

A hierarchical deep learning model for predicting pedestrian-level urban winds

Reda Snaiki ^a, Jiachen Lu ^{b,c}, Shaopeng Li ^d, Negin Nazarian ^{b,c,e,*}

^a Department of Construction Engineering, École de Technologie Supérieure, Université du Québec, Montréal, Québec, Canada

^b School of Built Environment, UNSW Sydney, Sydney, NSW 2052, Australia

^c ARC Centre of Excellence for Climate Extremes, University of New South Wales, Sydney, Australia

^d Department of Civil Engineering and Environmental Engineering, University of Louisiana at Lafayette, LA, USA

^e ARC Centre of Excellence for the 21st Century Weather, UNSW Sydney, Sydney, Australia

ARTICLE INFO

Keywords:

Urban wind flow
Urban climate modeling
Image-to-image translation
Hierarchical model
Deep learning

ABSTRACT

Deep learning-based surrogate models offer a computationally efficient alternative to high-fidelity computational fluid dynamics (CFD) simulations for predicting urban wind flow. However, conventional approaches usually only yield low-frequency predictions (essentially averaging values from proximate pixels), missing critical high-frequency details such as sharp gradients and peak wind speeds. This study proposes a hierarchical approach for accurately predicting pedestrian-level urban winds, which adopts a two-stage predictor-refiner framework. In the first stage, a U-Net architecture generates a baseline prediction from urban geometry. In the second stage, a conditional Generative Adversarial Network (cGAN) refines this baseline by restoring the missing high-frequency content. The cGAN's generator incorporates a multi-scale architecture with stepwise kernel sizes, enabling simultaneous learning of global flow structures and fine-grained local features. Trained and validated on the UrbanTALES dataset with comprehensive urban configurations, the proposed hierarchical model significantly outperforms the baseline predictor. An ablation study against a single-stage cGAN (mapping geometry directly to LES fields) and a hierarchical CNN refiner (trained only with L1 loss) shows that both the two-stage design and the adversarial objective are necessary to achieve the reported accuracy gains. With a marked qualitative improvement in resolving high-speed wind jets and complex turbulent wakes as well as wind statistics, the results yield quantitative enhancement in prediction accuracy (e.g., RMSE reduced by 76 % for the training set and 60 % for the validation set). This work presents an effective and robust methodology for enhancing the prediction fidelity of surrogate models in urban planning, pedestrian comfort assessment, and wind safety analysis. The proposed model will be integrated into an interactive web platform as Feilian Version 2.

1. Introduction

The accurate characterization of wind flow within the urban canopy is a longstanding challenge in environmental engineering and city planning, with direct implications for pedestrian comfort, pollutant dispersion, and the structural safety of buildings and urban infrastructures [1,2]. Historically, three primary methodologies have been employed to study this complex phenomenon: in-situ measurements, wind tunnel and reduced-scale experiments, and computational fluid dynamics (CFD) simulations [3–5]. In-situ measurements, which rely on data from fixed sensors such as anemometers, ultrasonic transducers, and local weather stations, provide highly accurate ground-truth data [6,7]. However, these point-based measurements are often too sparse to

resolve the intricate spatial variability of urban wind. Wind tunnel experiments offer a controlled environment for studying flow around scaled models but can be subject to scaling uncertainties and limitations in replicating real-world atmospheric conditions [8–10]. In recent decades, CFD simulations [11–15], particularly high-fidelity methods like Large-eddy Simulation (LES) [16], have become the state-of-the-art for generating comprehensive, three-dimensional datasets of urban wind flow ([17–19] and [20,21]). Despite their accuracy, the immense computational cost of these physics-based simulations, which often require hours or days of high-performance computing for a single scenario, renders them impractical for applications requiring rapid design iteration or large-scale analysis.

The computational bottleneck of CFD has motivated the

* Corresponding author.

E-mail address: n.nazarian@unsw.edu.au (N. Nazarian).

<https://doi.org/10.1016/j.buildenv.2026.114354>

Received 30 October 2025; Received in revised form 1 February 2026; Accepted 10 February 2026

Available online 10 February 2026

0360-1323/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

development of machine learning-based surrogate models that emulate the input-output mapping of physics-based solvers at a fraction of the cost [22–25]. Trained on large datasets generated from high-fidelity CFD simulations, these models have evolved rapidly, progressing from foundational methods like k-Nearest Neighbors (kNN) [1] to more complex deep learning frameworks [25–31]. Convolutional Neural Networks (CNNs), particularly the U-Net architecture, have since become standard for framing the task as an image-to-image translation problem, mapping urban geometry to wind field predictions [32,33]. More advanced paradigms include Fourier Neural Operators (FNOs), which learn resolution-invariant operators in the frequency domain for greater efficiency and generalizability [34–36], and Physics-Informed Neural Networks (PINNs) [37,38], which embed the governing Navier-Stokes equations directly into the learning process to enforce physical consistency and reduce reliance on large datasets. Furthermore, Graph Neural Networks (GNNs) are an emerging approach well-suited to the irregular topology of urban environments, having demonstrated performance comparable to CNNs in related forecasting tasks [39]. Particularly, in the field of urban wind prediction, Liu et al. [40] utilized a memory-scalable graph neural network, which can maintain consistent GPU memory usage as the simulation domain increases. Shao et al., [41] developed a physics-informed graph neural network that can satisfy the governing RANS equations in the CFD-based training data. Milla-Val et al. [42] used conditional generative adversarial network to predict detailed local wind fields with CFD-level resolution using the input of coarser numerical weather prediction data. Zhao et al., [43] trained graph neural networks to directly map steady RANS results to time-averaged LES targets. Calzolari and Liu [44] leveraged physics-constrained conditional generative adversarial networks to learn the mapping from early-stage flow fields to their statistically steady-state counterparts, which also showed the enhanced performance of graph-based model compared to convolutional models. Clarke et al. [26] introduced a multi-layer perceptron (MLP)-mixer architecture to embed critical 3D building data via 2D input channels, which outperformed conventional U-net models. More comprehensive analysis on existing literature can be found in the two review papers [45,46]. These advances have enabled prediction speeds several orders of magnitude faster than traditional solvers while retaining strong agreement with CFD-generated ground truth. Nonetheless, a persistent limitation remains [47]: pixel-wise regression models tend to produce overly smooth predictions, compressing the wind speed distribution to its mode over idealized urban arrays and even erasing local extrema over realistic urban neighborhoods, which are all critical for assessing wind extremes in the built environment.

To address this research gap, this paper introduces a novel two-stage, hierarchical predictor-refiner model that significantly enhances the fidelity of surrogate model predictions. The first stage employs a U-Net-based predictor to generate a stable, low-frequency baseline from urban geometry. The second stage applies a conditional Generative Adversarial Network (cGAN) as a refiner to restore high-frequency details absent in the baseline. The cGAN generator integrates a multi-scale architecture with stepwise kernel sizes, enabling it to capture both large-scale flow structures and small-scale turbulent features and wind extremes. The proposed hierarchical model is trained and validated on the high-resolution UrbanTALES dataset [48], which covers neutral simulations over both idealized building arrays and realistic urban neighborhoods. Its performance is evaluated relative to the original U-Net model of Lu et al., [47] and further benchmarked against two additional baselines: a single-stage cGAN that maps geometry directly to LES, and a hierarchical CNN refiner that is trained only with L1 loss. The proposed model, named Feilian Version 2, will be integrated into the interactive web platform at <http://feilian.climate-resilientcities.com/> as an improved successor to the U-Net-based Feilian Version 1.

2. Methodology

This section details the proposed methodology for enhancing surrogate model predictions of urban wind flow. The goal is to develop a hierarchical deep learning model capable of accurately and efficiently predicting the 2D pedestrian-level wind speed field (at a height of 1.5 m). This framework is designed to first generate a robust baseline prediction [47] and subsequently refine it using a generative adversarial network to emulate the high-fidelity detail of LES. The following subsections describe the data source to define the problem, the overall two-stage framework, the architecture of each component, and the final training protocol.

2.1. Data source and problem definition

The dataset used to train and validate the proposed framework is UrbanTALES, an extensive collection of 512 high-resolution Large-eddy Simulations (LES) of pedestrian-level wind. These simulations provide the high-fidelity ground-truth data for the surrogate modeling task. The dataset includes cases ranging from idealized building arrays to realistic neighborhoods extracted from OpenStreetMap [49], spanning diverse urban densities, height distributions, and wind orientations.

To contextualize the CFD ground truth used in this study, a brief overview of the underlying LES configuration is provided here. The UrbanTALES dataset [48] is derived from high-fidelity LES using the Parallelized LES Model (PALM) version r4554, as detailed in Lu et al. [47]. Key simulation parameters include a 1.5-order turbulence closure with a dynamic regularized gradient subgrid-scale model, Arakawa staggered C-grid for vertical discretization, second-order central differences for horizontal discretization, and a minimal storage scheme for time integration. The pressure solver employs the Fastest Fourier Transform in the West (FFTW), with periodic lateral boundary conditions, a Neumann surface boundary following the logarithmic law (roughness length $z_0 = 0.01$ m), and a Dirichlet (no-slip) top boundary. Topography is represented via 2D building height arrays in ASCII format. The domain height is fixed at ~ 130 m (5–8 times the mean building height) to ensure accurate resolution of above-canopy flows, with grid resolutions of 1 m for realistic configurations and 0.5 m for idealized ones, balancing computational efficiency and convergence of turbulent statistics. This setup builds on prior validations against direct numerical simulations [50], and wind tunnel experiments [51], demonstrating reliable capture of momentum and energy exchanges in neutral urban canopy flows.

The problem is framed as an image-to-image translation task. The primary input for the overall framework is the urban geometry, represented as a 2D map of building heights for a given scenario and a prevailing wind direction. The desired output is the corresponding 2D map of the time-averaged wind speed at a pedestrian-relevant height of 1.5 m, as provided by the LES ground-truth fields. The objective is to learn this mapping from geometry to the high-fidelity wind field with high computational efficiency and accuracy.

In this formulation, the model input consists solely of a raw 2D building-height map, with no explicit morphological indices (e.g., aspect ratio or frontal area index). Instead, the hierarchical convolutional architecture implicitly encodes urban morphology across multiple spatial scales. This encoding can be interpreted through the network's expanding receptive fields: the initial layers capture building-scale features (e.g., local acceleration around individual structures), intermediate layers encode street-scale features (e.g., street width and connectivity affecting flow channeling), and the deep bottleneck layers aggregate these into neighbourhood-scale descriptors analogous to planar and frontal area densities. This multi-scale feature extraction allows the model to reconstruct wind fields based on learned physical relationships, such as wake formation and channeling, supporting robust generalization to the unseen urban forms in the validation set.

2.2. Overall framework

The proposed methodology is founded on the principle of decomposing the complex, direct mapping from urban geometry to a high-fidelity wind field into two distinct, more manageable sub-tasks, an approach conceptually similar to hierarchical or cascaded learning methods [52,53]. This two-stage "predictor-refiner" framework, illustrated in Fig. 1, is designed to improve both the stability of the training process and the quality of the final prediction.

The first stage, the predictor, utilizes a pre-trained U-Net model based on the architecture of Lu et al. [47]. This model serves to establish a computationally efficient, physically plausible baseline prediction. Its primary role is to learn the low-frequency components of the flow field, effectively capturing the large-scale flow structures and dominant wind paths dictated by the urban morphology. The second stage, the refiner, introduces a conditional Generative Adversarial Network (cGAN). This network is not trained on the urban geometry directly; instead, its sole input is the baseline prediction from Stage 1, which is tasked to transform it into a high-fidelity output. The refiner learns to add the missing high-frequency details, such as sharp velocity gradients and peak wind speeds, thereby correcting the systematic smoothing artifacts inherent in the baseline model. By constraining the cGAN to this refinement task, the training process is stabilized and the network's capacity is focused entirely on learning these high-frequency components of the flow.

2.3. Stage 1: Baseline predictor model

The first stage of the proposed framework is the predictor, which is responsible for generating a rapid, baseline approximation of the wind field. For this role, the pre-trained U-Net surrogate model developed by Lu et al. [47] is adopted without any re-training. Its weights, originally learned on UrbanTALES using an 80%/20% random train-test split, are reused directly. This off-the-shelf baseline is retained because it has already been shown to efficiently produce physically plausible 2D pedestrian-level wind speed maps from urban geometry, and keeping Stage 1 fixed ensures that the performance gains reported here can be attributed to the Stage-2 refinement.

The predictor model functions as a direct mapping from the urban form to the resulting pedestrian-level wind field. The input consists of a 2D array representing urban building heights, coupled with the prevailing wind direction. The model's output is a 2D map of the time-averaged wind speed at a pedestrian-relevant height of 1.5 m. To handle varied urban layouts and wind angles, the model utilizes a sophisticated pre-processing scheme where, for non-orthogonal winds, the urban geometry is rotated to align the flow with the image axes. This canonization of the input simplifies the learning task for the network (see [47] for more information on the methodology).

The employed model is a U-Net, a type of fully convolutional neural network (CNN) particularly effective for image-to-image translation tasks, as illustrated in Fig. 2. Lu et al. [47] adapted the standard U-Net architecture for this specific physical problem by incorporating additional downsampling layers for hierarchical feature extraction, batch normalization for training stability, and a final ReLU activation function

to ensure the physically necessary constraint of non-negative wind speed outputs. The model was trained on the comprehensive UrbanTALES dataset using a Mean Absolute Error (MAE) loss function. In the context of this work, this pre-trained (based on the entire UrbanTALES dataset) and validated model is used directly to generate the baseline predictions that serve as the input for the subsequent generative refiner network.

2.4. Stage 2: Generative adversarial refiner network

The second stage of the framework employs a conditional Generative Adversarial Network (cGAN) to perform the critical refinement task. This network is designed to execute a sophisticated image-to-image translation, transforming the low-fidelity baseline prediction into a high-fidelity output that emulates the ground truth. The cGAN framework consists of two neural networks, a Generator and a Discriminator, which are trained simultaneously in an adversarial process [54,55]. The proposed implementation is based on the Pix2Pix architecture [56], which is highly effective for paired image translation problems.

The overall architecture of the proposed cGAN refiner is illustrated in Fig. 3. The U-Net-based generator takes as input the baseline wind field predicted by Stage 1 and outputs a refined, high-fidelity wind field. Both the generated and ground-truth wind fields are concatenated with the source image and passed to the PatchGAN discriminator [57], which classifies local patches as 'real' or 'fake'. The generator is trained using a combination of pixel-wise reconstruction loss (L1 loss) and adversarial loss, enabling it to produce outputs that are both quantitatively accurate and perceptually realistic [56,58,59].

2.4.1. Enhanced generator architecture

The generator is responsible for producing the final refined wind speed map from the baseline prediction. To achieve the required level of detail and accuracy, a standard U-Net is insufficient. Therefore, an enhanced generator architecture was developed, based on a U-Net structure but with significant modifications to improve its capacity for multi-scale feature learning. The overall role of the generator within the cGAN framework is shown in Fig. 3.

The core of the architecture is the incorporation of stepwise kernel sizes within an 8-layer encoder. Unlike traditional U-Nets that use a small, fixed-size kernel throughout, this design employs large kernels in the upper encoder to capture broad, low-frequency spatial relationships, and progressively smaller kernels in deeper layers to extract high-frequency details. Concretely, the encoder uses strided convolutions (stride 2) with kernel sizes 32, 32, 16, 16, 8, 8, 8, and 4 across successive levels, with instance normalization in all encoder blocks except the first and LeakyReLU ($\alpha = 0.2$) activations [60]. Channel width increases from 64 up to 512 and is capped at 512 in the deepest blocks, consistent with common U-Net practice for stable capacity scaling.

The decoder mirrors the encoder structure and restores resolution using transpose convolutions (kernel 4, stride 2, padding 1). Each decoder block applies instance normalization and ReLU activations, and skip connections concatenate encoder features at matching resolutions to preserve high-resolution spatial information critical for precise reconstruction. To mitigate overfitting while the network learns fine-




Fig. 1. Overview of the two-stage predictor-refiner framework. Stage 1: A U-Net baseline predictor maps urban geometry (2D array of building height shown by the grayscale color) and inflow conditions to a rough speed field. Stage 2: A conditional GAN refiner transforms the rough prediction into a high-fidelity output matching LES-level detail.




Fig. 2. Schematic of the Stage-1 U-Net baseline predictor, showing the encoder–decoder structure and skip connections used for rough estimation, adapted from Lu et al., [47].

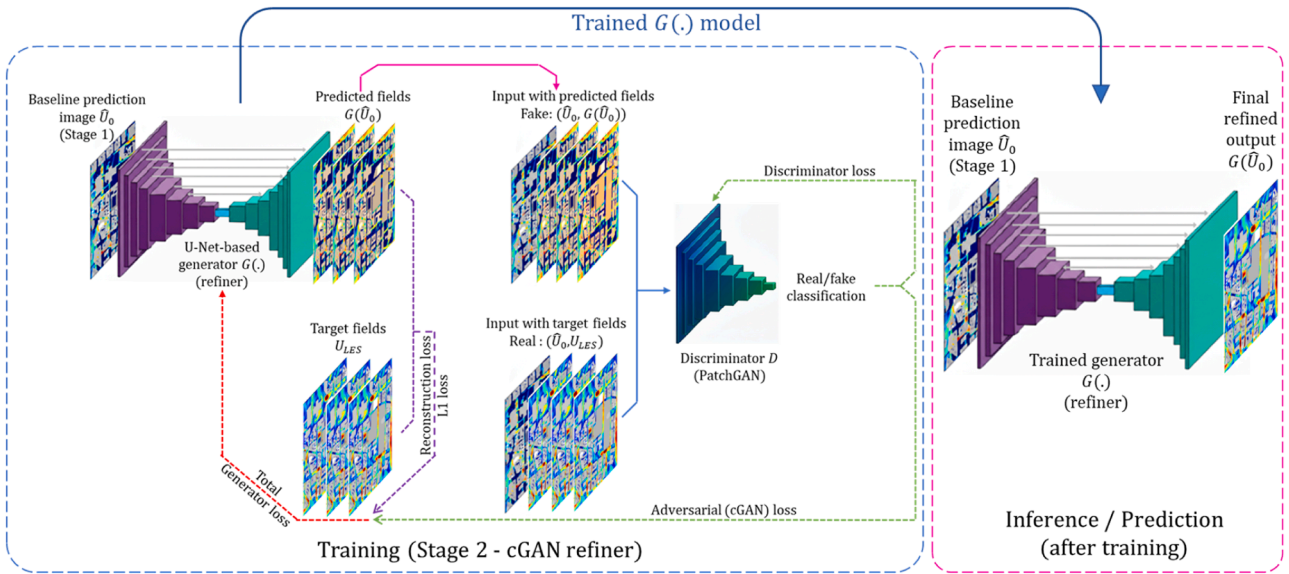


Fig. 3. Overall architecture of the proposed cGAN refiner, showing the generator, discriminator, and loss computation pathways.

scale corrections, dropout ($p = 0.5$) is applied in the three deepest decoder stages. The final layer is a transpose convolution followed by a Tanh activation, mapping the single-channel output to the normalized $[-1, 1]$ range required for stable GAN training and consistent with the data normalization used for targets.

2.4.2. Discriminator architecture

The discriminator drives the adversarial training process by distinguishing between refined predictions produced by the generator and real LES ground-truth data. Within the framework shown in Fig. 3, it receives as input a two-channel tensor created by concatenating the source image (the baseline U-Net prediction) with either the real target image (LES truth) or the generator’s refined prediction.

A PatchGAN architecture is used, consistent with the Pix2Pix framework. Rather than outputting a single scalar ‘real’ or ‘fake’ score for the entire image, the PatchGAN operates on overlapping image patches, classifying each one individually. This approach is computationally efficient and encourages the generator to produce realistic high-frequency textures and sharp details throughout the output. Concretely, the discriminator is implemented as a deep CNN composed of successive 4×4 strided convolutions (stride 2, padding 1) with instance normalization and LeakyReLU ($\alpha=0.2$) activations; the first block omits normalization for stability. A final 4×4 convolution (stride 1, padding

1) produces a 2-D logit map (no sigmoid), where each element scores the realism of a corresponding patch in the input; these logits are subsequently consumed by a BCE-with-logits objective.

2.5. Training objective and protocol

2.5.1. Training objective

Training the cGAN refiner follows the paired image-translation paradigm of Pix2Pix and uses a composite loss that combines an adversarial term (to encourage perceptual realism) with a pixel-wise reconstruction term (to preserve quantitative fidelity). Let x denote the Stage-1 baseline prediction (the generator input), y the LES ground truth, and $G(x)$ the generator output. Let D denote the PatchGAN discriminator operating on paired inputs; given a source x and a target y (or the generated output $G(x)$), it produces a pre-sigmoid logit map $f(x, y)$ with one logit per patch, and the corresponding probability map is $D(x, y) = \sigma(f(x, y))$, where σ is the elementwise sigmoid. All expectations below are averaged over the training minibatch and over patches. The conditional adversarial objective is:

$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}_{x,y}[\log D(x, y)] + \mathbb{E}_x[\log(1 - D(x, G(x)))] \quad (1)$$

which the discriminator D maximizes and the generator G minimizes. In

practice, this is implemented using a binary cross-entropy (BCE) loss. Since the PatchGAN discriminator outputs a map of predictions for each patch, the final loss is calculated by averaging the BCE loss across all patches.

To ensure that the refined field remains quantitatively close to the LES target, a pixel-wise reconstruction loss is added. Following Pix2Pix, the L1 norm is used, which is empirically found to encourage less blurring than L2:

$$\mathcal{L}_{L1}(G) = \mathbb{E}_{x,y}[\|y - G(x)\|_1] \quad (2)$$

The generator's total loss is a weighted sum of the adversarial loss (implemented as the negative log probability of the discriminator labeling the generated examples as real) and the L1 reconstruction term:

$$\mathcal{L}_G = \mathbb{E}_x[-\log D(x, G(x))] + \lambda \mathcal{L}_{L1}(G) \quad (3)$$

Consistent with Pix2Pix, $\lambda = 100$ is selected to strongly favor faithful reconstruction while still encouraging realistic high-frequency detail. Equivalently, optimization alternates (or concurrently updates) to solve:

$$\min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G) \quad (4)$$

The discriminator loss (implemented with BCE over real and fake patches) is:

$$\mathcal{L}_D = -\mathbb{E}_{x,y}[\log D(x, y)] - \mathbb{E}_x[\log(1 - D(x, G(x)))] \quad (5)$$

i.e., the negative of \mathcal{L}_{cGAN} so that minimizing \mathcal{L}_D drives D to classify real pairs as 1 and generated pairs as 0.

2.5.2. Training protocol

The refiner network requires paired examples of baseline predictions (from stage 1) and corresponding LES ground truth to learn the refinement task. To construct this dataset, the pre-trained Feilian U-Net of Lu et al. [47] is first run in inference mode on all UrbanTALES configurations to generate baseline wind fields from the building-height maps. Each Feilian prediction, paired with its corresponding LES field, defines one (baseline → LES) training example. Stage 1 is kept completely frozen throughout; no additional training or re-partitioning of its data is performed in this work. Any NaN values corresponding to building locations in the raw data arrays were replaced by zeros prior to resizing and normalization. Each paired example was resized to a uniform spatial dimension of 512×512 pixels to enable efficient batch training. Pixel values for both baseline inputs and LES targets were linearly scaled to the range $[-1, 1]$ prior to training (the generator uses a final tanh activation function). Crucially, the scaling factors (per-channel maxima) were computed on the training split only and then applied unchanged to the validation split to avoid leakage; diagnostic images saved during training were re-scaled back to physical units using these training-derived maxima. After resizing, the baseline channel was clamped to non-negative values to maintain consistency with the Stage-1 predictor's range.

To assess generalization, the available paired dataset of Feilian-LES examples was split into a training set (90 %) and a validation set (10 %). This 90 %-10 % split is defined solely over the Stage-1 prediction → LES pairs and is independent of the original 80 %/20 % train-test partition used to train Feilian in Lu et al. [47]. The refiner was trained exclusively on the Stage-2 training split, while the validation split was used only for checkpointing and early stopping.

The model was trained following the paired image-translation paradigm described in Section 2.5.1. The following settings were used: a generator learning rate of 2×10^{-4} , a discriminator learning rate of 1×10^{-4} , the Adam optimizer, and a batch size of 4 for a maximum of 500 epochs. The L1 wt λ was set to 100. The adversarial term was computed using BCEWithLogitsLoss for numerical stability. When updating the discriminator, generated samples were detached from the generator's computation graph to prevent backpropagation into G . Instance normalization was used in the majority of convolutional blocks to

stabilize training, and dropout was applied in the deepest decoder stages to reduce overfitting. An early stopping protocol with a patience of 20 epochs on the validation L1 loss was employed; the checkpoint achieving the minimum validation L1 was retained for all subsequent analyses.

3. Results

This section presents a comprehensive evaluation of the proposed two-stage predictor-refiner framework. The model's performance is assessed across a range of test cases, including diverse, real-world urban configurations and simplified, idealized geometries. The robustness of the model to variations in the inflow wind direction is also investigated. The evaluation combines qualitative visual comparisons against the LES ground truth with quantitative analysis using standard error metrics.

3.1. Training dynamics and model convergence

The cGAN refiner was trained following the Pix2Pix protocol described in Section 2.5.1, with a BCE-with-logits adversarial objective and an L1 reconstruction term weighted by $\lambda_{L1} = 100$. Inputs and targets were resized to 512×512 and linearly scaled to $[-1, 1]$ using statistics from the training split. For the supervised target channel, the training-set maximum was $\max_B = 2.7932 \text{ ms}^{-1}$. Optimization used Adam for the generator and discriminator with learning rates 2×10^{-4} and 1×10^{-4} , respectively, batch size 4, a budget of 500 epochs, checkpointing on validation L1, and early stopping with a patience of 20 epochs.

Validation error decreased steadily through mid-training and then entered the characteristic oscillatory regime of adversarial learning. The best checkpoint was obtained at epoch 146 with a validation L1 of 0.029 (normalized), which corresponds to a mean absolute error of approximately 0.041 ms^{-1} . At this epoch, the training L1 was 0.007 ($\approx 0.0103 \text{ ms}^{-1}$), yielding a generalization gap of 0.022 ($\approx 0.031 \text{ ms}^{-1}$). Training subsequently continued without surpassing this minimum and terminated at epoch 166 due to early stopping. In a ± 10 -epoch window around the best epoch, the standard deviation of the validation L1 was 0.004 ($\approx 0.0056 \text{ ms}^{-1}$), indicating a flat and stable minimum without late-epoch drift.

Adversarial dynamics were characterized by discriminator loss statistics across epochs and by decomposition of the generator objective at the selected checkpoint. The discriminator loss at epoch 146 was 0.6749, close to $\ln 2 \approx 0.693$, the theoretical value when real and generated samples are comparably hard to distinguish. Over the full trajectory, the median discriminator loss was 0.3248 with an inter-quartile range of 0.5494, reflecting intermittent phases of discriminator strength interleaved with balanced play. Measured as epoch fractions, 44.6 % of epochs lay in a balanced band (discriminator loss 0.4–1.0), 36.7 % showed discriminator saturation (loss < 0.2), and only 1.2 % indicated a weak discriminator (loss > 1.2). Decomposing the generator objective at epoch 146 yields an adversarial term $G_{adv} = 0.8658$ and a GAN-to-reconstruction ratio of 1.179, indicating that both the pixel term and the adversarial signal contributed meaningfully at selection time, consistent with sharper flow details without sacrificing pixel fidelity. Fig. 4 reports the per-epoch training and validation L1 curves. All downstream quantitative and qualitative evaluations use this best-validation checkpoint (epoch 146).

Beyond the training-curve and convergence diagnostics (Fig. 4), generalization is evaluated at dataset scale. The best-validation checkpoint is applied to all available cases (training and validation). Figs. 5–7 present distributions of per-case R^2 , RMSE, and MAE comparing the Baseline U-Net and the Hierarchical model for (i) all cases, (ii) real configurations, and (iii) idealized configurations. For readability, panels reporting R^2 exclude cases with negative baseline R^2 . The evaluation follows the same normalization and protocol described in Section 2.

Using the best-validation checkpoint, the Hierarchical model




Fig. 4. Training and validation L1 versus epoch.

outperforms the Baseline U-Net on both data splits and across configuration types. For the training split (Fig. 5), median RMSE decreases from 0.131 to 0.032, a 75.6 % reduction; median MAE decreases from 0.085 to 0.019, a 77.6 % reduction; median R^2 increases from 0.573 to 0.974, a 69.9 % relative increase. For the validation split (Fig. 5), median RMSE decreases from 0.139 to 0.056, a 59.7 % reduction; median MAE decreases from 0.091 to 0.034, a 62.6 % reduction; median R^2 increases

from 0.573 to 0.945, a 64.9 % relative increase. By configuration type, real urban cases (Fig. 6) show MAE decreasing from 0.108 to 0.036 (66.7 % reduction), RMSE from 0.169 to 0.062 (63.3 % reduction), and R^2 increasing from 0.537 to 0.943, a 75.6 % relative increase; idealized layouts (Fig. 7) show MAE decreasing from 0.070 to 0.014 (80.0 % reduction), RMSE from 0.120 to 0.022 (81.7 % reduction), and R^2 increasing from 0.589 to 0.982, a 66.7 % relative increase. Overall, the results indicate robust generalization and systematic error reduction under a single checkpoint and evaluation protocol.

The hierarchical model also demonstrates significant improvements in dataset-level wind-field statistics over the Baseline U-Net. On the training split, the MAE for the spatial mean dropped by 68.1 % (from 0.0166 m/s for the baseline to 0.0053 m/s for the hierarchical), while the MAE for spatial standard deviation fell by 79.2 % (from 0.0178 m/s to 0.0037 m/s). The error for maximum wind speed was also reduced by 35.5 %. These improvements show strong generalization to the validation split, which saw MAE reductions of 41.5 % for the mean, 57.7 % for the standard deviation, and 32.8 % for the maximum wind speed. Taken together with the case-wise error metrics, these dataset-level statistics indicate that the refiner consistently reduces bias in the bulk flow (mean), tightens variability estimates (std), and curbs extremes (max) across both splits.

3.2. Baseline architectures and comparative analysis

To isolate the individual contributions of the hierarchical model and the adversarial training strategy, two additional baseline models were implemented using the same data splits, normalization, and early-stopping protocol described in Section 2.5. The first baseline is a

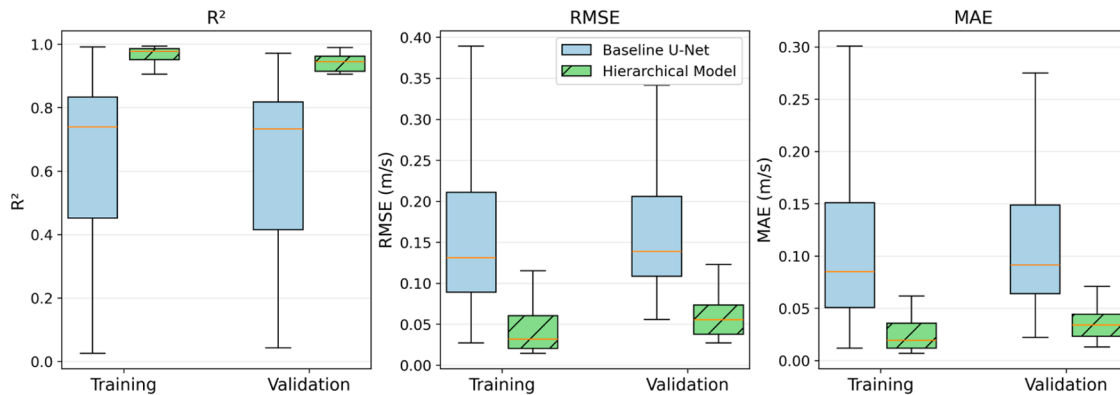


Fig. 5. Comparison of model performance on the complete dataset. The box plots illustrate the per-case distribution of (a) R^2 (left), (b) RMSE (middle), and (c) MAE (right) for the Baseline U-Net and the Hierarchical model, benchmarked against the LES ground truth data. In each box plot, the central line indicates the median across cases, the box spans the interquartile range (25th–75th percentiles), and the whiskers indicate dispersion beyond the interquartile range. Higher R^2 and lower RMSE/MAE indicate closer agreement with the LES reference.




Fig. 6. Performance evaluation on real urban configurations. The box plots present the same comparison as in Fig. 5, showing per-case distributions of (a) R^2 (left), (b) RMSE (middle), and (c) MAE (right) but restricted to the subset of realistic urban environments. Box-plot elements and metric interpretation follow Fig. 5.




Fig. 7. Performance evaluation on idealized configurations. The box plots present the same comparison as in Fig. 5, showing per-case distributions of (a) R² (left), (b) RMSE (middle), and (c) MAE (right) but restricted to the subset of idealized geometric layouts. Box-plot elements and metric interpretation follow Fig. 5.

single-stage cGAN that removes the hierarchical structure and attempts to learn the mapping directly from the urban geometry (building-height map) to the high-fidelity LES wind field. The generator architecture is identical to the Stage-2 refiner described in Section 2.4.1 (Enhanced U-Net), and the same PatchGAN discriminator as in Section 2.4.2 is used. Training follows the pix2pix-style objective of Eq. (4), combining a pixel-wise L1 reconstruction term with an adversarial loss. The only substantive difference with respect to the proposed framework is the input domain: the generator receives the raw geometry instead of the low-frequency baseline prediction produced by Stage 1.

The second baseline targets the specific contribution of the adversarial loss. The overall predictor-refiner structure is retained: the pre-trained Stage-1 U-Net (Section 2.3) first produces a baseline wind field from the urban geometry, and a second network refines this prediction toward the LES target. The refiner uses exactly the same Enhanced U-Net architecture as the Stage-2 generator described in Section 2.4.1, but the discriminator is removed and the network is trained purely as a convolutional regressor. Training minimizes the pixel-wise L1 loss between the refined prediction and the LES field (Eq. (2)), without any adversarial term. In practice, the Stage-1 model is kept fixed and is used off-line to generate baseline predictions for all cases; these predictions are then paired with the corresponding LES outputs to form the refiner training dataset. By removing the adversarial feedback loop while keeping the same hierarchical input and refiner architecture, this model isolates the performance gains that are strictly attributable to GAN-based training versus those arising from the hierarchical design and deep convolutional architecture alone.

The performance of the proposed hierarchical cGAN model is evaluated against these two baselines to quantify the benefits of the two-stage architecture and the adversarial training objective. Table 1 summarizes the R², RMSE, and MAE values for all three models on the training and validation splits. Values for the proposed model correspond to the median metrics reported in Section 3.1.

The single-stage cGAN exhibits the weakest performance among the evaluated architectures, with a validation R² of 0.518 and an RMSE of 0.199 m/s. Despite using the same generator capacity and adversarial loss as the proposed model, the absence of a guiding baseline leads to poor reconstruction of the global flow structure. This behaviour is

Table 1

Quantitative comparison of the proposed hierarchical cGAN and baseline models.

Model Architecture	Split	R ²	RMSE (m/s)	MAE (m/s)
Single-stage cGAN	Training	0.626	0.167	0.103
	Validation	0.518	0.199	0.125
Hierarchical CNN refiner	Training	0.953	0.059	0.032
	Validation	0.810	0.125	0.072
Proposed hierarchical cGAN	Training	0.974	0.032	0.019
	Validation	0.945	0.056	0.034

consistent with the intuition that directly mapping complex, irregular urban geometries to fully developed turbulent fields is ill-posed for a single network: the model must simultaneously learn low-frequency channeling effects and high-frequency fluctuations from geometry alone.

The hierarchical CNN refiner shows that the two-stage predictor-refiner design is structurally superior to the single-stage approach, improving the validation R² to 0.810 and reducing the validation RMSE to 0.125 m/s. However, this model exhibits a generalization gap: the training RMSE is 0.059 m/s, whereas the validation RMSE more than doubles, indicating sensitivity to overfitting when trained solely with an L1 reconstruction loss. In addition, the exclusive use of the L1 objective encourages predictions close to the conditional mean, yielding smoother fields that under-represent sharp gradients and peak wind speeds in the LES.

In contrast, the proposed hierarchical cGAN achieves the highest fidelity and robustness. By combining the Stage-1 baseline for large-scale structure with an adversarial Stage-2 refiner for small-scale texture, the model attains a validation RMSE of 0.056 m/s. This corresponds to a reduction of approximately 55 % relative to the hierarchical CNN refiner (0.125 m/s) and about 72 % relative to the single-stage cGAN (0.199 m/s). The generalization gap in RMSE is also considerably smaller (0.032 m/s in training vs 0.056 m/s in validation), compared with 0.059 m/s vs 0.125 m/s for the CNN refiner. These results indicate that the adversarial loss not only drives the generator to resolve high-frequency details but also acts as an effective regularizer, mitigating overfitting and improving generalization compared with a purely regression-based refiner.

3.3. Case studies: real urban configurations

The performance of the predictor-refiner framework is illustrated on a set of six realistic city geometries from the UrbanTALES dataset [48] that cover 1) a compact ($\lambda_p = 0.49$) commercial core district, 2) a sparse ($\lambda_p = 0.21$) suburban residential area, 3) a compact ($\lambda_p = 0.47$) modern residential grid, 4) a compact ($\lambda_p = 0.46$) dense residential blocks, 5) an intermediate ($\lambda_p = 0.28$) mixed-use urban blocks, and 6) a sparse ($\lambda_p = 0.14$) modern residential grid where the original Lu et al. [47] model performs poorly. The evaluation combines a qualitative visual analysis of the generated wind fields with a quantitative assessment using standard error metrics. Figs. 8 and 9 present a visual comparison of the model's performance for the six different selected city configurations under a 0° prevailing wind. In all cases, the baseline predictor (left panel in each figure) captures the broad spatial distribution of the wind, identifying the primary wind corridors and low-speed zones. However, the predictions are characterized by excessive smoothing and a failure to resolve localized, high-speed wind jets. In contrast, the refined predictions (middle panel) demonstrate a marked increase in fidelity. Across the six cities, the refiner network successfully sharpens the flow

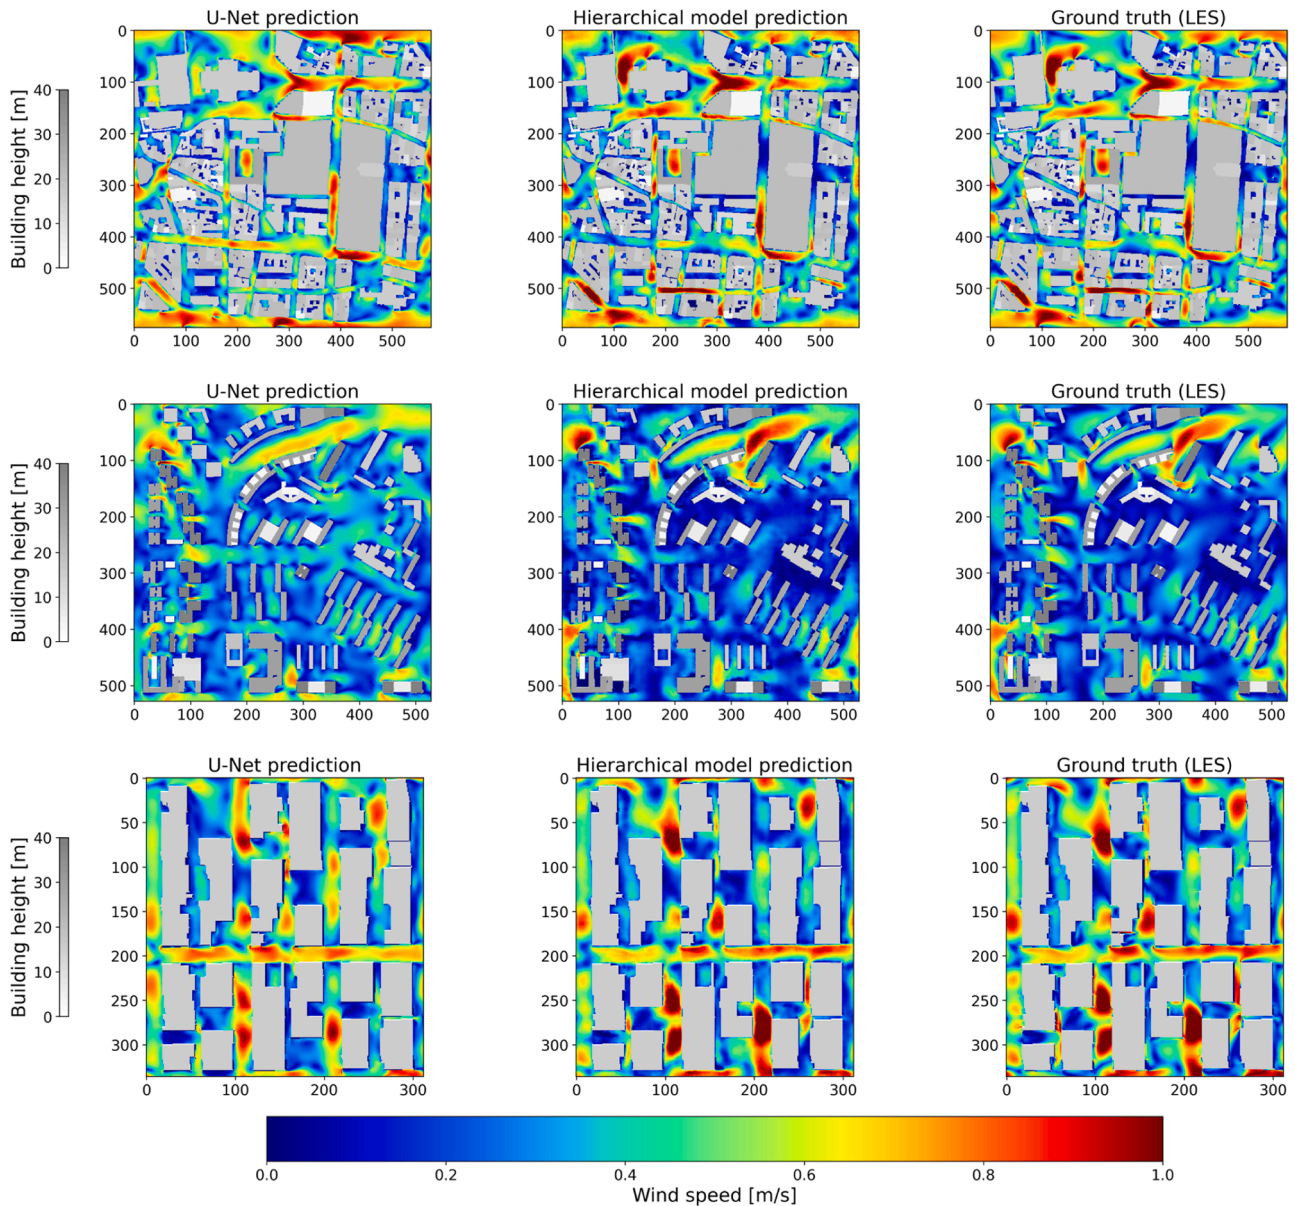


Fig. 8. Qualitative comparison (0° wind). Each row presents a different city: (top) FR-PA-V2_d00, (middle) ES-Bar-V1_d00, and (bottom) CA-Van-U2_d00. Columns compare the baseline U-Net prediction (left), the hierarchical model prediction (middle), and the LES ground truth (right). The naming convention can be found in Nazarian et al. [48].

field. It restores sharp gradients and adds back critical details in wake regions, producing a mean flow that closely matches the LES ground truth (right panel).

The visual improvements are corroborated by Fig. 10. Across the six test cities, mean RMSE decreases by 67 % and median RMSE decreases by 60 %. Per-city RMSE decreases are FR-PA-V2: 66.0 %, ES-Bar-V1: 66.3 %, CA-Van-U2: 59.7 %, BR-SP-U2: 81.6 %, CH-Bas-V1: 54.7 %, and US-Det-U9: 50.7 %. Variability also tightens markedly, with the RMSE standard deviation reduced by $\approx 90\%$. For R^2 , the mean increases by $\approx 141\%$ and the median increases by $\approx 130\%$; moreover, the share of cases with $R^2 \geq 0.90$ rises from 0 % to 100 %. Dispersion is similarly reduced, with the R^2 standard deviation dropping by $\approx 95\%$. Collectively, these percentage gains indicate strong generalization across morphologically diverse cities and reliable recovery of the fine-scale structure required for accurate pedestrian-level wind estimation.

3.4. Case studies: idealized configurations

The capability of the proposed model on regular geometries was assessed using three idealized building-array cases: UA53_d00, US16_d00, and VA16-CL-28_d00 as illustrated in Fig. 11. The Baseline U-Net captures the dominant channelized flow through the arrays but retains a smoothed appearance that blurs shear layers and wake recirculation. The Hierarchical model (predictor + refiner) sharpens these features, reinstating concentrated high-speed streaks along building edges and clarifying wake decay patterns visible in the LES reference.

For quantitative context, Fig. 12 summarizes the per-case RMSE and R^2 metrics for both models. The error magnitudes decrease substantially after refinement. For instance, the RMSE for case UA53_d00 was reduced by 90.5 % (from 0.3638 to 0.0346 m s^{-1}), while cases US16_d00 and VA16-CL-28_d00 saw reductions of 83.2 % (from 0.1363 to 0.0229 m s^{-1}) and 83.1 % (from 0.1441 to 0.0243 m s^{-1}), respectively. Across these three cases, the mean and median RMSE reductions were 85.6 % and 83.2 %. The goodness-of-fit also improved significantly, with R^2 values

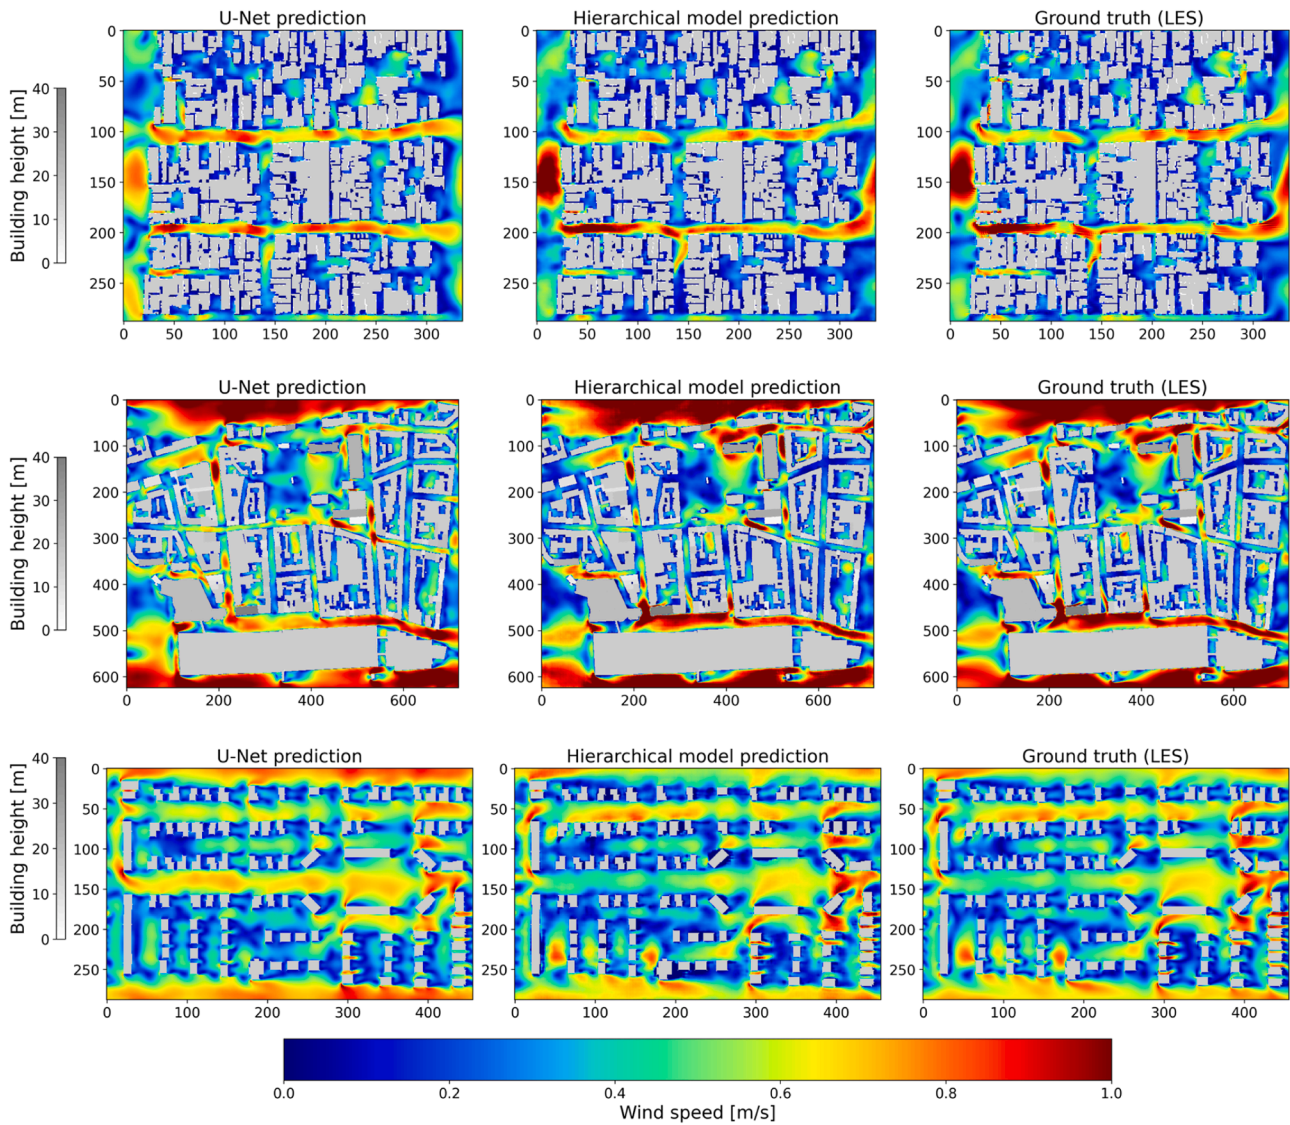


Fig. 9. Qualitative comparison (0° wind). Each row presents a different city: (top) BR-SP-U2_d00, (middle) CH-Bas-V1_d00, and (bottom) US-Det-U9_d00. Columns compare the baseline U-Net prediction (left), the hierarchical model prediction (middle), and the LES ground truth (right). The naming convention can be found in Nazarian et al. [48].




Fig. 10. Quantitative comparison of model performance across six diverse city configurations. The two subplots compare the performance of the baseline U-Net against the final refined model. (Left) Root Mean Square Error (RMSE); (Right) R^2 Score.

increasing from -1.135 to 0.981 (UA53_d00), 0.050 to 0.973 (US16_d00), and 0.254 to 0.979 (VA16-CL-28_d00).

3.5. The effect of wind direction

A key test of a surrogate model's robustness is its ability to generalize to different inflow directions, as this significantly alters the complex aerodynamic interactions within the urban canopy. To evaluate the

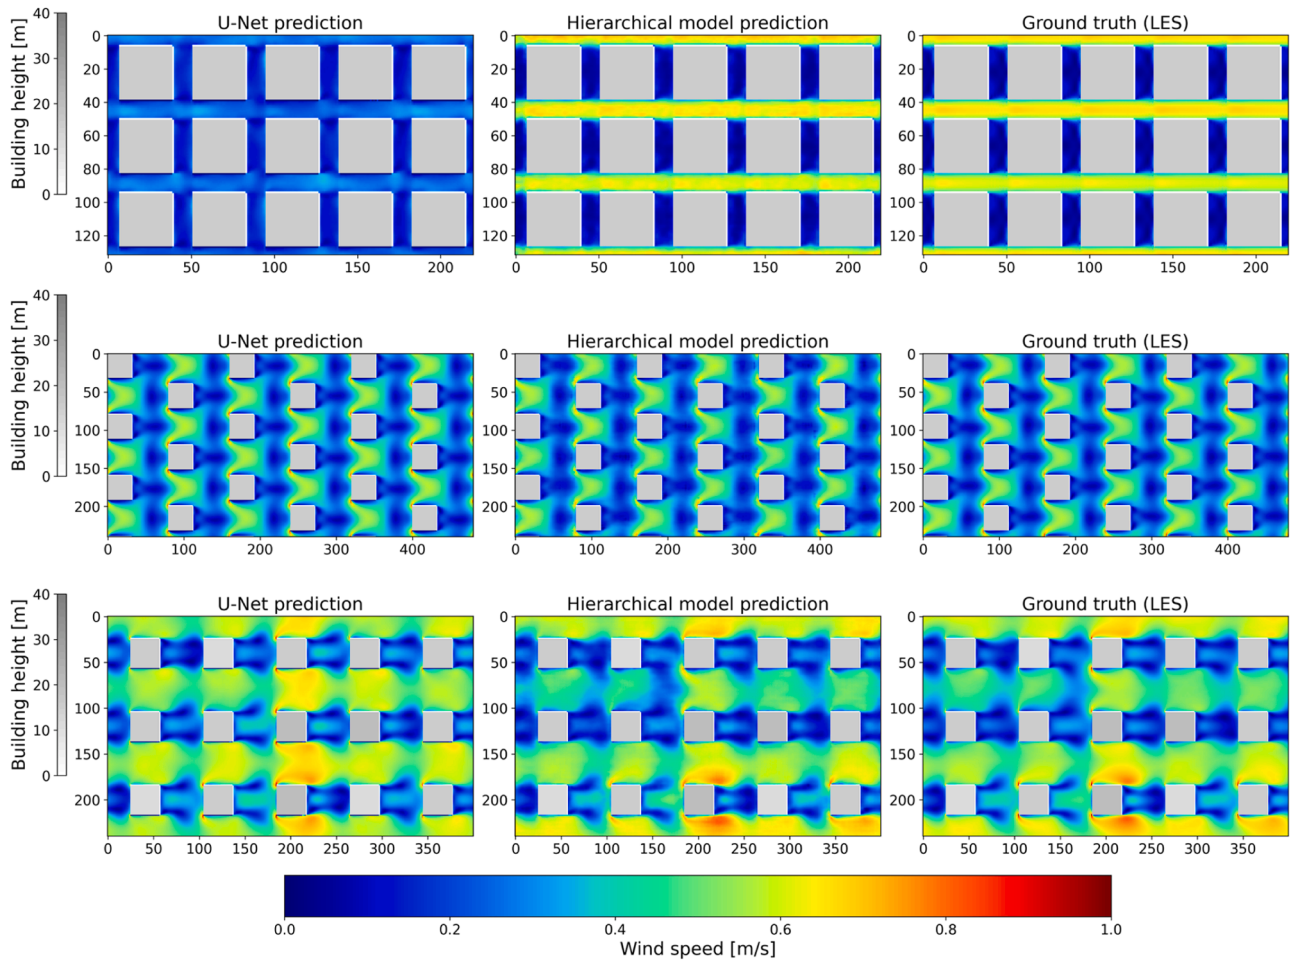


Fig. 11. Qualitative comparison for idealized urban geometries. Each row shows a different idealized test case under a 0° wind direction, comparing the baseline U-Net prediction (left), the final refined prediction (middle), and the LES ground truth (right). Cases shown: UA53_d00, US16_d00, and VA16-CL-28_d00.




Fig. 12. Quantitative comparison of model performance across three diverse idealized configurations. The two subplots compare the performance of the baseline U-Net against the final refined model. (Left) Root Mean Square Error (RMSE); (Right) R² Score.

framework’s directional robustness, the KO-SE-V11 test case was evaluated for five inflow angles: 0°, 15°, 30°, 45°, and 90°.

As shown in Fig. 13, the refiner network improves the KO-SE-V11 predictions for all five inflow angles (0°, 15°, 30°, 45°, 90°). The qualitative gains are most evident for oblique (15°–45°) and perpendicular (90°) winds: the baseline captures broad recirculation but misses narrow high-speed jets and underestimates wake intensities, whereas the refiner restores concentrated jet streaks along passages and corrects wake magnitudes, yielding patterns that align closely with the LES fields. The quantitative results mirror these visual trends. RMSE decreases from 0.1607 to 0.0441 m/s at 0° (–72.6%), 0.2853 to 0.0521 m/s at 15° (–81.7%), 0.3698 to 0.0629 m/s at 30° (–83.0%), 0.3591 to 0.0615 m/s

at 45° (–82.9%), and 0.2368 to 0.0847 m/s at 90° (–64.2%). Despite the strong angle dependence of the underlying flow physics, the refined model maintains a tight error band of 0.044–0.085 m/s across directions, indicating that the learned refinement is not tied to a single orientation and generalizes reliably to wind rotation.

4. Discussion

This study shows that decomposing the surrogate task into a two-stage predictor–refiner framework materially increases the fidelity of urban wind-field estimates. The Stage-1 U-Net delivers a fast, physically plausible baseline that preserves large-scale organization, while the

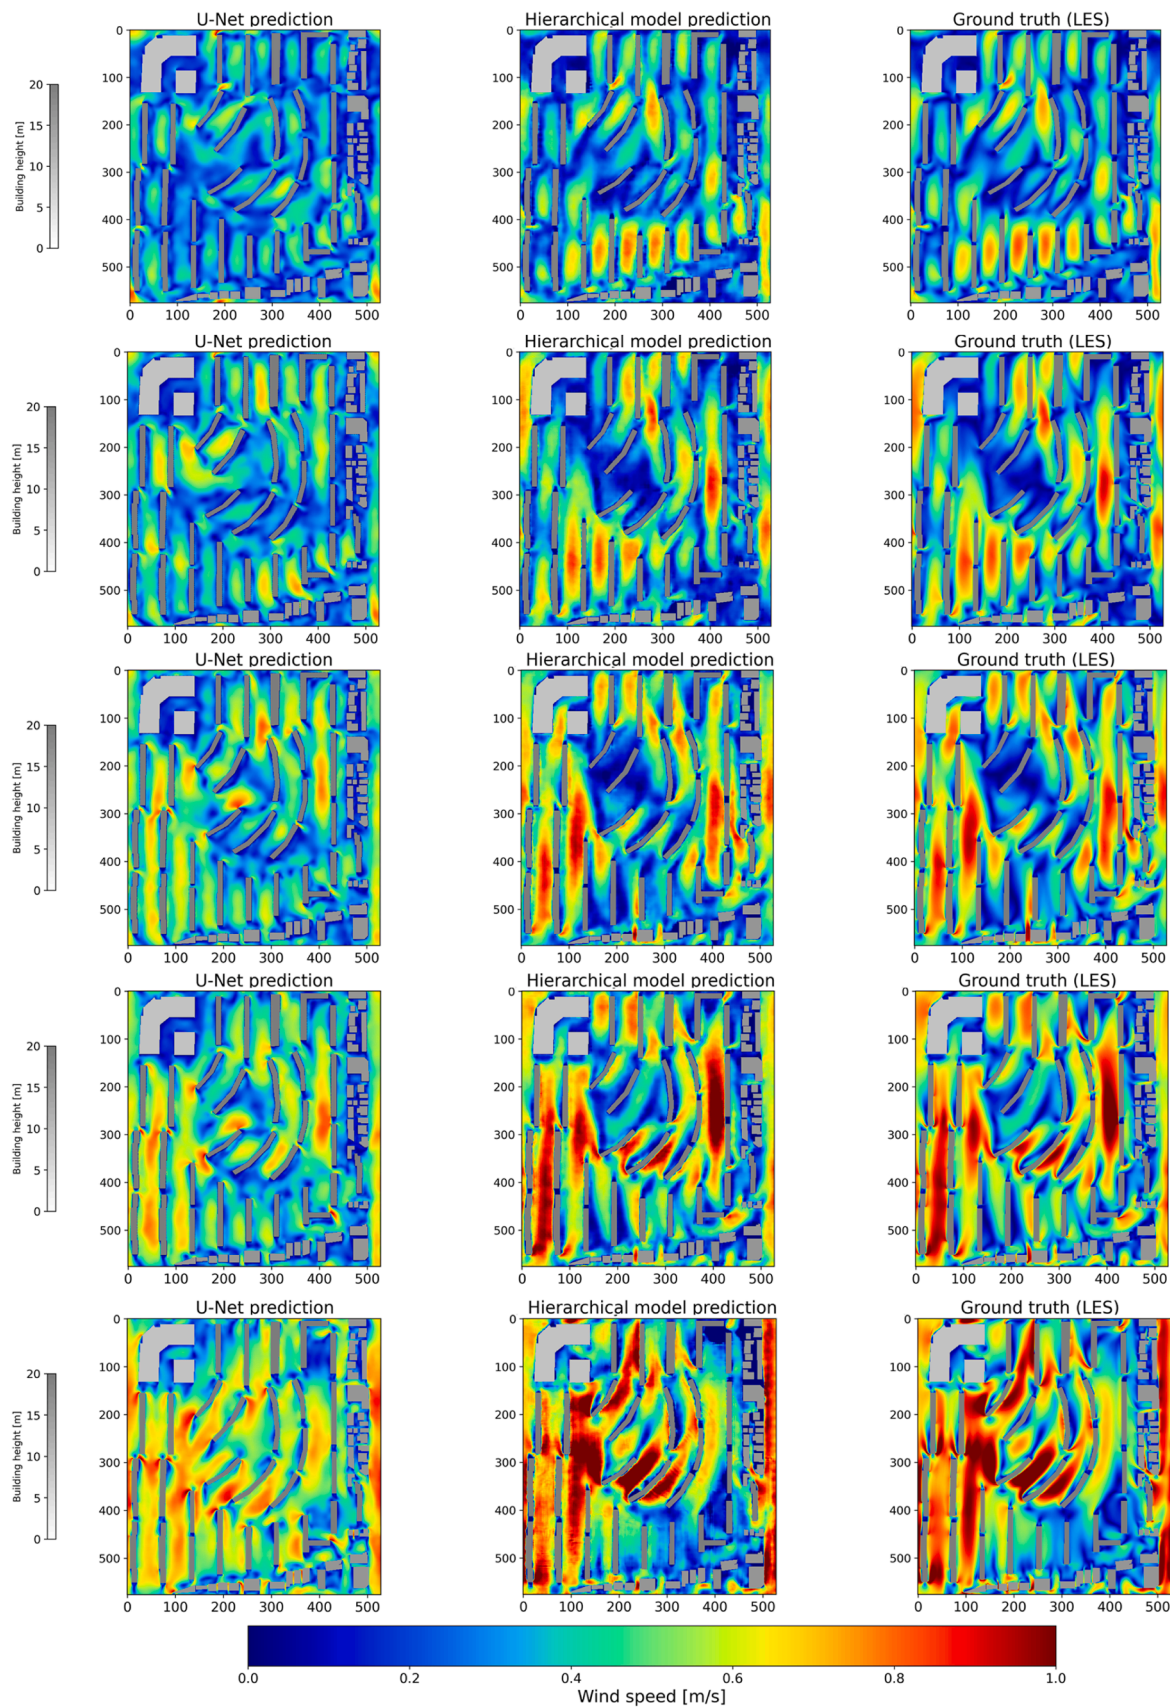


Fig. 13. Qualitative comparison of model performance under varying wind directions. Results for the KO-SE-V11 case under five different inflow angles: 0°, 15°, 30°, 45°, and 90°. Each row compares the baseline U-Net prediction (left), the final refined prediction (middle), and the LES ground truth (right).

Stage-2 conditional GAN (Pix2Pix with PatchGAN) specializes on the missing high-frequency content. From a computational standpoint, the hierarchical surrogate remains lightweight relative to the underlying LES solver. For a single 512×512 sample, the Stage-1 baseline predictor (Feilian U-Net) contains approximately 34.6 million trainable parameters and requires about 152.8 GFLOPs per forward pass, whereas the Stage-2 refiner (Enhanced U-Net generator) adds roughly 139.8 million parameters and 279.6 GFLOPs. In practice, these complexities translate to manageable inference times on standard hardware. On the desktop CPU used in this study, the Stage-1 baseline requires approximately 0.64 s per scenario, with the refinement stage adding an overhead of 0.84 s, resulting in a total end-to-end latency of ≈ 1.5 s. While this represents an increase of about $2.3 \times$ compared to the baseline alone, the sub-two-second total remains negligible compared to the hours or days required for the LES simulations that provide the ground truth data. In exchange for this modest overhead, the refiner consistently recovers sharp shear zones, narrow jet streaks, and corrected wake intensities across diverse urban morphologies, features that are essential for pedestrian-level wind assessment.

Training dynamics support the plausibility of these gains. Using BCE-with-logits adversarial loss and an L1 reconstruction term with $\lambda = 100$, early stopping on validation L1 selected a stable checkpoint (best at epoch 146) with a small generalization gap (Train L1 0.007, Val L1 0.029, gap 0.022). Discriminator health metrics remained in a productive regime (median \mathcal{D} -loss 0.325), and the generator's GAN-to-reconstruction contribution ratio at the selected checkpoint (≈ 1.18) indicates that sharpening from the adversarial game complemented, rather than overwhelmed, pixel fidelity. Validation L1 varied by only ± 0.004 around the minimum within a ± 10 -epoch window, suggesting convergence without late-epoch drift.

The quantitative evidence is consistent and strong. For six realistic cities, RMSE decreases by 50–82 % on a per-case basis (FR-PA-V2: 66 %, ES-Bar-V1: 66 %, CA-Van-U2: 60 %, BR-SP-U2: 82 %, CH-Bas-V1: 55 %, US-Det-U9: 51 %), with corresponding tightening of dispersion across cities. On canonical regular grids, where the baseline is already comparatively well calibrated, the refiner still removes most residual error, yielding ≈ 84 % RMSE reductions on average. Directional robustness is also evident: for KO-SE-V11, RMSE improves by 64–83 % over 0° , 15° , 30° , 45° , 90° , while maintaining a tight absolute error band of 0.044–0.085 m/s across angles. Together, these regimes indicate that the learned mapping is not narrowly tuned to a specific morphology or orientation, but generalizes to the altered flow topologies induced by wind rotation and layout variability.

The architectural choices in the refiner appear central to these outcomes. Step-wise large receptive fields in the encoder (32→16→8) provide multi-scale context that helps correct low-frequency bias while preserving high-frequency detail via skip connections. Instance normalization and moderated dropout in early decoder layers balance stability and capacity. Crucially, training-set-only scaling to $[-1, 1]$ avoids leakage, and denormalization restores predictions to physical units before all metric computations, ensuring that reported RMSE values are interpretable in m/s.

Comparison with the two alternative baselines highlights the distinct contributions of the hierarchical design and the adversarial objective. A single-stage cGAN that maps geometry directly to LES, despite sharing the same generator and discriminator as Stage 2, underperforms both in accuracy and in capturing coherent large-scale flow patterns, indicating that the hierarchical decomposition (predictor → refiner) is essential rather than merely adding GAN-based training. Introducing the hierarchical CNN refiner, retaining the two-stage structure but removing the discriminator and training the refiner purely with an L1 loss, already yields a substantial reduction in error relative to the single-stage cGAN, but still produces comparatively smooth fields and a larger generalization gap between training and validation. In contrast, the proposed hierarchical cGAN combines the structural advantages of the two-stage

predictor-refiner layout with adversarial supervision, further reducing errors, sharpening wakes and jets, and simultaneously narrowing the generalization gap, which suggests that the GAN loss serves both as a mechanism for recovering high-frequency content and as an effective regularizer.

From a practitioner perspective, the proposed framework is designed for rapid, design-iteration workflows rather than one-off high-fidelity simulation. In practice, the required inputs are minimal: (i) a 2D rasterized building-height map (readily obtained from GIS layers or exported from basic CAD/BIM massing models) and (ii) a target/prevaling wind direction; the preprocessing step rotates/canonizes the geometry to support arbitrary wind angles. The model outputs a 2D pedestrian-level ($z = 1.5$ m) time-averaged wind-speed field that can be used to quickly identify wind hot-spots (jets), sheltered/stagnation regions, and the spatial redistribution of these features across alternatives. This supports screening-level evaluation and “what-if” testing (e.g., adding/removing a tower, changing a footprint, opening a corridor) in seconds, enabling an interactive feedback loop within the Feilian Version 2 platform. For early-stage designs where massing and heights may be approximate, results should be interpreted primarily in a relative sense (comparison across alternatives), with higher confidence expected once geometry is refined and, when required by the application, complemented by targeted CFD or wind-tunnel validation.

While the model performs well across diverse cities and inflow angles, some limitations remain. The surrogate refines a single 2D plane (pedestrian height) of time-averaged wind speed, not the full 3D or unsteady flow; this enables tractable training and fast inference but omits vertical coupling (e.g., rooftop separation, lofting) and temporal variability relevant to comfort and safety. Only mean wind speed is supervised, therefore other key statistics of urban flow dynamics (turbulence intensity, TKE, gust factor, directional persistence, near-wall shear) are not guaranteed. In addition, the objective function combines adversarial and L1 terms without hard physics constraints (e.g., mass conservation or strict boundary adherence), so plausibility is learned statistically rather than enforced. As with any supervised surrogate, the model is most reliable when applied within the morphology distribution represented in UrbanTALES; predictions for urban fabrics far outside this envelope (e.g., markedly different roughness regimes, vegetation/porosity patterns, or extreme height-density combinations absent from the dataset) may degrade. In such cases, the surrogate should be used for screening only, and domain adaptation and active-learning extensions offer a practical pathway to systematically expand coverage.

The present framework is developed and validated using UrbanTALES LES under neutral atmospheric stability, driven by mechanically forced flow (constant pressure gradient with fixed friction velocity) and therefore does not account for buoyancy effects associated with stable/unstable stratification or surface heating. Consequently, the model should be interpreted as a surrogate for mechanically driven pedestrian-level winds, and its predictions may be biased in situations where thermal forcing materially alters near-ground flow and mixing. Extending the approach to non-neutral conditions would be of clear interest: future versions could augment the inputs with stability/thermal descriptors (e.g., near-surface temperature or heat-flux proxies, potentially as an additional “thermal channel”) and retrain on stratified LES datasets when available, and/or couple the surrogate with mesoscale/NWP outputs that provide background stability information while using the current model as a neutral baseline. While the proposed framework is data-driven, it delivers outputs that are directly interpretable in a planning and design context. The final hierarchical model produces a high-resolution pedestrian-level ($z = 1.5$ m) time-averaged wind-speed field, which can be used to visually and quantitatively identify wind hot-spots (jets), sheltered/stagnation zones, and wake-affected regions around buildings and street corridors. To support practitioner use, these predicted fields can be post-processed into design-oriented indicators such as hot-spot localization (spatial maxima), exceedance/percentile

maps, and area-based discomfort metrics, enabling rapid comparison across design alternatives and prioritization of mitigation measures. As an extension, explainable-AI methods (e.g., saliency/occlusion/Integrated Gradients applied to the input height map) and uncertainty/statistical hot-spot characterization can be incorporated to highlight influential geometric regions and quantify confidence, while maintaining the same simulation-free inference workflow at deployment. Future work could address these points and others via multi-target supervision (adding TI/TKE/gusts), lightweight physics priors (divergence penalties, boundary consistency), uncertainty quantification, explainability, domain adaptation and active learning for high-error regimes, and extensions to 3D or time-aware models that retain the practical efficiency demonstrated here.

5. Conclusion

This study introduced a two-stage deep learning model for urban wind prediction that couples a U-Net predictor with a Pix2Pix cGAN refiner operating on the predictor's output. The first stage provides a fast, physically plausible baseline that preserves large-scale flow structure, while the second stage focuses on reconstructing the high-frequency content that is systematically smoothed out by regression-based surrogates. Trained on the UrbanTALES dataset, the hierarchical model substantially improves accuracy relative to the Stage-1 U-Net across both realistic and idealized configurations. On dataset-level metrics, the refiner reduces RMSE and MAE by approximately 60–80 % on the validation split, and increases median R^2 into the 0.94–0.97 range. Case-wise analyses confirm that the refined fields recover sharp shear layers, narrow jet streaks, and wake deficits that are strongly underestimated in the baseline predictions, while also improving bulk statistics such as spatial mean, standard deviation, and maximum wind speed. The ablation baselines further clarify the source of these gains. A single-stage cGAN that maps geometry directly to LES underperforms the hierarchical approach, indicating that the two-stage predictor-refiner structure is critical for capturing global flow organization. A hierarchical CNN refiner trained only with L1 loss improves upon the baseline U-Net but exhibits a larger generalization gap and produces overly smooth corrections. In contrast, the proposed hierarchical cGAN combines the benefits of the two-stage design with adversarial training, yielding sharper, more physically realistic wind fields and more robust validation performance. Overall, the results demonstrate that a hierarchical, adversarially trained surrogate can bridge much of the gap between conventional deep surrogates and LES, providing wind fields that are both physically interpretable and accurate enough to support applications in urban planning, pedestrian wind comfort, and wind hazard screening.

CRedit authorship contribution statement

Reda Snaiki: Writing – original draft, Visualization, Validation, Methodology, Formal analysis, Conceptualization. **Jiachen Lu:** Writing – review & editing, Data curation, Conceptualization. **Shaopeng Li:** Writing – review & editing, Conceptualization. **Negin Nazarian:** Writing – review & editing, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was partly supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) [grant number CRSNG RGPIN 2022-03492].

Data availability

Data are available at: <https://urbantales.climate-resilientcities.com/>.

References

- [1] N. BenMoshe, E. Fattal, B. Leiti, Y. Arav, Using machine learning to predict wind flow in urban areas, *Atmos. (Basel)* 14 (6) (2023) 990.
- [2] H. Mittal, A. Sharma, A. Gairola, A review on the study of urban wind at the pedestrian level around buildings, *J. Build. Eng.* 18 (2018) 154–163.
- [3] C.S.B. Grimmond, T.S. King, M. Roth, T.R. Oke, Aerodynamic roughness of urban areas derived from wind observations, *Bound. Layer Meteorol.* 89 (1) (1998) 1–24.
- [4] N. Nazarian, M. Lipson, L.K. Norford, Multiscale Modeling Techniques to Document Urban Climate change. In *Urban Climate Change and Heat Islands*, Elsevier, 2023, pp. 123–164.
- [5] J. Zhong, J. Liu, Y. Zhao, J. Niu, J. Carmeliet, Recent advances in modeling turbulent wind flow at pedestrian-level in the built environment, *Archit. intell.* 1 (1) (2022) 5.
- [6] D. Fenner, A. Christen, S. Grimmond, F. Meier, W. Morrison, M. Zeeman, J. Barlow, J. Birkmann, L. Blunn, N. Chrysoulakis, M. Clements, Urbisphere-berlin campaign: investigating multiscale urban impacts on the atmospheric boundary layer, *Bull. Am. Meteorol. Soc.* 105 (10) (2024) E1929–E1961.
- [7] X. Lyu, Y. He, S. Yin, S. Wong, T.K. Tse, E. Ng, Evaluating urban design strategies for pedestrian-level ventilation improvement in a high-density urban living environment—A LiDAR and wind tunnel study, *Build. Environ.* 269 (2025) 112439.
- [8] Q. Li, J. Chen, X. Luo, Estimating omnidirectional urban vertical wind speed with direction-dependent building morphologies, *Energy Build.* 303 (2024) 113749.
- [9] E. Ng, C. Yuan, L. Chen, C. Ren, J.C. Fung, Improving the wind environment in high-density cities by understanding urban morphology and surface roughness: a study in Hong Kong, *Landsc. Urban. Plan.* 101 (1) (2011) 59–74.
- [10] G. Shen, S. Zheng, Y. Jiang, W. Zhou, D. Zhu, An improved method for calculating urban ground roughness considering the length and angle of upwind sector, *Build. Environ.* 266 (2024) 112144.
- [11] Z. Kaseb, M. Hafezi, M. Tahbaz, S. Delfani, A framework for pedestrian-level wind conditions improvement in urban areas: CFD simulation and optimization, *Build. Environ.* 184 (2020) 107191.
- [12] G. Kim, G. Lee, Validation of CFD models of urban microclimates under high temperature and humidity conditions during daytime heatwaves in dense low-rise areas, *Build. Environ.* 266 (2024) 112087.
- [13] A. Ricci, I. Kalkman, B. Blocken, M. Burlando, M.P. Repetto, Impact of turbulence models and roughness height in 3D steady RANS simulations of wind flow in an urban environment, *Build. Environ.* 171 (2020) 106617.
- [14] A. Ricci, M. Burlando, M.P. Repetto, B. Blocken, Static downscaling of mesoscale wind conditions into an urban canopy layer by a CFD microscale model, *Build. Environ.* 225 (2022) 109626.
- [15] T. Talwar, C. Yuan, Impact of natural urban terrain on the pedestrian wind environment in neighborhoods: a CFD study with both wind and buoyancy-driven scenarios, *Build. Environ.* 261 (2024) 111746.
- [16] T. Sanemitsu, W. Wang, N. Ikegaya, Impact of momentum sources on pedestrian-level peak wind predictions of an urban-like array using large-eddy simulations and statistical models, *Build. Environ.* 267 (2025) 112244.
- [17] B. Blocken, Computational Fluid Dynamics for urban physics: importance, scales, possibilities, limitations and ten tips and tricks towards accurate and reliable simulations, *Build. Environ.* 91 (2015) 219–245.
- [18] R. Buccolieri, J.L. Santiago, A. Martilli, CFD modelling: the most useful tool for developing mesoscale urban canopy parameterizations, in: *Building Simulation, 14*, Tsinghua University Press, Beijing, 2021, pp. 407–419.
- [19] J. Lu, N. Nazarian, M.A. Hart, E.S. Krayenhoff, A. Martilli, Novel geometric parameters for assessing flow over realistic versus idealized urban arrays, *J. Adv. Model. Earth Syst.* 15 (7) (2023) e2022MS003287.
- [20] J. Lu, N. Nazarian, M.A. Hart, E.S. Krayenhoff, A. Martilli, Representing the effects of building height variability on urban canopy flow, *Q. J. R. Meteorol. Soc.* 150 (758) (2024) 46–67.
- [21] U. Shaukat, J.B. Jakobsen, N. Ikegaya, K.E.T. Giljarhus, Validation of large eddy simulations in urban wind studies using a new overall area metric, *Build. Environ.* 267 (2025) 112285.
- [22] S.L. Brunton, B.R. Noack, P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid. Mech.* 52 (1) (2020) 477–508.
- [23] R. Snaiki, A. Merabtime, Recent advances on machine learning techniques for urban heat island applications: a review and new horizons, *Sustain. Cities Soc.* (2025) 106943.
- [24] T. Wu, R. Snaiki, Applications of machine learning to wind engineering, *Front. Built Environ.* 8 (2022) 811460.
- [25] P. Xie, T. Li, J. Liu, S. Du, X. Yang, J. Zhang, Urban flow prediction from spatiotemporal data using machine learning: a survey, *Inf. Fusion* 59 (2020) 1–12.
- [26] A. Clarke, K.E.T. Giljarhus, L. Oggiano, A. Saddington, K. Depuru-Mohan, Deep learning for urban wind prediction: an MLP-mixer approach with 3D encoding, *Build. Environ.* (2025) 113495.
- [27] P. Kastner, T. Dogan, A GAN-based surrogate model for instantaneous urban wind flow prediction, *Build. Environ.* 242 (2023) 110384.
- [28] J. Wehrle, C. Jung, M. Giometto, A. Christen, D. Schindler, Introducing new morphometric parameters to improve urban canopy air flow modeling: a CFD to

- machine-learning study in real urban environments, *Urban Clim.* 58 (2024) 102173.
- [29] Y. Wu, Z. Cao, Y. Lei, Y. Han, M. Yuan, L. Wang, X. Zhou, Multi-constraint physics-informed generative adversarial networks (PIGAN) enable small-data learning for urban wind field prediction, *Build. Environ.* (2026) 114202.
- [30] D. Xiao, C.E. Heaney, L. Mottet, F. Fang, W. Lin, I.M. Navon, Y. Guo, O.K. Matar, A. G. Robins, C.C. Pain, A reduced order model for turbulent flows in the urban environment using machine learning, *Build. Environ.* 148 (2019) 323–337.
- [31] Z. Zhao, H. Li, S. Wang, Machine learning-based surrogate models for fast impact assessment of a new building on urban local microclimate at design stage, *Build. Environ.* 266 (2024) 112142.
- [32] Y. Lu, X.H. Zhou, H. Xiao, Q. Li, Using machine learning to predict urban canopy flows for land surface modeling, *Geophys. Res. Lett.* 50 (1) (2023) e2022GL102313.
- [33] Vargiomezis, T. and Gorré, C., 2025. From large-eddy simulations to deep learning: a U-net model for fast urban canopy flow predictions. arXiv preprint arXiv: 2507.06533.
- [34] W. Peng, S. Qin, S. Yang, J. Wang, X. Liu, L.L. Wang, Fourier neural operator for real-time simulation of 3D dynamic urban microclimate, *Build. Environ.* 248 (2024) 111063.
- [35] S. Qin, D. Zhan, D. Geng, W. Peng, G. Tian, Y. Shi, N. Gao, X. Liu, L.L. Wang, Modeling multivariable high-resolution 3D urban microclimate using localized fourier neural operator, *Build. Environ.* 273 (2025) 112668.
- [36] G. Tian, S. Qin, D. Geng, D. Zhan, S. Yang, P. Liu, N. Gao, L.L. Wang, AI-enhanced large-eddy simulation for urban wind flows: accelerating spin-up with fourier neural operator initialization, *Build. Environ.* (2025) 114046.
- [37] Y. Chen, D. Wang, D. Feng, G. Tian, V. Gupta, R. Cao, M. Wan, S. Chen, Three-dimensional spatiotemporal wind field reconstruction based on LiDAR and multi-scale PINN, *Appl. Energy* 377 (2025) 124577.
- [38] H.M. Gråberg, Towards Physics-Informed Neural Networks for Urban Wind Flow Prediction, NTNU, 2022. Master's thesis.
- [39] Y. Yu, P. Li, D. Huang, A. Sharma, Street-level temperature estimation using graph neural networks: performance, feature embedding and interpretability, *Urban Clim.* 56 (2024) 102003.
- [40] Z. Liu, S. Zhang, X. Shao, Z. Wu, Accurate and efficient urban wind prediction at city-scale with memory-scalable graph neural network, *Sustain. Cities Soc.* 99 (2023) 104935.
- [41] X. Shao, Z. Liu, S. Zhang, Z. Zhao, C. Hu, PIGNN-CFD: a physics-informed graph neural network for rapid predicting urban wind field defined on unstructured mesh, *Build. Environ.* 232 (2023) 110056.
- [42] J. Milla-Val, C. Montañés, N. Fueyo, Adversarial image-to-image model to obtain highly detailed wind fields from mesoscale simulations in urban environments, *Build. Environ.* 266 (2024) 112123.
- [43] R. Zhao, S. Liu, J. Liu, N. Jiang, Q. Chen, A two-stage CFD-GNN approach for efficient steady-state prediction of urban airflow and airborne contaminant dispersion, *Sustain. Cities Soc.* 112 (2024) 105607.
- [44] G. Calzolari, W. Liu, Accelerating large eddy simulations of urban airflow with generative adversarial networks, *Build. Environ.* (2025) 113622.
- [45] G. Calzolari, W. Liu, Deep learning to replace, improve, or aid CFD analysis in built environment applications: a review, *Build. Environ.* 206 (2021) 108315.
- [46] T. Chen, R. Li, X. Hu, B. Zhang, Y. Liu, L.L. Wang, N. Gao, Machine learning as CFD surrogate models for rapid prediction of building-related physical fields: a review of methods and state-of-the-art, *Build. Environ.* (2025) 113667.
- [47] J. Lu, W. Li, S. Hobeichi, S.A. Azad, N. Nazarian, Machine learning predicts pedestrian wind flow from urban morphology and prevailing wind direction, *Environ. Res. Lett.* 20 (5) (2025) 054006.
- [48] N. Nazarian, J. Lu, M. Lipson, S. Liu, M.A. Hart, S. Kraysenhoff, L. Blunn, A. Martilli, UrbanTALES: a comprehensive LES dataset for urban canopy layer turbulence analyses and parameterization development. eartharxiv.org, 2025.
- [49] OpenStreetMap contributors, 2017. Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org>.
- [50] O. Coceal, A. Dobre, T.G. Thomas, S.E. Belcher, Structure of turbulent flow over regular arrays of cubical roughness, *J. Fluid. Mech.* 589 (2007) 375–409.
- [51] M.J. Brown, R.E. Lawson, Comparison of Centerline Velocity Measurements Obtained Around 2D and 3D Building Arrays in a Wind Tunnel (No. LA-UR-01-4131), Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), 2001.
- [52] P. Viola, and M. Jones, 2001, December. Rapid object detection using a boosted cascade of simple features. In Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001 (Vol. 1, pp. 1-1). Ieee.
- [53] C.N. Silla Jr., A.A. Freitas, A survey of hierarchical classification across different application domains, *Data Min. Know. Disc.* 22 (1–2) (2011) 31–72, <https://doi.org/10.1007/s10618-010-0175-9>.
- [54] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, *Adv. Neural Inf. Process. Syst.* (2014) 27.
- [55] Mirza, M. and Osindero, S., 2014. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784.
- [56] P. Isola, J.Y. Zhu, T. Zhou, A.A. Efros, Image-to-image translation with conditional adversarial networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 1125–1134.
- [57] Demir, U. and Unal, G., 2018. Patch-based image inpainting with generative adversarial networks. *arXiv preprint arXiv:1803.07422*.
- [58] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution. *European Conference on Computer Vision*, Springer International Publishing, Cham, 2016, pp. 694–711.
- [59] M.Z. Yousif, L. Yu, H.C. Lim, Super-resolution reconstruction of turbulent flow fields at various Reynolds numbers based on generative adversarial networks, *Phys. Fluids* 34 (1) (2022).
- [60] J. Xu, Z. Li, B. Du, M. Zhang, J. Liu, Reluplex made more practical: leaky ReLU, in: 2020 IEEE Symposium on Computers and Communications (ISCC), IEEE, 2020, pp. 1–7.