
WindMiL: Equivariant Graph Learning for Wind Loading Prediction

Themistoklis Vargiemezis¹
tvarg@stanford.edu

Charilaos Kanatsoulis²
charilaos@cs.stanford.edu

Catherine Gorlé¹
gorle@stanford.edu

¹Department of Civil & Environmental Engineering, Stanford, CA, USA

² Department of Computer Science, Stanford, CA, USA

Abstract

Accurate prediction of wind loading on buildings is crucial for structural safety and sustainable design, yet conventional approaches such as wind tunnel testing and large-eddy simulation (LES) are prohibitively expensive for large-scale exploration. Each LES case typically requires at least 24 hours of computation, making comprehensive parametric studies infeasible. We introduce WINDMiL, a new machine learning framework that combines systematic dataset generation with symmetry-aware graph neural networks (GNNs). First, we introduce a large-scale dataset of wind loads on low-rise buildings by applying signed distance function interpolation to roof geometries and simulating 462 cases with LES across varying shapes and wind directions. Second, we develop a reflection-equivariant GNN that guarantees physically consistent predictions under mirrored geometries. Across interpolation and extrapolation evaluations, WINDMiL achieves high accuracy for both the mean and the standard deviation of surface pressure coefficients (e.g., $\text{RMSE} \leq 0.02$ for mean C_p) and remains accurate under reflected-test evaluation, maintaining hit rates above 96% where the non-equivariant baseline model drops by more than 10%. By pairing a systematic dataset with an equivariant surrogate, WINDMiL enables efficient, scalable, and accurate predictions of wind loads on buildings.

1 Introduction

Wind loading on buildings is crucial for structural engineering and resilient urban planning. Accurate predictions of surface pressure loads and integrated loads guide both safety codes and sustainable building practices. Yet, traditional approaches such as wind tunnel testing and full-scale experiments are costly and time-consuming [Alrawashdeh and Stathopoulos, 2015, Richards and Hoxey, 2012, Vargiemezis and Gorlé, 2024], while high-fidelity CFD simulations, such as large-eddy simulation (LES), though accurate, remain prohibitively expensive for large-scale design exploration [Potsis et al., 2023, Blocken, 2015, Vargiemezis and Gorlé, 2025b,a]. These limitations highlight the need for efficient and accurate surrogates that can operate at scale.

Recent advances in geometric deep learning provide a powerful foundation for such surrogates. Graph neural networks (GNNs) naturally represent the irregular meshes common in CFD, with nodes as mesh points and edges as connectivity. Models such as MeshGraphNets [Pfaff et al., 2020] and Graph Network Simulators [Sanchez-Gonzalez et al., 2020] show the potential of graph-based learning for fluid and structural dynamics, while PolyGNN [Chen et al., 2024] highlights their use in building geometry reconstruction. These works highlight that graph learning can bridge building geometry, flow physics, and predictive modeling.

To fully realize the potential of graph learning in wind engineering, surrogate models must incorporate the physical symmetries that govern flows around buildings. Reflectional and rotational invariances are fundamental for ensuring consistent and reliable predictions. For example, when the flow is aligned with x -direction, y is the vertical, and z is the spanwise, the pressure distribution of a building at $+45^\circ$ wind incidence should be the reflected version of a building at -45° wind incidence with respect to the xy -plane, as shown in Fig. 1.

Equivariant graph learning embeds these symmetry constraints directly into the model, ensuring that its predictions transform consistently under reflections and rotations. Approaches such as group-equivariant CNNs [Cohen and Welling, 2016] and E(3)-equivariant GNNs like NequIP [Batzner et al., 2022] have already demonstrated substantial gains in efficiency and generalization in physical sciences. These advantages make equivariant GNNs a suitable modeling choice for wind loading surrogates.

Progress in this direction has been hindered by the lack of systematic datasets. While benchmark experiments such as the TPU database [Quan et al., 2007] remain invaluable, they cover only a small number of canonical building geometries. By contrast, other areas such as automotive aerodynamics have advanced through parametric dataset generation paired with CFD [Benjamin and Iaccarino, 2025]. No comparable dataset exists for wind loading on low-rise buildings, despite their importance to building safety.

Our contribution: We introduce WINDMIL, a new machine learning paradigm for wind loading prediction that combines systematic datasets with symmetry-aware graph learning. Specifically:

1. **Dataset.** We generate a large-scale dataset of wind loading on low-rise buildings, with systematically varying roof morphologies and wind directions, using LES. This extends databases such as the TPU database through controlled geometric interpolation.
2. **Model.** We develop a reflection-equivariant GNN that respects reflectional invariances, providing physically consistent surrogates. Our approach achieves error reduction by more than 10% on the symmetrical geometries compared to the non-equivariant baseline.

By releasing both the dataset and the model, we establish an ML paradigm for wind engineering that is data-driven, symmetry-aware, and scalable. This paradigm opens new opportunities for accurate, efficient, and physically consistent prediction of wind loads, advancing resilient and sustainable urban design.

2 Dataset generation

2.1 Geometry parameterization

In this work, we adapt the Signed Distance Function (SDF)-based interpolation method, originally developed for automotive geometries by Benjamin and Iaccarino [2025], to systematically generate diverse building shapes for wind loading analysis. We begin with three basis geometries derived from the TPU dataset [Quan et al., 2007]: a flat roof, a gable roof, and a hip roof. All buildings share the same footprint of 18 m length \times 12 m width in full scale, while their heights differ: 8 m for the flat roof, 16 m for the gable roof, and 24 m for the hip roof.

Each geometry is first converted into a binary grid using ray-tracing, where cells are marked as occupied or empty based on ray-surface intersections. The binary grids are then transformed into signed distance functions (SDFs), which assign to each grid point the signed distance to the building surface. By convention, the SDF equals zero on the surface, is positive outside, and negative inside. This

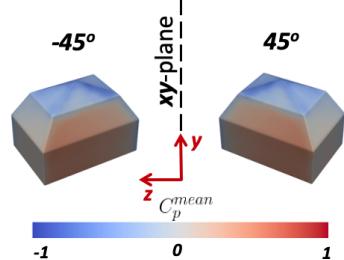


Figure 1: Contour plots of mean C_p at -45° and -45° wind incidence.

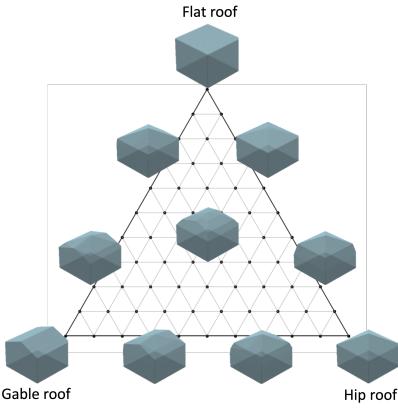


Figure 2: Convex hull of the three basis geometries. Points inside the convex hull correspond to interpolated buildings.

continuous representation enables smooth and consistent interpolation of shapes. To generate new buildings, we construct a convex hull in the SDF space spanned by the three basis geometries, as shown in Fig. 2. Barycentric interpolation within this convex hull produces intermediate SDFs that smoothly transition between the flat, gable, and hip roof cases. For the wind incidence at 0° , this yields 66 unique interpolated buildings. To build a complete database, each interpolated building is rotated in 15° increments from 0° to 90° , resulting in $66 \times 7 = 462$ cases, with one convex hull generated for each wind direction.

Formally, interpolation between two SDFs $\phi_1(\vec{x}_i)$ and $\phi_2(\vec{x}_i)$ on a structured grid, where (\vec{x}_i) are spatial coordinates on a discrete grid, can be written as

$$\phi_3((\vec{x}_i)) = \alpha\phi_1((\vec{x}_i)) + (1 - \alpha)\phi_2((\vec{x}_i)), \quad (1)$$

where α is a scalar weight. Varying α yields intermediate geometries between the basis cases. Applying Eq. (1) to a unit sphere and unit cube, one can produce intermediate shapes by varying α , an example of which is shown in Fig. 3

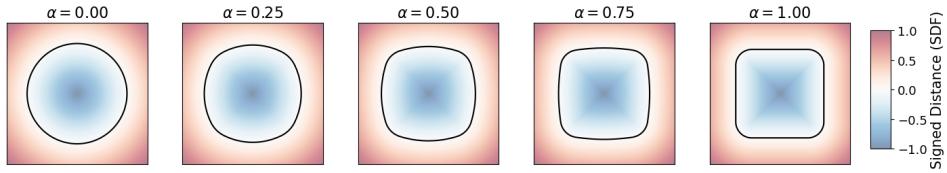


Figure 3: SDF interpolation between a sphere and a cube, per Eq. (1) for different α .

Finally, the interpolated SDFs are reconstructed into surface meshes using the marching cubes algorithm [Benjamin and Iaccarino, 2025], followed by Laplacian smoothing to refine geometry quality. This ensures the resulting buildings are physically suitable for downstream tasks such as CFD simulations or data-driven surrogate modeling. By combining convex-hull interpolation with systematic rotations, the method enables the creation of a large dataset of building geometries.

2.2 Large-eddy Simulations

To generate the dataset, LES simulations are performed using the CharLES code, a low-Mach, isentropic solver developed by Cadence Design Systems [2022]. CharLES implements a finite volume approach with an automated body-fitted meshing technique based on 3D-clipped Voronoi diagrams, and uses numerical schemes with second-order accuracy in space and time to solve the governing equations [Ambo et al., 2020]. In addition, the Vreman turbulence model is used to model the unresolved part of the stress tensor [Vreman, 2004]. CharLES has been extensively validated against experimental data in a range of wind engineering applications. For instance, it has been validated against wind tunnel measurements of wind-induced pressure loads on high- and low-rise buildings [Ciarlatani et al., 2023, Vargiemezis and Gorlé, 2025a,b], as well as against measurements of wind-driven natural ventilation [Hwang and Gorlé, 2022]. Comparisons against field measurements have also been performed for wind pressures on the Space Needle [Hochschild and Gorlé, 2024] and for natural ventilation flow in a dense urban environment [Hwang and Gorlé, 2023].

The mesh uses 2.1×10^6 control volumes, with refinement near the no-slip surfaces of the building, as shown in Fig. 4. The resolution on the building surfaces is $\Delta/H_{ref} = 0.0098$, where Δ is the mesh size locally, and H_{ref} is the characteristic height of the building. The domain extents $5H_{ref}$ upstream, $15H_{ref}$ downstream, and $5H_{ref}$ in the lateral and vertical directions, following the proposed guidelines [Franke et al., 2011]. At the inlet, a logarithmic mean velocity profile is prescribed, and artificially generated turbulent fluctuations are superimposed. The turbulent velocity field is generated using the divergence-free digital filter method proposed by [Kim et al., 2013]. At the outlet, a zero gradient condition is applied, while the two lateral boundaries are periodic, and a slip condition is applied at the top boundary. Finally, at the ground, a rough wall function for a neutral atmospheric boundary layer with a fixed roughness of 0.0027 m is specified. More details regarding the setup can be found in the study of a different isolated building, since the same setup is used [Vargiemezis and Gorlé, 2025b].

The simulation is run for a total time of 20 seconds, and statistics are collected for the last 15 seconds, after the initial burn-in period of 5 seconds. The total time for one simulation on 64 CPU's requires 24 hours. The quantities of interest (QoIs) of the simulation are the time-averaged (mean) and the

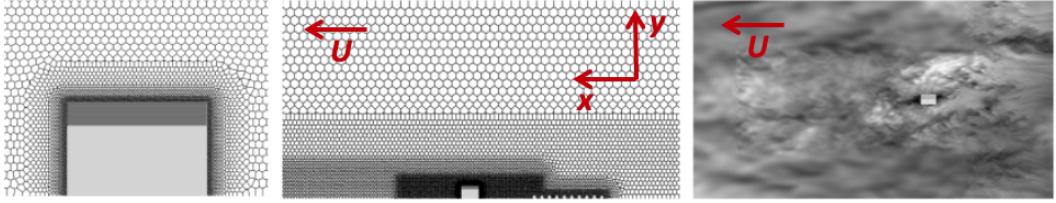


Figure 4: View of the mesh used for the large-eddy simulations, with side views of refinement zones shown: (left) near the building and (middle) in the surrounding area. (right) Top view of instantaneous velocity magnitude contours showing the turbulent flow around the building.

standard deviation (std) of the pressure coefficient. The pressure coefficient at a surface node i is defined as

$$C_{p,i} = \frac{p_i - p_\infty}{\frac{1}{2}\rho U_\infty^2}, \quad (2)$$

where p_i is the local surface pressure at location i on the building surface, p_∞ is the freestream pressure, ρ is the density of air, and U_∞ is the reference velocity.

3 Model architecture

We represent each building surface as a point cloud, from which we construct a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where nodes $v \in \mathcal{V}$ correspond to sampled surface points and edges $e \in \mathcal{E}$ are defined based on spatial proximity. Each node is associated with a set of geometric features $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$. In particular, each node v has features $\mathbf{x}_v \in \mathbb{R}^d$, with $d = 6$: normalized coordinates (x, y, z) with respect to building height H_{ref} and surface normal components (n_x, n_y, n_z) , which are unit vectors oriented outward of building surface. The learning task is to predict the per-node mean or std pressure coefficients C_p . To process the graph structure and the geometrical features of each building, we employ a message-passing GNN $f_\theta(\mathcal{G}, \mathbf{X})$, [Kipf and Welling, 2016, Xu et al., 2019, Hamilton et al., 2017], defined by the following recursive formula:

$$\mathbf{x}_v^{(l)} = h^{(l-1)} \left(\mathbf{x}_v^{(l-1)}, g^{(l-1)} \left(\left\{ \mathbf{x}_u^{(l-1)} : u \in \mathcal{N}(v) \right\} \right) \right), \quad (3)$$

where $\mathcal{N}(v)$ represents the 1-hop neighborhood of vertex v , and g, h are permutation equivariant operators. In our implementation, $f_\theta(\mathcal{G}, \mathbf{X})$ is modeled as a stack of GraphSAGE layers [Hamilton et al., 2017] with residual connections, layer normalization, and a global skip from the input projection.

Wind loading exhibits a reflection symmetry with respect to the horizontal (xy) plane, as shown in Fig. 1. For instance, the surface pressure distribution at a wind incidence angle of $+45^\circ$ is the reflected version of that at -45° , when the reflected building is taken across the xy plane. To encode this invariance, we define a reflection operator \mathcal{R} that flips the vertical coordinate and normal component:

$$\mathcal{R} : (z, n_z) \mapsto (-z, -n_z),$$

while leaving (x, y, n_x, n_y) unchanged. The model is designed such that predictions are invariant under this transformation. To enforce symmetry, both the original and reflected features are passed through the same encoder f_θ . The resulting embeddings are then averaged:

$$\mathbf{z}_v = \frac{1}{2} \left(f_\theta(\mathcal{G}, \mathbf{X})[v] + f_\theta(\mathcal{G}, \mathcal{R}(\mathbf{X}))[v] \right).$$

The overall model $F(\cdot)$, which maps inputs to predictions, then satisfies the reflection-equivariant property $F(\mathcal{R}(\mathbf{x})) = F(\mathbf{x})$. Finally, embeddings \mathbf{z}_v are passed through a feed-forward predictor g_ϕ , yielding the node-level outputs $C_{p,v} = g_\phi(\mathbf{z}_v)$. This architecture ensures predictions that respect the underlying physical symmetry of wind loading on buildings.

4 Experiments

We compare the reflection-equivariant GNN, WINDMIL, with the baseline model GRAPH SAGE, which does not incorporate geometric symmetries. Both models are evaluated in two cases: (a)

interpolation and (b) extrapolation in the shape space, see also Fig. 5. Interpolation corresponds to a random train/dev/test split of the points in the convex hull, while extrapolation is assessed on unseen building configurations. In extrapolation, geometries located on the convex hull boundary of the shape space are left out for testing, while the remaining points are randomly divided into train/dev.

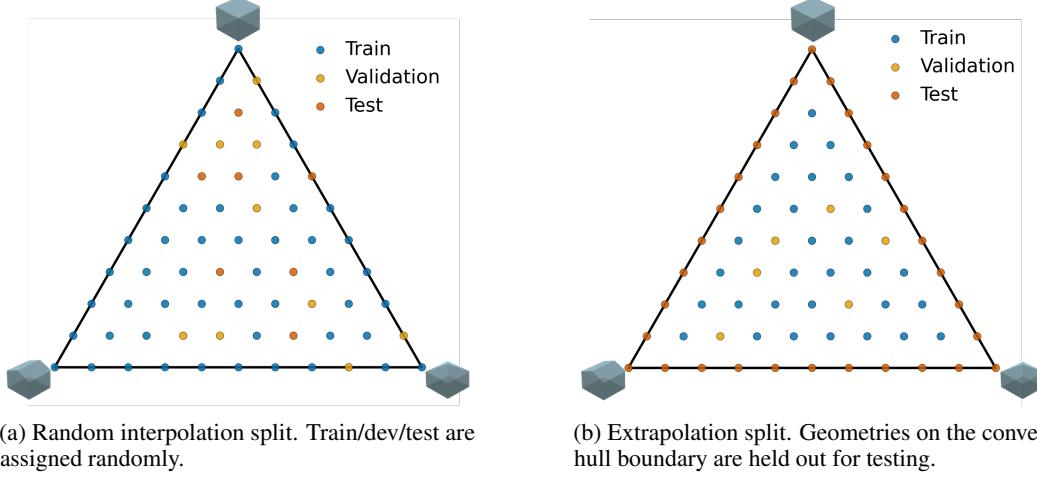


Figure 5: Visualization of dataset splits in the barycentric shape space. The convex hull is spanned by the three basis roof geometries (flat, gable, hip), and each point corresponds to one interpolated building.

For both split methods, we report results with and without including reflected data in the test set. It is important to note that training is always performed using only the original (non-reflected) geometries. In the “+Sym” evaluation setting (see also Table 1), we extend the test set to include the reflections of the held-out test geometries. This setup evaluates whether the models generalize correctly to symmetrical geometries, without having seen them during training. We evaluate the predictive accuracy using five metrics: root mean square error (RMSE), mean squared error (MSE), mean absolute error (MAE), coefficient of determination (R^2), and hitrate (%). The hitrate is defined as the percentage of predictions within a tolerance of ± 0.10 for mean C_p and ± 0.05 for std C_p , which correspond to the 10% of their maximum values observed on the buildings.

4.1 Interpolation performance

Table 1 summarizes the interpolation results. When evaluated on the non-reflected test set, both GRAPHsAGE and WINDMiL achieve high predictive accuracy for the mean and std C_p . Errors remain small, with RMSEs below 0.02 for mean C_p and 0.003 for std C_p , and hit rates exceeding 95% and 97%, respectively. This confirms that both models successfully capture the pressure variations within the interpolation regime.

	Method	RMSE	MSE	MAE	R^2	Hitrate (%)
Mean C_p	GRAPHsAGE	0.019	0.211	0.032	0.980	95.95
	GRAPHsAGE + Sym	0.055	0.273	0.053	0.943	85.41
	WINDMiL	0.020	0.213	0.033	0.979	95.80
	WINDMiL + Sym	0.019	0.207	0.031	0.981	96.53
Std C_p	GRAPHsAGE	0.003	0.130	0.012	0.958	97.98
	GRAPHsAGE + Sym	0.008	0.169	0.020	0.882	91.63
	WINDMiL	0.003	0.137	0.013	0.949	97.62
	WINDMiL + Sym	0.003	0.135	0.013	0.953	97.86

Table 1: Interpolation performance on mean and std C_p . +Sym indicates that the test set additionally includes the reflections of the test set geometries.

When the test set is extended with reflected geometries (+Sym), however, clear differences are observed. The baseline GRAPHsAGE shows a substantial drop in performance, with the hitrate decreasing from 95.95% to 85.41% for mean C_p and from 97.98% to 91.63% for std C_p . This

indicates that the non-equivariant model does not generalize consistently to symmetry-transformed cases and effectively treats reflections as unseen geometries. In contrast, WINDMIL maintains high accuracy under reflection, achieving 96.53% and 97.86% hitrates for mean and std C_p , respectively, with similar RMSE and R^2 values to the original test set. These results show that explicitly encoding reflection symmetry not only enforces physically consistent predictions but also improves robustness to unseen symmetric transformations.

4.2 Extrapolation performance

Table 2 reports results for extrapolation, where building geometries on the convex hull boundary of the shape space are held out for testing. As expected, errors increase compared to interpolation since the models are evaluated on unseen geometries outside the training distribution. For the non-reflected test set, RMSE increases to ~ 0.05 for mean C_p and ~ 0.019 for std C_p , yet both GRAPHsAGE and WINDMIL maintain high predictive accuracy with R^2 values between 0.94–0.98 and hitrates above 94% and 97% for mean and std C_p , respectively. These results indicate that both models preserve their accuracy in the more challenging extrapolation task.

	Method	RMSE	MSE	MAE	R^2	Hitrates (%)
Mean C_p	GRAPHsAGE	0.047	0.227	0.033	0.976	94.98
	GRAPHsAGE + Sym	0.076	0.581	0.054	0.939	84.67
	WINDMIL	0.050	0.251	0.035	0.974	94.16
	WINDMIL + Sym	0.049	0.242	0.034	0.974	94.69
Std C_p	GRAPHsAGE	0.019	0.0365	0.012	0.944	97.35
	GRAPHsAGE + Sym	0.029	0.086	0.020	0.870	91.57
	WINDMIL	0.019	0.037	0.013	0.942	97.30
	WINDMIL + Sym	0.019	0.037	0.012	0.944	97.33

Table 2: Extrapolation performance on mean and std C_p . +Sym indicates that the test set additionally includes the reflections of the test set geometries.

When the test set is extended with reflected geometries (+Sym), differences between the models are observed. The baseline GRAPHsAGE shows a clear decrease in performance, with hitrates dropping from 94.98% to 84.67% for mean C_p and from 97.35% to 91.57% for std C_p , along with corresponding reductions in R^2 . In contrast, WINDMIL remains accurate under symmetry transformations, maintaining hit rates of 94.69% for mean C_p and 97.33% for std C_p , with nearly unchanged RMSE values. This consistency highlights the benefit of explicit reflection equivariance; the model generalizes more accurately to unseen symmetric configurations, avoiding the errors observed in the non-equivariant baseline.

4.3 Qualitative Analysis

To further assess predictive performance, we present contour plots of mean and std C_p on the building surfaces for the interpolation and extrapolation cases in Fig. 6 and Fig. 7, respectively. Predictions from the baseline GRAPHsAGE and our proposed WINDMIL are compared against the LES targets.

For the standard interpolation test set in Fig. 6, both models capture the main spatial patterns of C_p . They correctly reproduce high- and low-pressure regions across roof and wall surfaces. The close visual agreement with LES is consistent with the quantitative metrics in Table 1, where both models achieve low errors and hitrates above 95%. When the reflected geometries +Sym are included, differences become more visible; the GRAPHsAGE tends to overpredict both mean and std C_p on roof surfaces, while WINDMIL maintains close agreement with the LES ground truth.

For the extrapolation split in Fig. 7, both models remain accurate for the mean C_p and std C_p , which is confirmed with the quantitative analysis of Table 2, where both models achieved hitrates of more than 94%. When the +Sym extrapolation is considered, both models remain accurate for mean C_p , but differences are visible for std C_p . The GRAPHsAGE significantly underpredicts the variability at the building corners, specifically where the ground truth indicates high std values. WINDMIL, by comparison, captures these regions much more closely. These qualitative observations show that that reflection symmetry is particularly important for predicting higher-order statistics of pressure loads.

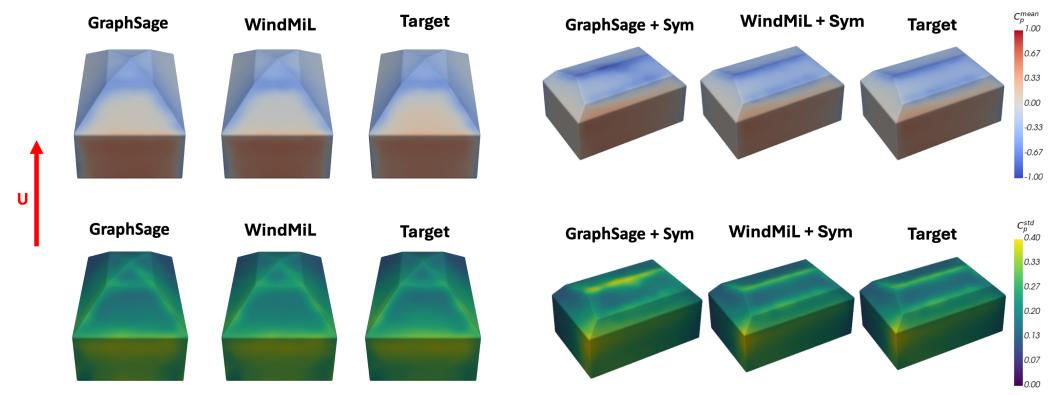


Figure 6: Contour plots of mean and std C_p for the interpolation split. Columns show GRAPHsAGE, WINDMiL, and LES ground truth; rows correspond to mean C_p and std C_p .

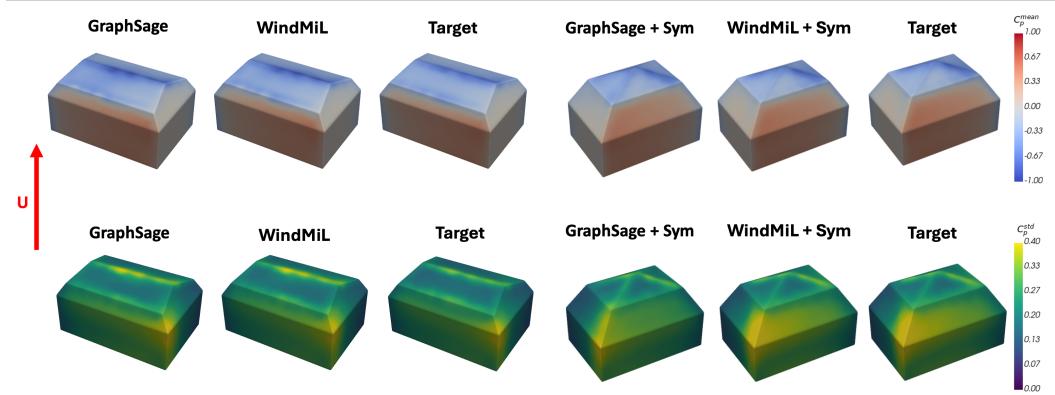


Figure 7: Contour plots of mean and std C_p for the interpolation split. Columns show GRAPHsAGE, WINDMiL, and LES ground truth; rows correspond to mean C_p and std C_p .

5 Conclusions

We introduced WINDMiL, a symmetry-aware graph learning framework for predicting wind loading on low-rise buildings, together with a systematically generated LES dataset of 462 buildings with varying roof types and wind directions. The dataset is built by interpolating between canonical roof geometries using signed distance functions and running high-fidelity LES for each configuration. A single LES simulation requires on the order of 24 hours on 64 CPU cores. On top of this dataset, we proposed a reflection-equivariant GNN surrogate that encodes the physical symmetry, such that pressures at $+45^\circ$ wind incidence should mirror those at -45° . WINDMiL achieves high accuracy for both mean and standard deviation of the surface pressure coefficient C_p in both interpolation and extrapolation settings with hitrates above 95%. WINDMiL achieves similar accuracy when the reflected (+Sym) geometries are included in the test set, while the standard GRAPHsAGE degrades by more than 10 percentage points in hitrate. These results show that explicitly enforcing symmetry in the model architecture improves physical consistency without sacrificing accuracy, and enables reliable surrogate predictions at a fraction of LES cost. Future work will extend this framework to multi-building urban areas, additional inflow conditions, and the prediction of integrated loads for structural design.

References

Hatem Alrawashdeh and Ted Stathopoulos. Wind pressures on large roofs of low buildings and wind codes and standards. *Journal of Wind Engineering and Industrial Aerodynamics*, 147:212–225, 2015.

Kei Ambo, Hiroaki Nagaoka, David A Philips, Chris Ivey, Guillaume A Brès, and Sanjeeb T Bose. Aerodynamic force prediction of the laminar to turbulent flow transition around the front bumper of the vehicle using dynamic-slip wall model les. In *AIAA Scitech 2020 Forum*, page 0036, 2020.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):2453, 2022.

Mark Benjamin and Gianluca Iaccarino. A systematic dataset generation technique applied to data-driven automotive aerodynamics. *APL Machine Learning*, 3(1), 2025.

Bert Blocken. Computational fluid dynamics for urban physics: Importance, scales, possibilities, limitations and ten tips and tricks towards accurate and reliable simulations. *Building and Environment*, 91:219–245, 2015.

Cadence Design Systems. Fidelity charles, 2022. URL <https://community.cadence.com/>. Accessed: April 14, 2025.

Zhaiyu Chen, Yilei Shi, Liangliang Nan, Zhitong Xiong, and Xiao Xiang Zhu. Polygnn: Polyhedron-based graph neural network for 3d building reconstruction from point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 218:693–706, 2024.

Mattia Fabrizio Ciarlatani, Zhu Huang, David Philips, and Catherine Gorlé. Investigation of peak wind loading on a high-rise building in the atmospheric boundary layer using large-eddy simulations. *Journal of Wind Engineering and Industrial Aerodynamics*, 236:105408, 2023.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.

Jorg Franke, Antti Hellsten, K Heinke Schlunzen, and Bertrand Carissimo. The cost 732 best practice guideline for cfd simulation of flows in the urban environment: a summary. *International Journal of Environment and Pollution*, 44(1-4):419–427, 2011.

William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017.

John Hochschild and Catherine Gorlé. Comparison of measured and les-predicted wind pressures on the space needle. *Journal of Wind Engineering and Industrial Aerodynamics*, 249:105749, 2024. ISSN 0167-6105. doi: <https://doi.org/10.1016/j.jweia.2024.105749>. URL <https://www.sciencedirect.com/science/article/pii/S0167610524001120>.

Yunjae Hwang and Catherine Gorlé. Large-eddy simulations of wind-driven cross ventilation, part1: validation and sensitivity study. *arXiv preprint arXiv:2204.00786*, 2022.

Yunjae Hwang and Catherine Gorlé. Large-eddy simulations to define building-specific similarity relationships for natural ventilation flow rates. *Flow*, 3:E10, 2023.

Yusik Kim, Ian P Castro, and Zheng-Tong Xie. Divergence-free turbulence inflow conditions for large-eddy simulations with incompressible flow solvers. *Computers & Fluids*, 84:56–68, 2013.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International conference on learning representations*, 2020.

Theodore Potsis, Yoshihide Tominaga, and Ted Stathopoulos. Computational wind engineering: 30 years of research progress in building structures and environment. *Journal of Wind Engineering and Industrial Aerodynamics*, 234:105346, 2023.

Yinggeer Quan, Yukio Tamura, M. Matsui, S.Y. Cao, and A. Yoshida. Tpu aerodynamic database for low-rise buildings. *Proceedings of 12th International Conference on Wind Engineering*, pages 1615–1622, 01 2007.

PJ Richards and RP Hoxey. Pressures on a cubic building—part 1: Full-scale results. *Journal of Wind Engineering and Industrial Aerodynamics*, 102:72–86, 2012.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.

Themistoklis Vargiemezis and Catherine Gorlé. Wind tunnel experiments on realistic low-rise building in urban area, 2024. URL <https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-2571/#detail-4070145275427230191-242ac118-0001-012/?version=2>.

Themistoklis Vargiemezis and Catherine Gorlé. Predicting wind-induced interference effects on a low-rise building in a realistic urban area using large-eddy simulations. *Journal of Wind Engineering and Industrial Aerodynamics*, 2025a.

Themistoklis Vargiemezis and Catherine Gorlé. A predictive large-eddy simulation framework for the analysis of wind loads on a realistic low-rise building geometry. *Journal of Wind Engineering and Industrial Aerodynamics*, 256:105950, 2025b.

AW Vreman. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of fluids*, 16(10):3670–3681, 2004.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6ia5Km>.