Physics of Fluids

RESEARCH ARTICLE | DECEMBER 24 2025

Wake reconstruction of a building using sparse data of virtual Doppler wind lidar

Haoyu Yang (杨浩宇) ; Jinlong Yuan (袁金龙) ; Yixiang Chen (陈逸翔) ; Haiyun Xia (夏海云); Xu Zhou (周旭) ; Tianwen Wei (魏天问); Yuanhao Gu (顾元豪); Yiming Zhang (张一鸣); Hao Wang (王浩)

Physics of Fluids 37, 125187 (2025)
https://doi.org/10.1063/5.0297817

Physics of Fluids ARTICLE pubs.aip.org/aip/pof

Wake reconstruction of a building using sparse data of virtual Doppler wind lidar

Cite as: Phys. Fluids **37**, 125187 (2025); doi: 10.1063/5.0297817
Submitted: 21 August 2025 · Accepted: 5 December 2025 ·
Published Online: 24 December 2025

Haoyu Yang (杨浩宇),[1,2] Jinlong Yuan (袁金龙),[1,2,3,a)] Yixiang Chen (陈逸翔),[1,2] Haiyun Xia (夏海云),[1,2,4] Xu Zhou (周旭),[3] Tianwen Wei (魏天问),[1,2] Yuanhao Gu (顾元豪),[1,2] Yiming Zhang (张一鸣),[5] and Hao Wang (王浩)[5]

AFFILIATIONS

[1]State Key Laboratory of Climate System Prediction and Risk Management, Nanjing University of Information Science and Technology, Nanjing 210044, China
[2]School of Atmospheric Physics, Nanjing University of Information Science and Technology, Nanjing 210044, China
[3]CMA Weather Modification Center, Beijing 100081, China
[4]School of Earth and Space Science, University of Science and Technology of China, Hefei 230026, China
[5]Key Laboratory of Concrete and Prestressed Concrete Structures of Ministry of Education, Southeast University, Nanjing 211189, China

[a)]Author to whom correspondence should be addressed: yuanjinlong@nuist.edu.cn

ABSTRACT

Understanding how wind flows in urban environments is vital for the low-altitude aviation economy, aviation safety, and building monitoring. However, even Doppler lidar, currently the most effective remote sensing tool for wind field measurement, faces challenges in urban environments due to insufficient spatial resolution and signal blockage by buildings. To solve these problems, we propose UrbanFlow-Dif, a novel artificial intelligence framework designed to reconstruct fine-scale wind information from sparse lidar data. The model learns from high-resolution simulations of airflow around buildings generated by computational fluid dynamics and uses a diffusion-based generative model, which can remove noise and recover detailed flow structures. Subsequently, the sensitivity of model performance with different scanning configurations is studied to reveal the influence mechanism of lidar position angle, spatial resolution, and scanning step angle. These results demonstrate that UrbanFlow-Dif successfully reconstructs the complex wake structures, achieving a root mean square error below 1.284 m/s even under sparsity conditions (finer than 30 m/1°). Ensemble averaging can further improve the reconstruction performance. This work provides a new foundation for fine-scale urban wind field detection and reconstruction.

I. INTRODUCTION

The fine detection of urban wind fields at high spatiotemporal resolution is critical for the low-altitude aviation economy,[1] aviation safety,[2–5] and building monitoring.[6–9] Coherent Doppler wind lidar, as a remote sensing instrument, actively emits pulsed laser beams and retrieves radial wind velocity by detecting the Doppler frequency shift of the backscattered signal from aerosol particles. It has the advantage of a flexible scanning mode and been widely used in atmospheric wind field detection.[10–13] However, the spatial resolution of existing lidar systems is insufficient to meet the requirements for capturing meter-scale turbulent characteristics within the urban building wake.[8] Although some pulse modulation methods have emerged in recent years to improve spatial resolution,[14,15] these high-resolution lidars face the challenge of pulse crosstalk, spectrum broadening, longer accumulation time and poor stability of the pulse coding system. Indiscriminately enhancing the range resolution can lead to a substantial increase in hardware costs and technical complexity.[16] In addition, the laser propagation is blocked by low-altitude buildings, resulting in a phenomenon known as beam blockage.[17] This results in a lack of detection information in the core region of wake. Therefore, for the detection of fine-scale building wake, lidar still faces the challenge related to sparse spatial sampling and beam blockage.

In recent years, the rapid advancement of deep learning methods has provided new opportunities for reconstructing building wake, owing to their strong capability to address highly nonlinear problems. In urban environments, background wind fields often exhibit fully developed turbulence characterized by high Reynolds numbers (Re). This complexity poses substantial challenges for effective flow field

reconstruction, especially when relying on sparse *in situ* measurements. However, most studies have focused on the wake with low Reynolds numbers (Re) and random sampling.[18,19] Data-driven frameworks have been developed for reconstructing flow fields from *in situ* measurements, exploiting the organized, viscosity-dominated characteristics of low Reynolds number flows that enable effective encoder–decoder architectures for accurate reconstruction.[20–23] A deep learning framework based on generative adversarial networks (GANs) is developed to reconstruct global turbulent wake from randomly distributed *in situ* measurements, demonstrating effective performance even with chaotic high Reynolds number flow behavior.[24] Although this model has been widely applied, it is constrained by its deterministic architecture, which is designed to produce a single optimal reconstruction for each set of input data during training.[25,26] However, data from *in situ* measurements are insufficient to fully constrain the entire physical system. In practice, it is difficult to obtain high-density *in situ* measurement data. This sampling method tends to capture the average state of the wind field, but is less effective at capturing turbulent characteristics in the building wake. As a result, reconstruction becomes a typical one-to-many problem, in which the same sparse input can yield multiple physically reasonable complete wind field solutions. GANs struggle to generate multiple solutions that comply with physical constraints.

Diffusion models (DMs) can capture the spatial distribution of atmospheric states, providing a foundation for inferring the current state from limited observation.[27,28] As a class of probabilistic generative models, DMs are suitable for learning complex, high-dimensional distributions—from precipitation[29] to earth's atmospheric states.[30] This makes them well-suited for capturing the intricate patterns of building wake. By using Markov chains to generate data samples, diffusion models avoid the unstable adversarial training required by GANs. DMs can serve as learned priors for Bayesian inference, where their probabilistic foundations enable principled uncertainty quantification. The key feature of DMs is their inherent stochasticity, which allows them to produce multiple plausible outputs from the same input, all consistent with the learned distribution. Urban wind field reconstruction depends on simulating *in situ* measurements (such as wind velocimeters).[31,32] This sampling method tends to capture the mesoscale and microscale wind field, but it is difficult to capture turbulent characteristics in the building wake. Even with DMs, the diversity of their solution sets may be limited by the information bottleneck inherent in the input data itself. Fortunately, lidar can provide more uniform and structured spatial wind field information. This makes it possible to reconstruct the fine wind field characteristics that existing methods cannot observe.

To address these challenges of insufficient spatial resolution and beam blockage, we propose UrbanFlow-Dif, a novel artificial intelligence framework that reconstructs fine-scale wind fields from sparse lidar data. Trained on high-resolution CFD simulations of building airflows, the model employs a diffusion-based generative approach that learns to infer detailed flow structures from noisy, incomplete observations. First, the wind field of building wake flows are simulated using computational fluid dynamics (CFD), ensuring the generated data are physically constrained and representative of realistic urban flow patterns with a high Re number. Second, the virtual lidar scanning process is applied to project these high-fidelity CFD fields into a sparse set of observations for radial wind fields. Datasets of sparse inputs of lidar are built, corresponding to ground truth fields. Finally, a framework leveraging a conditional diffusion model named UrbanFlow-Dif is demonstrated. This study is organized as follows. Section II describes the generation of the datasets based on CFD simulations and virtual lidar. Section III presents the UrbanFlow-Dif framework, including the conditional diffusion model architecture and training configuration. Section IV evaluates the model's performance and uncertainty, as well as its influence factor. A conclusion is drawn in Sec. V.

II. DATA

This section details the generation of the dataset used for training and evaluating the UrbanFlow-Dif model. The research framework is illustrated in Fig. 1. In Step 1, high-fidelity wind fields of building wakes are simulated using large eddy simulation (LES). In Step 2, radial velocity fields are calculated from eight lidar position angles, and 1600 snapshots are generated. In Step 3, training and testing sets are constructed using random sampling and structured sampling with beam blockage, respectively. In Step 4, the UrbanFlow-Dif model is trained to reconstruct wind fields from sparse inputs. In Step 5, model performance is evaluated using quantitative metrics on the testing set.

A. CFD simulation of building wakes

The data are derived from a high-fidelity CFD simulation of the wake around a single building with dimensions 10 m (length) × 10 m (width) × 50 m (height). The characteristic dimension D is defined as 10 m, representing the building width. The simulations are performed using LES, a technique well-suited for capturing the transient and turbulent characteristics of urban wind fields. The computational domain is defined as a rectangular volume of 500 (length) × 500 (width) × 200 m (height). The simulations are configured to model a neutrally stratified atmospheric boundary layer flow. The reference speed for the boundary condition of the velocity inlet is set to 10 m/s, with a Reynolds number of $\mathrm{Re} \approx 6.7 \times 10^5$. This setup ensures the development of a complex, turbulent wake downstream of the obstacle, which serves as the primary phenomenon for the reconstruction task. To accurately capture the wake structures while maintaining computational efficiency, a non-uniform grid was employed for the simulation, with a fine resolution of 0.4 m in the core region of the wake.

The two-dimensional slices of three-dimensional wind fields are extracted at a height of $z = 15$ m. This height is selected as a representative level that is sufficiently high to mitigate ground surface friction effects while still capturing the primary wake structure. A square domain of 256 × 256 m is selected. This domain is mapped onto a grid with a spatial resolution of 1 m, resulting in a 256 × 256 grid using cubic interpolation. This resolution is fine enough to adequately resolve the key turbulent structures within the wake of the 10 m building, ensuring physical fidelity. Second, it produces a uniform 256 × 256 pixel grid, which serves as an ideal input format for our model. Each grid contains the three Cartesian velocity components (u, v, w) on this grid, representing an instantaneous, high-fidelity snapshot of the building wake. From the continuous time-series data generated by the CFD simulation, 200 distinct snapshots were extracted at different time steps to ensure a diverse representation of the flow evolution.

B. Radial velocity calculation

With the snapshot of the building wake established, the next step is to simulate the physical process of lidar detection. This is accomplished by projecting the three-dimensional vector fields onto the

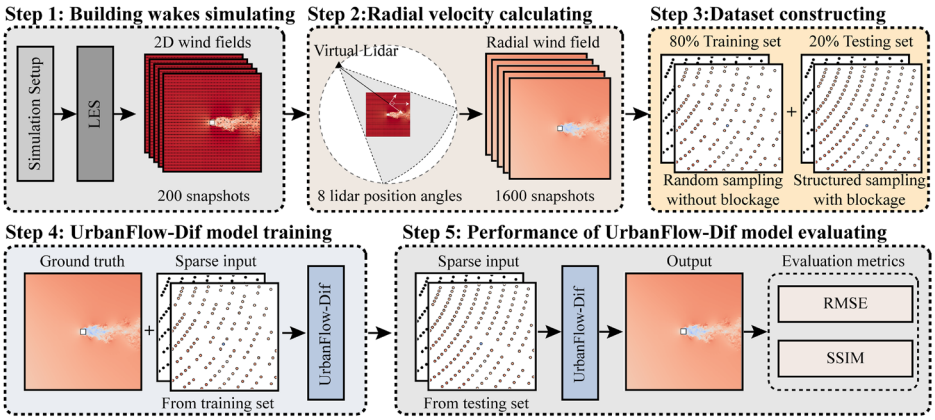

FIG. 1. Schematic of the research framework for wake reconstruction. Step 1: high-fidelity building wake fields are generated using large eddy simulation (LES). Step 2: radial velocity fields from eight viewing angles generate 1600 snapshots. Step 3: training and test datasets are constructed via random sampling (no blockage) and structured sampling (with blockage). Step 4: UrbanFlow-Dif learns the mapping from sparse inputs to ground truth fields. Step 5: Performance is evaluated using RMSE and SSIM metrics.

lidar's radial velocity fields, as shown in Fig. 2. These fields then serve as the input for the subsequent sparse sampling procedure.

To simulate the observation, a virtual Doppler wind lidar is configured to perform a plan position indicator (PPI) scan, where the elevation angle is fixed at 0° to capture the horizontal wind field.[16] The lidar position is defined by its distance (R) from the domain's center and its lidar position angle (α) for the mean flow direction, as shown in Fig. 2(j). In real physical scenes, the installation position of lidar is usually fixed, and the wind direction often changes, leading to changes in the wake distribution and blockage level caused by buildings. Here, these scenarios are simulated by changing the lidar position angle relative to the fixed wake axis. The transformation of the lidar position angle is used to reflect the influence of direction variations in this work, which simplified the dataset. For this study, the distance is held constant at R = 400 m, while the lidar position angle is varied across eight discrete values: {0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°}. By varying lidar position angles, data are generated from multiple lidar position angles.

For each ground truth field and a given virtual lidar position, a corresponding radial velocity field (Vr) is computed. The radial velocity at any grid point is the scalar projection of the three-dimensional Cartesian velocity components (u, v, w). This is calculated using the following equation:

$$V_r = (u \cdot \sin(\varphi) + v \cdot \cos(\varphi)) \cdot \cos(\theta) + w \cdot \sin(\theta), \quad (1)$$

where u, v, and w are the zonal, meridional, and vertical wind components. φ is the azimuth angle and θ is the elevation angle of the virtual lidar observation. In the lidar coordinate convention, the azimuth angle φ is defined with 0° corresponding to true north, and scanning proceeds clockwise as shown in Fig. 2(j). The lidar position angle α describes where the lidar is physically located, while the azimuth angle φ describes the direction the lidar scans toward. And $\alpha = 0°$ corresponds to the lidar positioned on the right side of the building, while $\alpha = 180°$ corresponds to the left side. Radial velocity is defined as positive when directed away from the lidar and negative when directed toward the lidar. This projection replicates how a physical lidar instrument transforms the three-component wind vector into the radial wind fields. The refresh rate of lidar can be significantly increased by reducing pulse accumulation time for short-range detection. And emerging multi-beam systems like the transient wind camera can

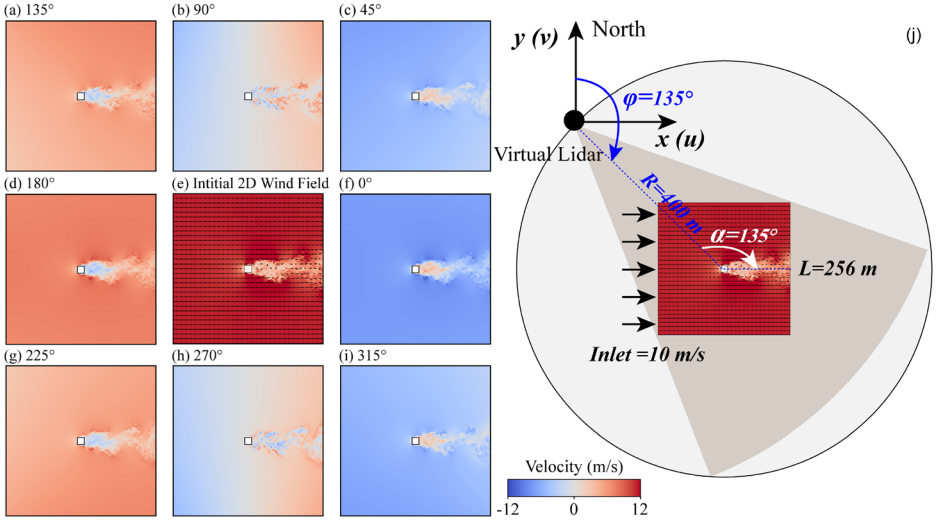

FIG. 2. Composite visualization of two-dimensional wind field and radial velocity components from multiple lidar position angles. (a)–(d) and (f)–(i) Radial velocity fields (V_r) were observed from virtual lidar positions with lidar position angles of 0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°, respectively. (e) Initial 2D wind field (u, v). (j) Lidar scanning geometry schematic showing lidar positioned at distance R = 400 m and lidar position angle (α) = 135°. The azimuth angle (φ) is defined with 0° corresponding to north, and the scan proceeds clockwise. u and v represent the zonal and meridional wind components, corresponding to azimuth angles of 90° and 0°, respectively. The inlet velocity is set to 10 m/s from the left side. All fields represent simulations over a 256 × 256 m domain.

acquire data almost simultaneously.[33] Therefore, the influence of scanning time is neglected in this study. The influence of this geometric projection is considerable, as demonstrated in Figs. 2(a)–2(d) and 2(f)–2(i), modifying the lidar position angle leads to dramatic changes in the resulting radial velocity patterns.

When $u = 10$, $v = 0$ m/s, and $\theta = 0°$, Eq. (1) simplifies to $Vr = 10 \cdot \sin\varphi$ (m/s). The sign and magnitude of the background radial velocity field are determined by the sin φ term, which varies with the azimuth angle. For example, when the lidar position angle is 0°, the sector PPI scan encompasses an azimuth range of $\varphi = 270° \pm 25.2°$ (based on the 256 m domain width at R = 400 m), corresponding to $\varphi \in [244.8°$ and $295.2°]$. The term of sin φ remains consistently negative across the entire scan sector, resulting in negative background radial velocities, as observed in Fig. 2(f). When the lidar position angle is 90°, the azimuth range shifts to $\varphi = 180° \pm 25.2°$, corresponding to $\varphi \in [154.8°$ and $205.2°]$. In this configuration, sin φ transitions from approximately −0.43–0.43, causing radial velocities to transition smoothly from negative values on the left side to positive values on the right side, as shown in Fig. 2(b).

C. Dataset construction

This section describes the construction of sparse input datasets from radial velocity fields. The 200 snapshots are divided into training and test datasets, comprising 80% and 20%, respectively. The training dataset is designed to compel the model to learn the underlying flow physics, independent of beam blockage. Random sampling is employed to generate sparse data points from Vr fields, creating an idealized reconstruction task that avoids overfitting to specific scanning patterns. For each training snapshot, Vr fields of eight different lidar position angles are generated. Beam blockage is ignored to ensure fully observable fine-scale wake characteristics. The training set is generated by randomly sampling 100, 500, and 1000 points to enhance model robustness across varying data sparsity. The sampling procedure is repeated five times for each Vr field, yielding 6400 unique training pairs. This approach enables flow field reconstruction from sparse data without bias from specific scanning patterns or beam blockage. By training on randomly sampling data, the model develops the ability to infer complete flow physics from limited observations, improving its generalization when applied to lidar scanning configurations with beam blockage.

The test dataset is designed to evaluate model performance under realistic conditions by virtual lidar scanning configurations. For each test snapshot, two key lidar parameters are systematically varied: radial spatial resolution and scanning step angle. Each parameter combination defines a unique scanning configuration that produces distinct and structured sparse inputs. For each configuration, the test dataset includes 40 unique pairs. During scanning, beam blockage effects are incorporated when scan lines intersect building obstacles. Larger radial spatial resolution and scanning step angle result in sparser sampling points with beam blockage, as shown in Fig. 3.

Each sparse input comprises two channels, both containing 256×256 matrices representing grid points in the horizontal and vertical directions of the radial wind field, with a grid resolution of 1 m. The first channel contains the radial wind velocity at sampling points, while the second channel provides a position mask for the sampling points. This second channel is designed to effectively utilize and highlight sensor location information. This approach provides flexibility in sensor configuration and directly encodes sensor locations into the model through their spatial coordinates.

III. METHODOLOGY

This section presents the methodology for reconstructing the fine-scale wake. We first introduce the architecture of our proposed framework, UrbanFlow-Dif, which is based on a conditional denoising diffusion probabilistic model (DDPM). Next, we define the evaluation metrics used to quantitatively assess the reconstruction performance. Finally, we describe the specific training configuration and hyperparameters employed in this study.

A. Denoising diffusion probabilistic model

For this study, we selected the denoising diffusion probabilistic model (DDPM) architecture, prioritizing maximal reconstruction fidelity over the faster sampling speeds offered by variants, such as DDIM. The robustness and proven performance of the DDPM in generating high-quality, physically plausible results make it the most suitable framework for the scientific task of precise flow field reconstruction.

The UrbanFlow-Dif, based on the DDPM, is proposed to reconstruct complete radial wind fields from sparse lidar data. The DDPM framework, first introduced by Ho *et al.*,[34] comprises a fixed forward and learnable reverse process. In the forward process, Gaussian noise is progressively added to the radial wind field, transforming it into unstructured pure noise. Then, a reverse process is learned to iteratively remove the noise and generate complete wind fields from sparse lidar data, as illustrated in Fig. 4.

The fixed forward process is defined as a Markov chain that progressively adds a small amount of Gaussian noise to the data over T

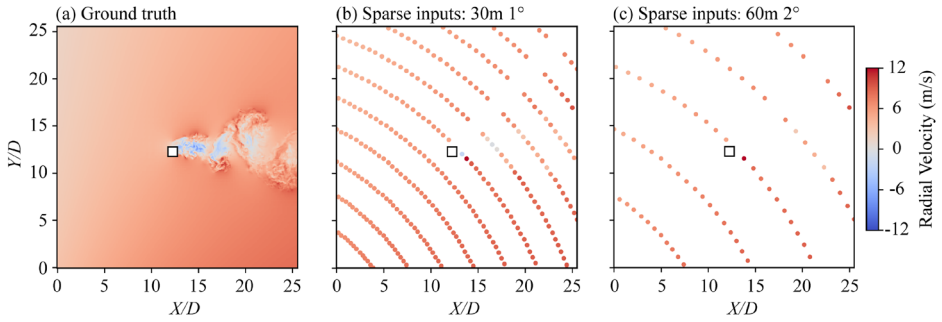

FIG. 3. Visualization of lidar scanning simulation, illustrating the impact of scanning configuration and beam blockage. (a) Ground truth radial velocity field in a resolution of 1 m. (b) and (c) Resulting structured sparse grids generated with two different scanning configurations: (b) fine resolution with 30 m radial spatial resolution and 1° scanning step angle and (c) coarse resolution with 60 m radial spatial resolution and 2° scanning step angle.

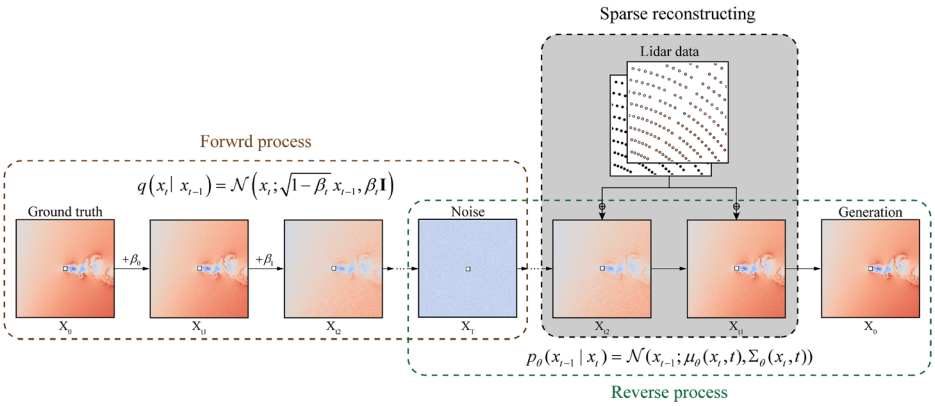

FIG. 4. Schematic diagram of the UrbanFlow-Dif model for sparse reconstruction of building wake based on the DDPM framework. The forward process progressively transforms a radial wind field (X_0) into an isotropic Gaussian noise sample (X_T) over T timesteps. The reverse process starts from X_T and is conditionally guided by sparse lidar data to iteratively denoise and reconstruct the complete radial wind field. In the equations, q denotes the fixed forward, while noising process, and p_θ represents the learnable reverse denoising process parameterized by a neural network with parameters θ. x_t is the data sample at time step t, β_t is the pre-defined noise variance at that step, and I is the identity matrix. The terms μ_ϑ and $\sum_\vartheta$ represent the mean and covariance of the reverse transition, which are predicted by the neural network.

discrete timesteps. At each time step t, q represents the probability distribution for progressively adding noise to the data[34]

$$q(x_t|x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1-\beta_t}\,x_{t-1}, \beta_t \mathbf{I}\right), \tag{2}$$

where x_t is the noisy data sample at time step t, and $\beta_t \in (0, 1)$ is a predefined, small positive constant representing the noise variance scheduled for that step. The term $\mathcal{N}$ denotes a Gaussian distribution, and $\mathbf{I}$ is the identity matrix.

This can be further simplified by directly sampling x_t from x_0[34]

$$q(x_t|x_0) = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, \epsilon \sim \mathcal{N}(0,1), \tag{3}$$

where $\alpha_t = 1 - \beta_t$, and $\bar{\alpha}_t = \prod_{t=1}^{t} a_i$, ϵ is a standard Gaussian noise.

The reverse process is approximated with a learned Gaussian transition parameterized by a neural network ϑ. This model is trained to predict the less noisy sample x_{t-1} from the noisier sample x_t[34]

$$p_\vartheta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\vartheta(x_t,t), \Sigma_\vartheta(x_t,t)), \tag{4}$$

where the μ_ϑ and $\sum_\vartheta$ are predicted by a neural network based on x_t and t. The full generative process starts with a sample x_T from the Gaussian prior and iteratively applies this learned denoising step for $t = T, T-1, \ldots, 1$ to generate a clean sample x_0. The function approximator for the mean, $\mu_\vartheta(x_t, t)$, is a standard U-Net architecture conditional on the time step t and the sparse lidar data.[34,35] The time step t is incorporated via sinusoidal positional embeddings, while the sparse lidar data are concatenated with the noisy input x_t along the channel dimension to guide the denoising process. Following DDPM parameterization, the training loss $\mathcal{L}(\vartheta)$ simplifies to a mean-squared error between the ground truth noise and the predicted noise at time step t[34]

$$\mathcal{L}(\vartheta) = E_{t,x_0,\epsilon}\left[\|\epsilon - \epsilon_\vartheta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, t)\|^2\right], \tag{5}$$

where a ground truth x_0, a time step t and standard Gaussian noise ϵ are randomly selected, and x_t is obtained through Eq. (4). x_t and t are fed into the U-Net architecture to predict the noise ϵ_ϑ.. Then, the mean squared error between the ground truth noise ϵ and the predicted noise ϵ_ϑ is calculated. The objective is to minimize this error by optimizing the neural network parameters ϑ.

B. Evaluation metrics

In order to quantitatively assess the quality of fine-scale wind characteristics, the root mean square error (RMSE) and structural similarity index measure (SSIM) are introduced. RMSE measures the discrepancy between generated radial wind fields and ground truth, serving as a reconstruction accuracy metric defined as

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(v_i^{output} - v_i^{real}\right)^2}, \tag{6}$$

where v_i^{output} and v_i^{real} represent reconstructed radial velocities and ground truth of grid point i, n is the total number of grid points.

In addition to RMSE, the SSIM is employed as a complementary metric. While RMSE is sensitive to absolute velocity magnitudes, SSIM focuses on structural similarity and pattern fidelity, making it particularly valuable for comparing reconstruction quality across scenarios with different background wind speeds. The SSIM incorporates factors, such as luminance, contrast, and structure, aiming to compare the structural similarity. The SSIM is given by

$$SSIM(x,y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}, \tag{7}$$

where μ_x, μ_y are the average pixel values and σ_x, σ_y, and σ_{xy} are the variances and covariance of the image. c_1 and c_2 are constants introduced to avoid division by zero. An SSIM value closer to 1 indicates higher structural similarity between the reconstructed and the ground truth, while a lower RMSE value signifies smaller error.

C. Training configuration

The UrbanFlow-Dif is implemented in PyTorch. Training utilizes randomly sampled data, while evaluation employs test sets with virtual

lidar scanning patterns. All radial velocity fields are processed at 256×256 pixels resolution, with the training dataset comprising 6400 unique pairs. The diffusion process employs $T = 1000$ time-steps with a linear β schedule increasing from 0.0001 to 0.02. Training is conducted for 400 epochs with a batch size of 64. The Adam optimizer is used with a learning rate of 0.0001 and momentum parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$).[36] All experiments are performed on a NVIDIA A6000 GPU.

IV. RESULTS

This section evaluates the UrbanFlow-Dif model's performance across a range of simulated lidar scanning configurations and lidar position angles. A sensitivity analysis is then conducted to evaluate the model's robustness and generalizability across the entire test dataset. Finally, the model's ensemble generative capabilities for uncertainty quantification are explored.

A. Reconstruction performance under different lidar position angles

The performance of the UrbanFlow-Dif model is qualitatively assessed under the different lidar position angles with a series of scanning configurations. Three representative lidar position angles (0°, 90°, and 180°) are selected, each resulting in a distinctly different beam blockage pattern. For each lidar position angle, a series of scanning configurations with varying levels of sparsity is applied. This comprehensive setup is applied to examine the model's ability to reconstruct complex flow structures from sparse data.

For the upstream lidar position angle of 0°, the model's performance on two different ground truth cases is presented in Fig. 5. The results for two typical flow field patterns (cases 1 and 2), are shown for four different data sparsity levels, achieved by adjusting the radial spatial resolution from 5 to 30 m and the scanning step angle from 0.5° to 1.0°. Each case shows three rows: inputs, outputs, and the difference between outputs and ground truth. Inputs are generated through structured sampling from the ground truth based on predefined scanning parameters, including the scanning step angle and radial spatial resolution. Regions where beam blockage occurs in the input are filled with white sectors. The building is marked by a black square. Performance within the core region of wake in the white dashed square is quantified using RMSE and SSIM. In the dense configurations of 5 m/0.5°, where inputs contained wake characteristics, outputs are improved through sharpened fine-scale patterns and restored spatial continuity, which shows high consistency with the ground truth. This

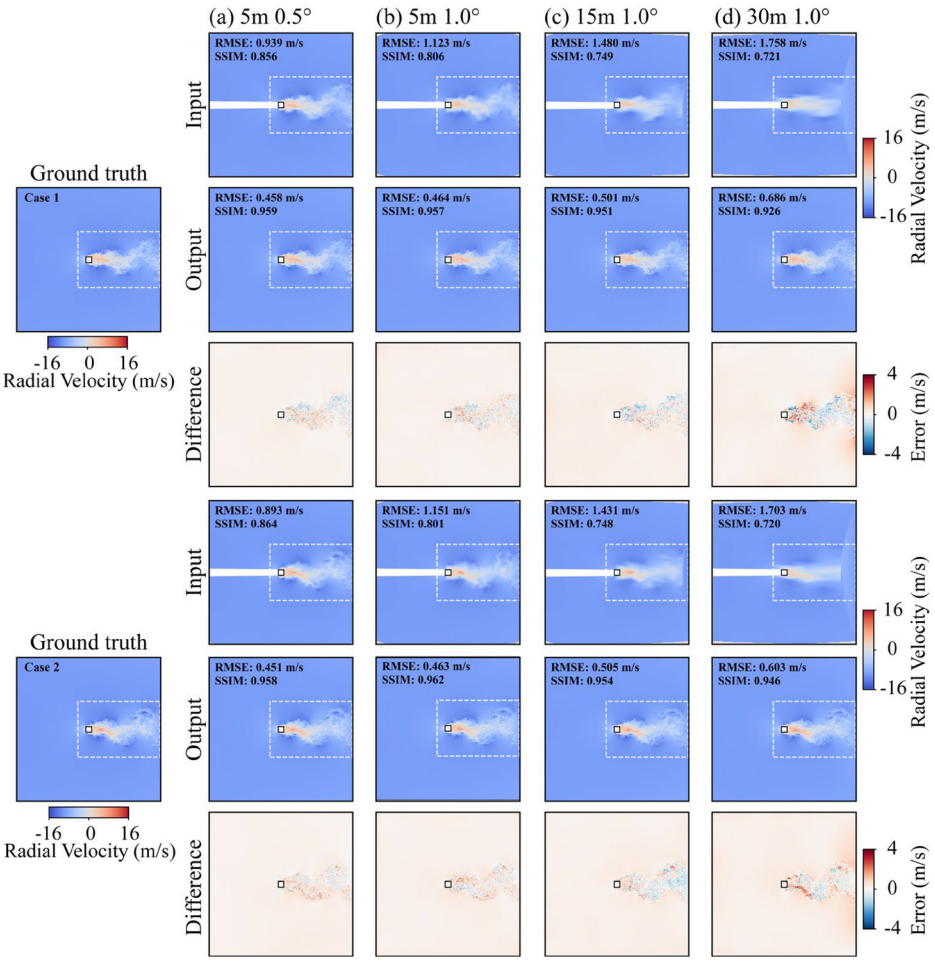

FIG. 5. UrbanFlow-Dif reconstruction performance for upstream lidar position angle of 0°. Two ground truth cases (top/bottom panels) are tested under four sparsity levels (a)–(d) by varying spatial and angular resolutions. Each case shows input interpolated from sparse data, reconstructed output, and difference field (ground truth-output). RMSE and SSIM values are provided for input and output fields, respectively. White regions indicate building locations.

improvement is quantitatively confirmed by the reduction in RMSE from 0.939 to 0.458 m/s for case 1 and from 0.893 to 0.427 m/s for case 2. Conversely, in the sparse configurations of 30 m/1°, where inputs suffered from severe blurring, the model can still reconstruct the dominant wake pattern. For instance, in case 2, RMSE decreased from 1.703 to 0.603 m/s, while SSIM improved from 0.720 to 0.946. In general, the difference field increases as the input becomes sparser.

Figure 6 presents reconstruction results for the side-on view of 90°. To further test the model's ability to handle strong background velocity gradients. This scenario evaluates the model's ability to separate building airwake from the large-scale background flow. In the dense configurations of 5 m/0.5°, the model sharpened the partial wake characteristics captured in the inputs, which led to RMSE reductions of 0.446 and 0.377 m/s for the two cases, respectively. In the sparse configurations of 30 m/1°, although the inputs lost detail, a partial reconstruction of the wake was achieved, maintaining the radial velocity dipole structure. The difference fields in the background wind regions remained near zero.

The model's robustness in the downstream lidar position angle of 180°, with maximum beam blockage, is shown in Fig. 7. In this scenario, the core region of airwake could not be directly sampled due to the beam blockage. However, fine-scale wake characteristics are reconstructed across all scanning configurations. In the dense configurations of 5 m/0.5°, fine-scale airwake characteristics are well reconstructed based on the edge information of beam blockage from inputs. It is quantitatively confirmed by the reduction in RMSE from 1.229 to 0.485 m/s for case 1 and from 1.128 to 0.487 m/s for case 2. Even in the configuration of 30 m/1.0°, minimal information about the core region of the wake is provided, thus increasing the uncertainty of outputs. However, the model can still generate reliable wake patterns consistent with the surrounding flow.

B. Model robustness and generalizability

To evaluate the model's robustness and generalizability, a sensitivity analysis is performed across the entire test dataset. Figure 8 presents a quantitative analysis of model reconstruction performance. The spatial distribution of RMSE between reconstructions and ground truth is shown in Fig. 8(a), averaged across the test dataset. Under identical lidar position angles, sparser sampling configurations result in higher RMSE values in the core region of wakes, while background wind field RMSE outside this region approaches zero.

A two-dimensional Fourier transform is applied to both the ground truth and the generated outputs. The radially averaged power

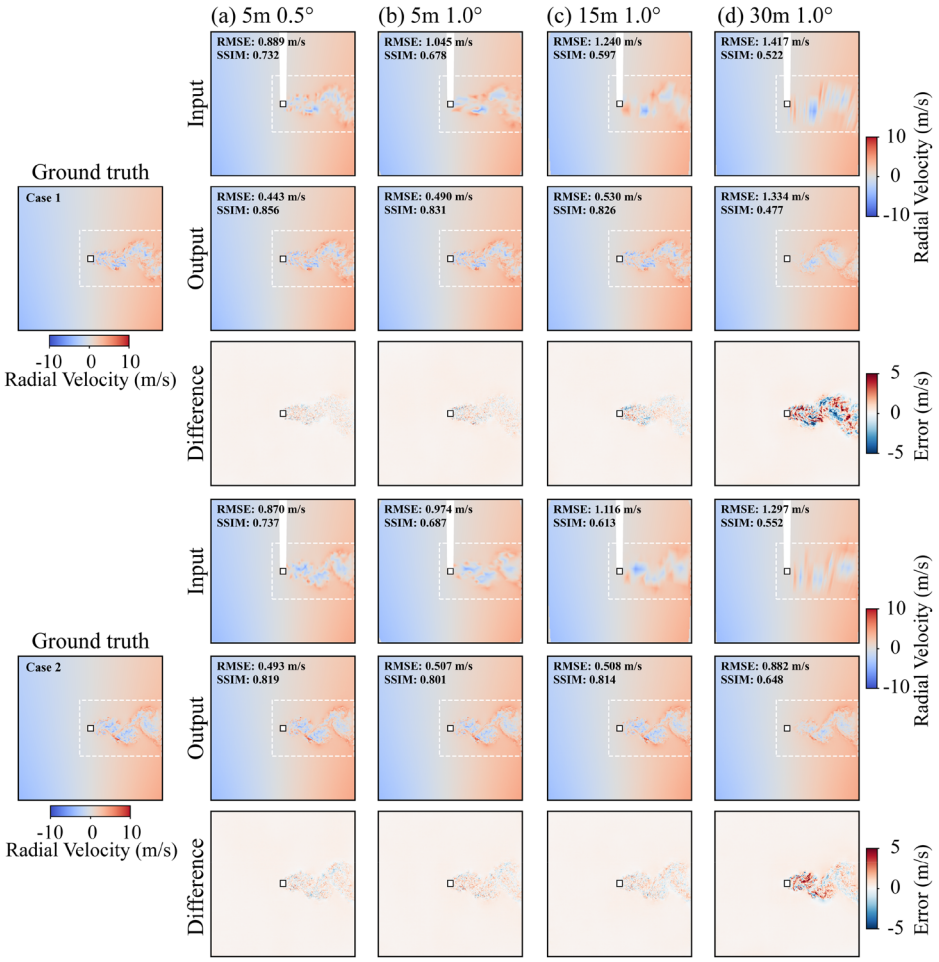

FIG. 6. UrbanFlow-Dif reconstruction performance for side-on lidar position angle of 90°. Two ground truth cases (top/bottom panels) are tested under four sparsity levels (a)–(d) by varying spatial and angular resolutions. Each case shows an input interpolated from sparse data, reconstructed output, and difference field (ground truth-output). RMSE and SSIM values are provided for input and output fields, respectively. White regions indicate building locations.

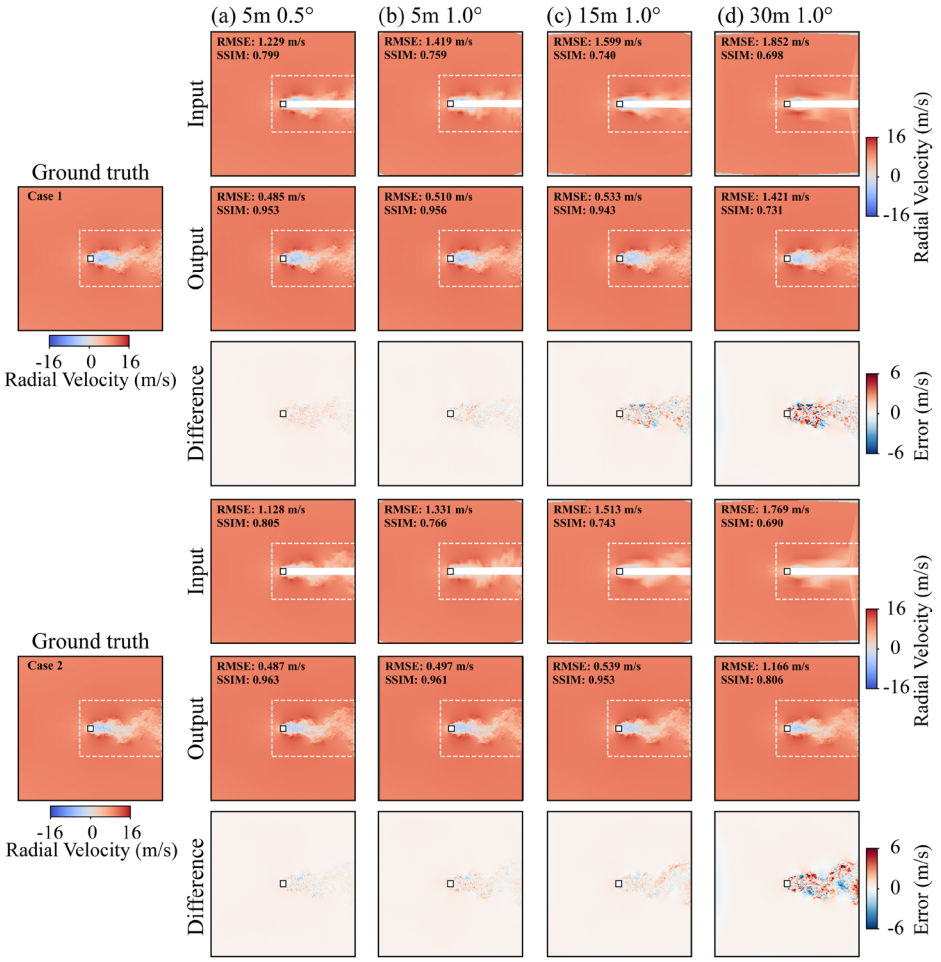

FIG. 7. UrbanFlow-Dif reconstruction performance for the downstream view of 180°. Two ground truth cases (top/bottom panels) are tested under four sparsity levels (a)–(d) by varying spatial and angular resolutions. Each case shows input interpolated from sparse data, reconstructed output, and difference field (ground truth-output). RMSE and SSIM values are provided for input and output fields, respectively. White regions indicate building locations.

spectral density (RAPSD), as termed by Harris *et al.* in their analysis of SR precipitation fields,[37] is computed to quantify the energy distribution across spatial scales, from large-scale flow to fine-scale wakes. The results are shown in Fig. 8(b). When the wave number is below 0.3 rad/m, the RAPSD curve of the outputs is consistent with ground truth at lidar position angles of 0° and 180°, indicating that such a scale of wind field can still be reconstructed using configurations of 5/0.5° and 60 m/1°. For the wave number ranged from 0.3 to 3 rad/m, corresponding to spatial scales of 2.1–20.9 m, the outputs of 60 m/1° underestimate spectral energy. When the wave number exceeds 3 rad/m, the RAPSD curves at lidar position angles of 0° and 180° start showing underestimation even under the dense configuration. Excessively sparse sampling configurations (e.g., 60 m/1°) are insufficient to reconstruct fine wake characteristics, but the main flow pattern can still be restored.

Figure 9 presents a quantitative evaluation of model performance with the core region of wake using average SSIM and RMSE. Reconstruction performance, as measured by higher SSIM and lower RMSE, shows a strong positive correlation with the sampling density. For instance, at the lidar position angle of 180°, under configurations finer than 30 m and 1°, the RMSE is below 1.167 m/s and SSIM exceeded 0.828. However, under the sparsest configuration of 60 m and 3°, the RMSE value reaches 3.506 m/s.

C. Ensemble generation for uncertainty quantification

The model's ensemble generative capabilities are tested for the 180° lidar position angle at varying lidar distances (R = 200, 400, and 1000 m), as shown in Fig. 10. It is worth noting that although R is different, only projections within the sub-domain of 256 × 256 m are considered. The ground truth is displayed in the first row for comparison. For each sparse input, 100 independent reconstructions are generated to form an ensemble average output. The ensemble average output, representing a more stable and physically robust estimate, is displayed in the second row. The third row shows the ensemble standard deviation of the difference between ensemble-averaged results and single reconstructions. Performance within the core wakes region in the white dashed square is quantified using RMSE and SSIM. In the dense configuration of 5 m/0.5°, ensemble average output errors increase with increasing R, where RMSE is elevated from 0.383 to 0.414 m/s, and SSIM is reduced from 0.979 to 0.972. The sparser configuration of 15 m/1° exhibits pronounced error increases, with RMSE

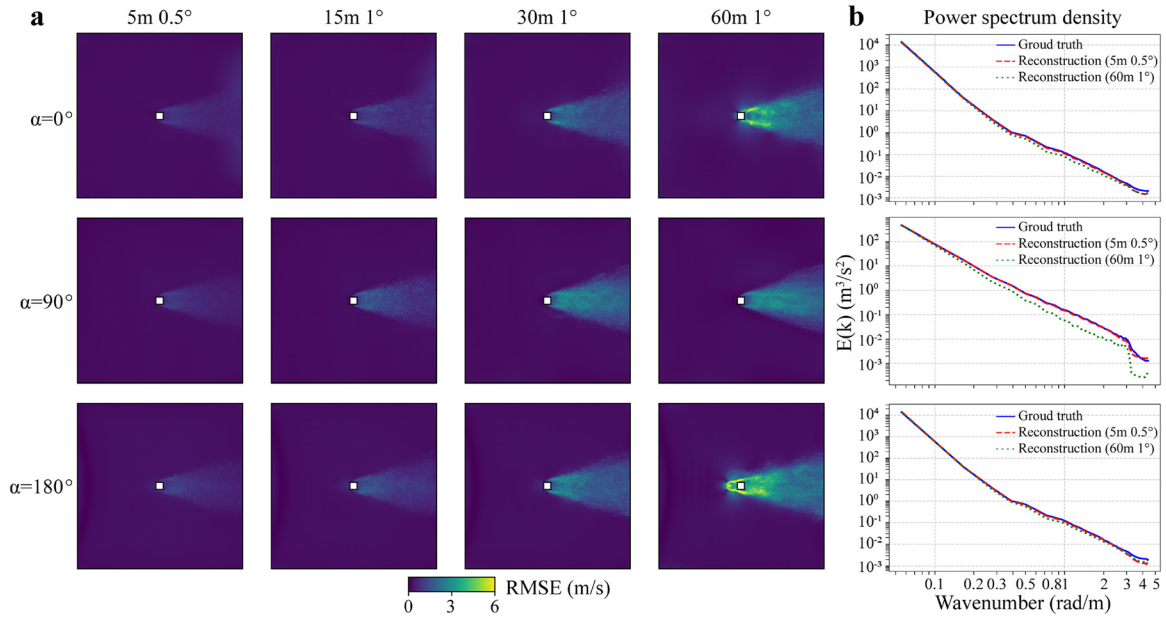

FIG. 8. Quantitative analysis of reconstruction performance. (a) Spatial distribution of root mean square error (RMSE) between reconstructions and ground truth for three representative lidar position angles of 0°, 90°, and 180°. Each column represents a different scanning configuration, indicating varying sampling density. (b) Radially averaged power spectral density (RAPSD) analysis comparing reconstruction results with ground truth under two configurations (5 m 0.5° and 60 m 1°). This quantifies the energy distribution across spatial scales, from large structures (low wave numbers) to fine-scale wakes (high wave numbers).

rising from 0.416 to 0.710 m/s and SSIM declining from 0.972 to 0.939. At R = 1000 m, the ensemble average output of the sparser configuration exhibits fine-scale wake distortion, and the ensemble standard deviation reveals extensive high-uncertainty regions approaching 2.8 m/s. However, the primary pattern of wake can still be reconstructed with SSIM reaching 0.939. Increasing distance results in reduced tangential spatial resolution, making it difficult to capture wind field information at the edges of beam blockage.

A direct comparison between single and ensemble-averaged reconstructions is detailed in Table I. For each sparse input, 100 single reconstructions are generated to form an ensemble average output. The results show that ensemble reconstruction consistently reduced errors compared to single reconstructions, with average RMSE being reduced by 3.48%–22.9% relative to single reconstruction and average SSIM being improved by 1.29%–8.22%. Under dense configurations of 5 m/0.5°, the average RMSE is decreased from 0.507 to 0.435 m/s through ensemble reconstruction, with an improvement of 14.2% being achieved, while the average SSIM is increased from 0.961 to 0.973. Under sparse configurations of 60 m/3°, the average RMSE is reduced from 3.506 to 3.384 m/s, with a little improvement of 3.48%.

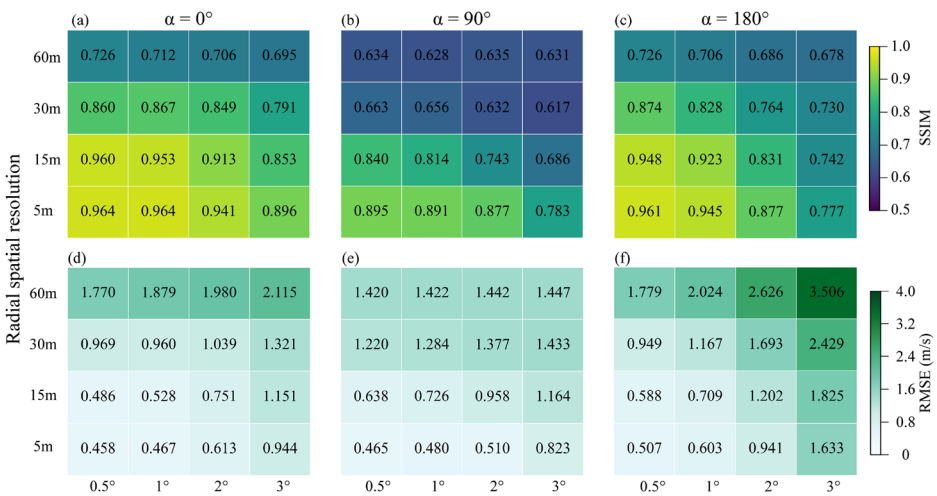

FIG. 9. Sensitivity analysis of the model reconstruction performance under different scanning parameters and lidar position angles. All evaluations were performed within the core region of wakes. Performance was quantified using two metrics: the average structural similarity index (SSIM) and the average root mean square error (RMSE).

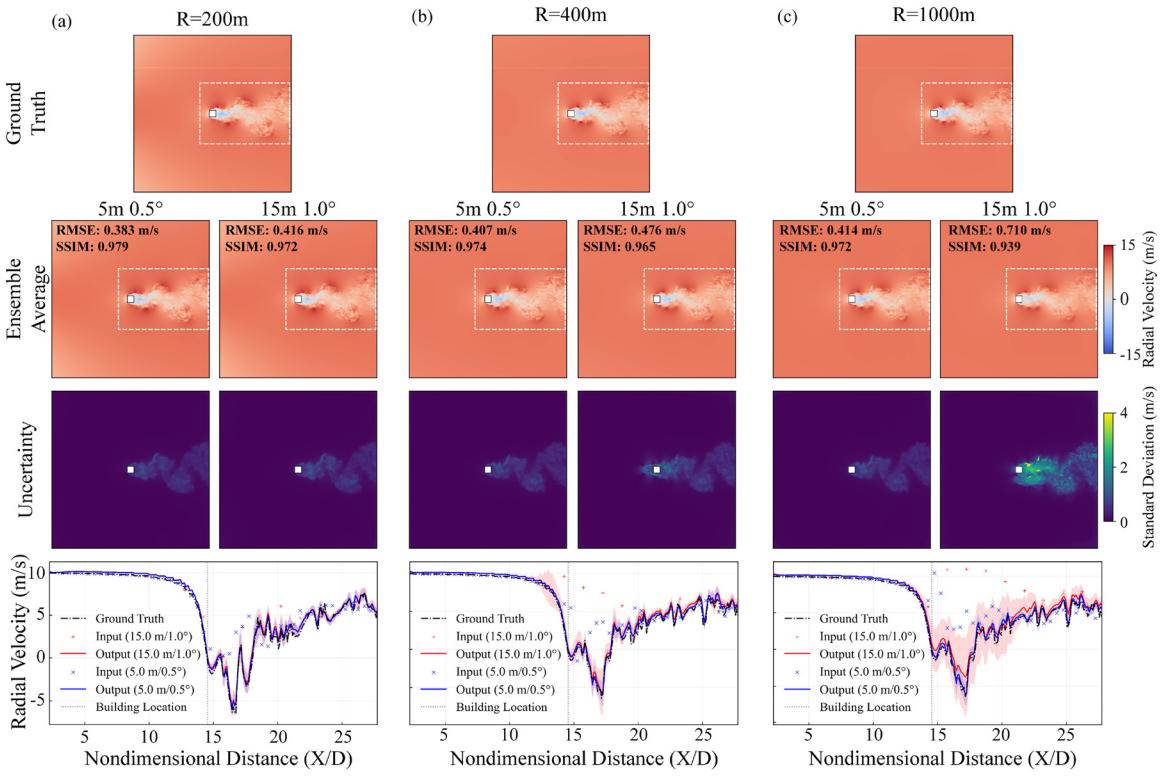

FIG. 10. Ensemble reconstruction and uncertainty quantification for fully blocked wake (lidar position angle = 180°). From top to bottom row: ground truth, ensemble average (100 reconstructions), ensemble standard deviation (uncertainty), and radial velocity profiles with 95% confidence intervals along horizontal cut (gray dashed line). Core wake performance (white dashed box) quantified by RMSE and SSIM at R = 200, 400, and 1000 m. D = 10 m represents the building width.

TABLE I. Performance comparison of ensemble-averaged reconstruction compared to single reconstruction under the lidar position angle of 180°. Boldface denotes the results of the ensemble-averaged reconstruction, and underlined values indicate the relative change (RC) compared to the single reconstruction.

Scanning configuration	RMSE (m/s)				SSIM			
	Input	Single	Ensemble	RC	Input	Single	Ensemble	RC
5 m/0.5°	1.209	0.507	**0.435**	−14.2%↓	0.822	0.961	**0.973**	+1.29%↑
15 m/1°	1.534	0.709	**0.586**	−17.3%↓	0.772	0.923	**0.946**	+2.49%↑
30 m/2°	3.130	1.693	**1.304**	−22.9%↓	0.706	0.764	**0.809**	+8.22%↑
60 m/3°	4.871	3.506	**3.384**	−3.48%↓	0.559	0.678	**0.689**	+1.62%↑

V. DISCUSSION

This study demonstrates the efficacy of the proposed UrbanFlow-Dif model in reconstructing fine-scale building wakes from sparse and incomplete lidar measurements. Reconstruction quality is primarily governed by three interacting factors: sampling density, lidar position angle (including beam blockage and projection effects), and ensemble averaging.

A. The critical role of sampling density

The reconstruction performance exhibits a strong positive correlation with sampling density across all scenarios (Figs. 5–7 and 9). Dense configurations (e.g., 5 m/0.5°) generate good reconstructions with minimal error, validating the model's baseline capability. More significantly, UrbanFlow-Dif maintains strong performance even under sparse conditions. Under a sparse configuration of 30 m/1°, key wake characteristics and dominant flow patterns are successfully captured. Improvements in radial spatial resolution, scanning step angle, and tangential spatial resolution contribute to enhanced performance in reconstructing wake characteristics from edge information. This indicates that diffusion priors effectively compensate for missing spatial information when coarse-scale coherent structures are partially resolved.

B. Impact of lidar position angle and beam blockage

The lidar position angle introduces distinct challenges by changing radial wind projections and the extent of hard occlusions from the beam blockage. The upstream lidar position angle of 0° represented an ideal scenario with minimal blockage, allowing the model to effectively recover wake patterns. The side-on lidar position angle of 90° was characterized by strong background velocity gradients in the radial projection; the model's ability to isolate the localized wake from this large-scale background flow demonstrates its robustness. The downstream lidar position angle of 180° presents the most severe challenge, as the building completely obscures the core wake region from the direct sampling. In this scenario, the model's task transitions from super-resolution to a complex physical inpainting problem. The model's ability to recover the general wake pattern demonstrates its capacity to leverage information at the edges of the blockage region to infer flow structures within the occluded zone. This capability to infer occluded flow based on learned physical priors represents a key advantage of the generative diffusion approach over traditional interpolation methods.

C. Probabilistic reconstruction and ensemble averaging

While single reconstructions generate satisfactory results, employing ensemble averaging further enhances performance, particularly in reducing error and enhancing physical robustness. As detailed in Table I, ensemble outputs improve wake reconstruction performance by consuming extra time, with RMSE reduced by 3.48%–22.9% and SSIM improved by 1.29%–8.22%. This improvement stems from the inherent stochastic nature of the diffusion model, which generates a distribution of physically plausible wind fields consistent with sparse constraints rather than a single estimate. The ensemble average, thus, represents a high-probability, stable estimate of the mean flow. However, ensemble-averaged results perform better; correspondingly, higher computational costs are incurred. Single reconstruction can be employed in operational scenarios with highly stringent timeliness requirements, while ensemble averaging methods can be utilized in applications where temporal constraints are more relaxed.

Furthermore, point-wise error metrics such as RMSE are inadequate for assessing reconstruction performance under sparse configurations. Sparse sampling inherently eliminates most wake texture structures from lidar detection; while RMSE is difficult to capture this structural loss, SSIM exhibits markedly reduced values that reflect the degradation.

D. Limitations and future work

Several limitations warrant discussion. First, the model is trained exclusively on a 15 m-height slice dataset for reconstruction. However, the proposed method is applicable to reconstructing wake slices at different heights with corresponding datasets. Second, RAPSD analysis [Fig. 8(b)] reveals a spectral bias: while large-scale energy is accurately recovered, reconstructions from sparse inputs underestimate spectral energy at high wave numbers (>0.3 rad/m). Under severe sparsity, the model does not perfectly reconstruct the turbulent energy of fine-scale building wakes. Improving fine-scale wake characteristics recovery remains a priority for future work. Second, this study focused on an isolated building; extending the model to realistic urban clusters represents an important next step. Third, the study assumed a fixed wind direction while varying lidar position angles to assess reconstruction robustness. It is important to note that changes in lidar position angle and wind direction are not equivalent: oblique wind angles change the building's effective width and modify the wake structure itself. While the current design isolated the model's performance under varying observational geometries, future work should incorporate wake data from multiple wind directions. This would enable evaluation of the model's ability to reconstruct different wake patterns from a fixed lidar position, reflecting another common operational scenario. Finally, we will extend the model to accommodate building cluster wake reconstruction and validate the approach through real-LIDAR observations.

VI. CONCLUSION

This study presents UrbanFlow-Dif, a novel diffusion-based model for reconstructing fine-scale building wakes from sparse wind lidar data, successfully addressing the inherent challenges of insufficient spatial resolution, beam blockage, and deterministic models. First, the density of sparse sampling determines the reconstruction quality. It is demonstrated that a practical configuration of 30 m and 1° can still capture key wake characteristics and flow patterns. This finding provides crucial guidance for optimizing scanning configurations within the spatial resolution limitations of existing lidar systems and minimum sampling requirements. Second, reconstruction performance is affected by variations in lidar position angles due to changes in beam blockage and radial wind field projections. Even at the most challenging downstream lidar position angle, where the wake core region is completely blocked, the general wake pattern can be effectively reconstructed by leveraging edge information. Finally, single reconstruction has achieved satisfactory results, we find that ensemble averaging improves wake reconstruction performance. Although this will consume extra time. The effective reconstruction of fine-scale wake characteristics from sparse wind lidar data provides a robust foundation for complex urban wind field detection, which is essential for aviation economy, aviation safety, and building monitoring.

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China (Grant No. 42405137), the Natural Science Foundation of Jiangsu Province (Grant No. BK20230434), and the Chinese Aeronautical Establishment (Grant No. 202300220R2001).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Haoyu Yang: Methodology (equal); Software (equal); Visualization (equal); Writing – original draft (lead). **Jinlong Yuan:** Conceptualization (lead); Funding acquisition (lead). **Yixiang Chen:** Data curation (lead). **Haiyun Xia:** Writing – review & editing (lead). **Xu Zhou:** Writing – review & editing (equal). **Tianwen Wei:** Writing – review & editing (equal). **Yuanhao Gu:** Data curation (equal). **Yiming Zhang:** Resources (equal); Writing – review & editing

(supporting). **Hao Wang:** Resources (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1]X. Huang, "The small-drone revolution is coming—Scientists need to ensure it will be safe," Nature **637**(8044), 29–30 (2025).
[2]H. Gao, C. Shen, X. Wang, P.-W. Chan, K.-K. Hon, and J. Li, "Interpretable semi-supervised clustering enables universal detection and intensity assessment of diverse aviation hazardous winds," Nat. Commun. **15**(1), 7347 (2024).
[3]P. W. Chan, "LIDAR-based turbulence intensity calculation using glide-path scans of the Doppler LIght Detection And Ranging (LIDAR) systems at the Hong Kong International Airport and comparison with flight data and a turbulence alerting system," Meteorol. Z. **19**(6), 549–563 (2010).
[4]T.-C. Wu and K. Hon, "Application of spectral decomposition of LIDAR-based headwind profiles in windshear detection at the Hong Kong International Airport," Meteorol. Z. **27**, 33–42 (2017).
[5]P.-W. Chan, P. Cheung, K.-K. Lai, J.-L. Xie, and Y.-Y. Leung, "Some interesting observations of cross-mountain east-to-southeasterly flow at Hong Kong International Airport and their numerical simulations," Atmosphere **16**(7), 810 (2025).
[6]E. Cheynet, J. B. Jakobsen, J. Snaebjörnsson, J. Reuder, V. Kumer, and B. Svardal, "Assessing the potential of a commercial pulsed lidar for wind characterisation at a bridge site," J. Wind Eng. Ind. Aerodyn. **161**, 17–26 (2017).
[7]M. Nafisifard, J. B. Jakobsen, J. T. Snaebjörnsson, M. Sjöholm, and J. Mann, "Lidar measurements of wake around a bridge deck," J. Wind Eng. Ind. Aerodyn. **240**, 105491 (2023).
[8]N. E. Theeuwes, J. F. Barlow, A. Mannisenaho, D. Hertwig, E. O'Connor, and A. Robins, "Observations of tall-building wakes using a scanning Doppler lidar," Atmos. Meas. Tech. **18**(6), 1355–1371 (2025).
[9]X. Lyu, Y. He, S. Yin, S. Wong, T. K. T. Tse, and 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**, 112439 (2025).
[10]J. Yuan, H. Xia, T. Wei, L. Wang, B. Yue, and Y. Wu, "Identifying cloud, precipitation, windshear, and turbulence by deep analysis of the power spectrum of coherent Doppler wind lidar," Opt. Express **28**(25), 37406–37418 (2020).
[11]S. Wu, X. Zhai, and B. Liu, "Aircraft wake vortex and turbulence measurement under near-ground effect using coherent Doppler lidar," Opt. Express **27**(2), 1142–1163 (2019).
[12]H. Gao, J. Li, P. W. Chan, K. K. Hon, and X. Wang, "Parameter-retrieval of dry-air wake vortices with a scanning Doppler lidar," Opt. Express **26**(13), 16377–16392 (2018).
[13]J. Li, C. Shen, H. Gao, P. W. Chan, K. K. Hon, and X. Wang, "Path integration (PI) method for the parameter-retrieval of aircraft wake vortex by lidar," Opt. Express **28**(3), 4286–4306 (2020).
[14]Y. Zhang, J. Yuan, Y. Wu, J. Dong, and H. Xia, "Sub-meter wind detection with pulsed coherent Doppler lidar," Phys. Rev. Fluids **8**(2), L022701 (2023).
[15]Y. Zhang, Y. Wu, and H. Xia, "Spatial resolution enhancement of coherent Doppler wind lidar using differential correlation pair technique," Opt. Lett. **46**(22), 5550–5553 (2021).
[16]J. Liu, J. Yuan, H. Yang, A. Chen, Y. Gu, T. Wei, and H. Xia, "Research on impact of range resolutions on wind detection performance using virtual lidar," Opt. Express **33**(15), 32386–32397 (2025).
[17]H. Yang, J. Yuan, L. Guan, L. Su, T. Wei, and H. Xia, "Reconstruction for beam blockage of lidar based on generative adversarial networks," Opt. Express **32**(8), 14420–14434 (2024).
[18]G. Jing, H. Wang, X. Li, G. Wang, and Y. Yang, "An airflow velocity field reconstruction method with sparse or incomplete data using physics-informed neural network," J. Build. Eng. **88**, 109231 (2024).
[19]C. Hou, L. Marra, G. Y. C. Maceda, P. Jiang, J. Chen, Y. Liu, G. Hu, J. Chen, A. Ianiro, S. Discetti, A. Meilán-Vila, and B. R. Noack, "Machine-learned flow estimation with sparse data—Exemplified for the rooftop of a UAV vertiport," Phys. Fluids **36**, 125198 (2024).
[20]N. B. Erichson, L. Mathelin, Z. Yao, S. L. Brunton, M. W. Mahoney, and J. N. Kutz, "Shallow neural networks for fluid flow reconstruction with limited sensors," Proc. R. Soc. A **476**(2238), 20200097 (2020).
[21]P. Dubois, T. Gomez, L. Planckaert, and L. Perret, "Machine learning for fluid flow reconstruction from limited measurements," J. Comput. Phys. **448**, 110733 (2022).
[22]X. Liu, X. Li, Y. Zhang, T. Guo, D. Song, and M. Bao, "A physics-informed convolutional network based on feature fusion for high-resolution flow field reconstruction from sparse and noisy data," Phys. Fluids **37**(7), 073622 (2025).
[23]B. Zhang, L. Wang, J. Ge, Z. Luo, J. Yuan, Z. Wang, and J. Xu, "Advanced wake modeling in wind farm: A physics-informed framework with virtual LiDAR measurements," Phys. Fluids **37**(6), 065133 (2025).
[24]P. Xie, R. Li, Y. Chen, B. Song, W.-L. Chen, D. Zhou, and Y. Cao, "A physics-informed deep learning model to reconstruct turbulent wake from random sparse data," Phys. Fluids **36**(6), 065145 (2024).
[25]G. Luo, M. Blumenthal, M. Heide, and M. Uecker, "Bayesian MRI reconstruction with joint uncertainty estimation using diffusion models," Magn. Resonance Med. **90**(1), 295–311 (2023).
[26]P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2021), Vol. 34, pp. 8780–8794.
[27]S. Yang, C. Nai, X. Liu, W. Li, J. Chao, J. Wang, L. Wang, X. Li, X. Chen, B. Lu, Z. Xiao, N. Boers, H. Yuan, and B. Pan, "Generative assimilation and prediction for weather and climate," arXiv:2503.03038 (2025).
[28]X. Shao, M. Li, L. Wang, H. Han, Z. Liu, J. Huang, and Z. Qi, "An accurate reconstruction method for indoor bioaerosol concentration field from asynchronous and sparse LiDAR measurements based on latent diffusion models," Build. Environ. **282**, 113309 (2025).
[29]S. Sun, C. Nai, B. Pan, W. Li, X. Li, E. Foufoula-Georgiou, and Y. Lin, "Fusion of multi-source precipitation records via coordinate-based generative model," arXiv:2506.11698 (2025).
[30]J. Chao, B. Pan, Q. Chen, S. Yang, J. Wang, C. Nai, Y. Zheng, X. Li, H. Yuan, X. Chen, B. Lu, and Z. Xiao, "Learning to infer weather states using partial observations," J. Geophys. Res. Mach. Learn. Comput. **2**(1), e2024JH000260, https://doi.org/10.1029/2024jh000260 (2025).
[31]H. Gao, G. Hu, D. Zhang, W. Jiang, K. T. Tse, and B. R. Noack, "SiGAN: A 3D sensor importance deep generative model for urban wind flow field monitoring," Build. Environ. **262**, 111787 (2024).
[32]H. Gao, G. Hu, D. Zhang, W. Jiang, K. T. Tse, K. C. S. Kwok, and A. Kareem, "Urban wind field prediction based on sparse sensors and physics-informed graph-assisted auto-encoder," Comput-Aided. Civ. Infrastruct. Eng. **39**(10), 1409–1430 (2024).
[33]J. Yuan, J. Qiu, J. Hu, J. Hu, and H. Xia, "Transient wind camera using 1(seed) 10(amplifier) CDWL with a frame rate of 2 Hz and a range of 5 km," Opt. Lett. **50**(17), 5458–5461 (2025).
[34]J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (NeurIPS Foundation, 2020), Vol. 33, pp. 6840–6851.
[35]J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning* (PMLR, 2015), pp. 2256–2265.
[36]D. P. Kingma, "Adam: A method for stochastic optimization," arXiv:1412.6980 (2014).
[37]N. J. Annau, A. J. Cannon, and A. H. Monahan, "Algorithmic hallucinations of near-surface winds: Statistical downscaling with generative adversarial networks to convection-permitting scales," Artif. Intell. Earth Syst. **2**(4), e230015 (2023).