Yu, and Zhuo Wang. Attention-based interpretable multiscale graph neural network for MOFs. *Journal of Chemical Theory and Computation*, 21(3):1369–1381, 2025. doi: 10.1021/acs.jctc.4c01525.

Yuchao Lin, Keqiang Yan, Youzhi Luo, Yi Liu, Xiaoning Qian, and Shuiwang Ji. Efficient approximations of complete interatomic potentials for crystal property prediction. In *International Conference on Machine Learning*, 2023.

Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *Advances in Neural Information Processing Systems*, 2019.

Bryce Meredig, Ankit Agrawal, Scott Kirklin, James E Saal, Jeff W Doak, Alan Thompson, Kunpeng Zhang, Alok Choudhary, and Christopher Wolverton. Combinatorial screening for new materials in unconstrained composition space with machine learning. *Physical Review B*, 2014.

Seyed Mohamad Moosavi, Aditya Nandy, Kevin Maik Jablonka, Daniele Ongari, Jon Paul Janet, Peter G Boyd, Yongjin Lee, Berend Smit, and Heather J Kulik. Understanding the diversity of the metal-organic framework ecosystem. *Nature communications*, 2020.

Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Rik Sarkar. Low distortion Delaunay embedding of trees in hyperbolic plane. In *Graph Drawing*, volume 7034, pp. 355–366. Springer, 2011. doi: 10.1007/978-3-642-25878-7_34.

Bing Wang, Wei Zhou, Yong sheng Ren, Jia jia Xu, Shu Zhan, and Wenhui Ma. Local angle information propagation model based on dual scale for crystal property prediction. *Computational Materials Science*, 2025.

Jingqi Wang, Jiapeng Liu, Hongshuai Wang, Musen Zhou, Guolin Ke, Linfeng Zhang, Jianzhong Wu, Zhifeng Gao, and Diannan Lu. A comprehensive transformer-based approach for high-accuracy gas adsorption predictions in metal-organic frameworks. *Nature Communications*, 15:1904, 2024. doi: 10.1038/s41467-024-46276-x.

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*, 2012.

Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 2018.

Keqiang Yan, Yi Liu, Yuchao Lin, and Shuiwang Ji. Periodic graph transformers for crystal material property prediction. *Advances in Neural Information Processing Systems*, 2022.

Keqiang Yan, Cong Fu, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. Complete and efficient graph transformers for crystal material property prediction. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=BnQY9XiRAS.

Shichao Zhu, Shirui Pan, Chuan Zhou, Jia Wu, Yanan Cao, and Bin Wang. Graph geometry interaction learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/551fdbb810aff145c114b93867dd8bfd-Abstract.html.

A ADDITIONAL EMPIRICAL ANALYSES

These analyses provide geometric, statistical, and computational context for the main results. Each addresses a distinct evaluation question and is interpreted within its corresponding comparison.

A.1 ATOMIC-GRAPH GEOMETRY

The Lorentz encoder operates on an atom-indexed covalent graph with SBU/linker information attached to atoms. The diagnostic uses normalized Gromov hyperbolicity

$$\delta_{\mathrm{norm}}(G) = \frac{\delta(G)}{\mathrm{diam}(G)}, \tag{16}$$

where $\delta(G)$ is the four-point hyperbolicity constant of the shortest-path metric. Lower values indicate closer agreement with a tree metric under this diagnostic; they do not establish that a real MOF graph is a tree. We evaluated 1,377 MOFs and compared each real graph with a degree-matched rewired control. The kNN-16 and spanning-tree graphs provide additional reference points. Tables 7–9 in the main text report the full diagnostic results.

The paired real-minus-rewired difference is -0.0278, with a bootstrap 95% confidence interval of $[-0.0313, -0.0221]$. The real graph has the lower value in 63.7% of pairs; a paired Wilcoxon test gives $p = 1.3 \times 10^{-30}$. This is a relative difference from matched controls, not evidence for a universal hyperbolicity threshold for stable MOFs.

These embeddings came from encoders with different capacity and edge-feature designs; Table 8 is a descriptive geometric comparison rather than an isolated test of manifold choice.

Within this diagnostic collection, the aggregate MAE reductions relative to the atom-plus-block baseline are 27.2% for CH_4 and 31.4% for CO_2. These subset-specific reductions are distinct from the primary ablation in Table 6. The association with per-structure gain is small and task-dependent. We use normalized hyperbolicity as a structural diagnostic, while chemical and pore descriptors may provide additional information for explaining individual prediction errors.

A.2 MOF-SPECIFIC COMPARISON: SCOPE

Table 12: MOF-specific comparison under matched training from scratch, without external pretraining. Values are MAE; column headings give pressures in bar. This supplementary comparison is reported separately from the primary benchmark; comparisons are made within the table.

	CH_4			CO_2		
Model	0.05	0.9	4.5	0.01	0.1	2.5
MSAIGNN	0.024	0.139	0.267	0.050	0.163	0.488
MOFNet	0.025	0.128	0.238	0.058	0.169	0.484
Uni-MOF	0.052	0.382	0.561	0.107	0.452	1.219
GeoDuet	**0.018**	**0.119**	**0.223**	**0.049**	**0.152**	**0.471**

Table 12 reports the supplementary comparison with MSAIGNN, MOFNet, and Uni-MOF under a matched from-scratch setting. No external pretraining is used. This comparison cannot characterize Uni-MOF's published pretrained performance (Wang et al., 2024). Hierarchical and multiscale MOF modeling predates this work (Chen et al., 2022a; Li et al., 2025); the contribution tested here is the coupled representation.

The primary benchmark and the supplementary from-scratch comparison are reported as separate evaluations. In particular, the GeoDuet entries for CH_4 at 4.5 bar are 0.164 and 0.223, respectively. Relative improvements are therefore computed against the baselines in the same table; the two GeoDuet entries are not combined into a single estimate. Across these six supplementary conditions, GeoDuet has the lowest MAE, with reductions of 2.0% to 25.0% relative to the strongest competing entry in each column. These improvements describe the supplementary evaluation only.

A.3 Repeated-run adsorption comparison

The three paired runs are reported in Table 5 in the main text.

GeoDuet has lower MAE in all three runs. To distinguish two common aggregation conventions, let b_s and m_s be the ComFormer and GeoDuet errors in run s. Using the displayed run-level values,

$$1 - \frac{\sum_s m_s}{\sum_s b_s} = 1 - \frac{0.469}{0.521} = 9.98\%, \tag{17}$$

$$\frac{1}{3}\sum_s \left(1 - \frac{m_s}{b_s}\right) = 10.01\%. \tag{18}$$

The first is the relative reduction of the mean MAE; the second is the mean of per-run relative reductions. Computing from the rounded summary means, $1 - 0.156/0.174$, instead gives 10.34% and loses precision. We use the first convention. Standard deviations in the table are consistent with sample standard deviations of the displayed runs. The distinction prevents rounding of the summary means from changing the reported improvement. Three runs do not support a significance claim for every benchmark condition.

A.4 Accessible surface area

Table 10 reports the complete MAE and R^2 comparison in the main text.

The paired bootstrap MAE reduction is 11.3 m^2/g (95% interval [2.8, 19.7]) relative to the atom-plus-block baseline and 168.6 m^2/g (95% interval [151.0, 186.5]) relative to iComFormer. The task remains a MOF pore-related property, so it does not establish universal transfer across crystal properties.

A.5 Computational cost

Table 11 reports parameters, batch and epoch times, and peak memory in the main text.

Relative to iComFormer, GeoDuet takes 4.5% longer per batch and uses 35.4% less peak memory in this setup. The TF32 row measures throughput under a different arithmetic setting and does not establish equivalent predictive accuracy. Total pipeline cost additionally depends on one-time building-block extraction and hardware configuration.

B Geometric definitions and operations

B.1 Hyperbolicity as a diagnostic

Let d_G be the shortest-path metric of a finite connected graph. For four vertices a, b, c, d, order the three sums

$$\{d_G(a,b) + d_G(c,d),\ d_G(a,c) + d_G(b,d),\ d_G(a,d) + d_G(b,c)\} \tag{19}$$

as $s_1 \le s_2 \le s_3$. Under the four-point convention,

$$\delta(G) = \frac{1}{2}\max_{a,b,c,d}(s_3 - s_2), \qquad \delta_{\mathrm{norm}}(G) = \frac{\delta(G)}{\mathrm{diam}(G)}. \tag{20}$$

The normalized quantity is defined for positive finite diameter. The maximum ranges over all vertex quadruples; restricting the maximum to a sampled set gives a lower bound on the exact constant (Fournier et al., 2015; Cohen et al., 2015). This graph statistic is distinct from the curvature parameter K of the learned representation.

The diagnostic supplies conditional motivation: if a graph metric exhibits tree-like structure, results on hyperbolic tree embeddings provide a reason to consider negatively curved representations (Sarkar, 2011). We do not assert that an arbitrary periodic MOF graph satisfies the assumptions of a tree-embedding theorem. Nor do we infer supervised generalization bounds from its measured hyperbolicity.

B.2 Exponential and logarithmic maps

The tangent space at $\mathbf{x} \in \mathbb{H}^{d,K}$ is

$$T_{\mathbf{x}}\mathbb{H}^{d,K} = \{\mathbf{v} \in \mathbb{R}^{d+1} : \langle \mathbf{x}, \mathbf{v} \rangle_{\mathcal{L}} = 0\}. \tag{21}$$

For a nonzero tangent vector, let $\|\mathbf{v}\|_{\mathcal{L}} = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle_{\mathcal{L}}}$. Then

$$\mathrm{Exp}_{\mathbf{x}}(\mathbf{v}) = \cosh\left(\frac{\|\mathbf{v}\|_{\mathcal{L}}}{\sqrt{K}}\right)\mathbf{x} + \sqrt{K}\sinh\left(\frac{\|\mathbf{v}\|_{\mathcal{L}}}{\sqrt{K}}\right)\frac{\mathbf{v}}{\|\mathbf{v}\|_{\mathcal{L}}}. \tag{22}$$

For distinct $\mathbf{x}, \mathbf{y} \in \mathbb{H}^{d,K}$, define $\alpha = -\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}/K \geq 1$. The inverse map is

$$\mathrm{Log}_{\mathbf{x}}(\mathbf{y}) = \frac{\mathrm{arcosh}(\alpha)}{\sqrt{\alpha^2 - 1}}(\mathbf{y} - \alpha\mathbf{x}). \tag{23}$$

The continuous extensions are $\mathrm{Exp}_{\mathbf{x}}(\mathbf{0}) = \mathbf{x}$ and $\mathrm{Log}_{\mathbf{x}}(\mathbf{x}) = \mathbf{0}$. At the origin, the time coordinate of a tangent vector is zero, so its spatial coordinates form the Euclidean vector used for fusion. Finite-precision implementations must handle the limiting cases and the domain of arcosh explicitly.

B.3 Concatenation, normalization, and pooling

For N points $\mathbf{x}_r = (t_r, \mathbf{s}_r) \in \mathbb{H}^{d_r,K}$ with the same K, concatenation is

$$\bigoplus_{\mathcal{L},r=1}^{N} \mathbf{x}_r = \left(\sqrt{\sum_{r=1}^{N} t_r^2 - (N-1)K}, \mathbf{s}_1 \| \cdots \| \mathbf{s}_N\right). \tag{24}$$

Indeed, $t_r^2 = K + \|\mathbf{s}_r\|_2^2$, so the squared time coordinate is $K + \sum_r \|\mathbf{s}_r\|_2^2$ and the Lorentz norm is $-K$.

For a *future-directed time-like* vector, i.e., $\mathbf{x}_0 > 0$ and $\langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} < 0$, radial normalization is

$$\Pi_{\mathcal{L}}(\mathbf{x}) = \sqrt{\frac{K}{-\langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}}}}\,\mathbf{x}. \tag{25}$$

This map cannot project an arbitrary ambient vector onto the upper hyperboloid: multiplying a space-like vector by a scalar preserves the sign of its Lorentz norm. Taking an absolute value in the denominator does not fix that problem. A distinct way to construct a manifold point from arbitrary spatial coordinates $\mathbf{s}$ is $(\sqrt{K + \|\mathbf{s}\|_2^2}, \mathbf{s})$; it should not be confused with radial normalization.

The mean of upper-hyperboloid points is future-directed time-like, so Equation 25 is valid for the pooling operation in Equation 12. To see this, each $t_i > \|\mathbf{s}_i\|_2$, and hence $\sum_i t_i > \|\sum_i \mathbf{s}_i\|_2$. This normalized ambient mean is the pooling operator used in the mathematical description; we do not claim that it is the minimizer of squared geodesic distances.

C Implementation and reproducibility details

C.1 Building-block extraction and atom assignment

Building-block extraction uses MOFstructure[1] and MOFfragmentor[2] to identify chemical units and their constituent atoms. The resulting membership associates block descriptors with atomic nodes. With binary assignment matrix $S \in \{0,1\}^{n \times M}$, $S_{im} = 1$ denotes assignment of atom i to unit B_m. If assignments are disjoint, $\sum_m S_{im} = 1$ and

$$\mathbf{z}_i = \sum_{m=1}^{M} S_{im} F_{\mathcal{B}}(B_m). \tag{26}$$

[1] https://github.com/bafgreat/mofstructure
[2] https://github.com/kjappelbaum/moffragmentor

This expression makes explicit the feature assignment in Equation 4. The matrix S maps descriptors to atoms; it is not an adjacency matrix for a separate block-level GNN. Consequently, adding block information changes the node inputs while preserving the atom-indexed message-passing formulation.

Table 13: Architecture and training configuration of GeoDuet.

Component	Setting
Atomic descriptor	92-dimensional CGCNN features
Block descriptor	SBU/linker information assigned to constituent atoms
Local neighborhoods	hMOF: radius to the 16th nearest neighbor; CoRE: cap of 10 neighbors
Inverse-distance feature	$\eta(r) = c/r$, $c = -0.75$
Orientation	Cosines of edge–lattice angles for all three lattice vectors
Basis expansion	RBF centers from -5 to 0; SBF order 10, 12 radial functions, envelope exponent 5
Encoders	Three message-passing layers; node and edge hidden dimensions 256
Graph readout	Euclidean mean; normalized Lorentz mean
Fusion	Four attention heads
Training	Batch size 64; 100 epochs; RiemannianAdam; one-cycle learning-rate schedule
Hardware	Two NVIDIA H800 GPUs with 80 GB each

C.2 Datasets and evaluation

The source hMOF collection has 137,953 structures (Wilmer et al., 2012); CoRE MOF 2019 contains experimentally reported structures prepared for computation (Chung et al., 2019). The CoRE collection considered here contains 14,142 structures. The reported hMOF pressures are 0.05, 0.5, 0.9, 2.5, and 4.5 bar for CH_4, and 0.01, 0.05, 0.1, 0.5, and 2.5 bar for CO_2. CoRE Ar pressures are 0.001, 0.01, and 0.1 bar. CoRE training targets undergo a $\log(1 + y)$ transform.

For CoRE, MAE is evaluated in the same transformed target space used for training:

$$\mathrm{MAE}_{\log} = \frac{1}{N} \sum_{i=1}^{N} \left| \widehat{y}'_i - \log(1 + y_i) \right|, \tag{27}$$

where y_i is the reference uptake and $\widehat{y}'_i$ is the prediction of its transformed value. The CoRE results in Table 4 use this metric for all compared methods. They are not inverse-transformed to mol kg^{-1}. In contrast, the hMOF adsorption tables report MAE on the original uptake scale. Absolute error magnitudes should therefore be compared within each benchmark and target, while the relative reductions in the main text use the common evaluation scale within each table.

C.3 Poincaré-ball comparison

Table 14: Reported implementation comparison: hMOF MAE in mol kg^{-1}. The two models represent the same hyperbolic geometry in different coordinates.

Implementation	CH_4, 4.5 bar	CO_2, 2.5 bar
Poincaré ball	0.215	0.421
Lorentz (GeoDuet)	**0.164**	**0.330**

The measurements favor the tested Lorentz implementation. The comparison is specific to the evaluated parameterizations and training settings. Because the two models are isometric representations of hyperbolic space, their error difference does not establish a difference in underlying geometric expressivity or identify a unique optimization mechanism.

D SCOPE AND LIMITATIONS

The empirical scope is MOF adsorption and accessible surface area. Repeated training covers one adsorption condition, and the geometry diagnostics do not match encoder capacities or edge features, so they do not isolate curvature as the cause of prediction gains. The from-scratch comparisons do not evaluate published pretrained configurations. Evaluation does not cover unseen topologies, new building units, framework flexibility, or extraction failures. Training-cost measurements exclude building-block preprocessing, and offline prediction accuracy does not establish prospective screening efficiency or synthesis success.