DEVELOPMENT OF A NONLINEAR REDUCED-ORDER MODEL FOR SPATIO-TEMPORAL URBAN CANYON AIRWAKES

Jeremy Seeyave
Georgia Institute of Technology
Atlanta, GA, USA

Marilyn J. Smith
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

An improved nonlinear reduced-order model (ROM) has been developed for application in Advanced Air Mobility (AAM) flight within urban canyon turbulent airwakes. The new ROM outputs time-accurate three-dimensional (3D) flowfields from a fluid mask and a timestamp as input. The ROM accurately predicts the training dataset, concurrently providing accurate interpolated intermediates. This capability has been validated through comparisons with both time-averaged and instantaneous flowfields, and is consistent with the lattice-Boltzmann model (LBM) simulations serving as truth data as well as experimental flow visualization. Furthermore, a diffusion model is employed to reintegrate the missing fine-scale turbulence into the ROM, thereby enhancing the fidelity of the predictions.

INTRODUCTION

Vehicle operations in the Advanced Air Mobility (AAM) framework will fly highly customized routes, adapting to the unique characteristics of each city and local weather conditions. AAM in this context includes both passenger vehicles, such as air taxis and air ambulances, as well as smaller uncrewed air vehicles (UAVs) for cargo delivery, observation, or fire suppression.

The unique aspects of AAM occur because of the region of the atmospheric boundary layer (ABL) where flight will occur, especially in urban environments. Figure 1 illustrates the different ABL regions that are relevant in AAM. The local character of the ABL (stable, unstable) defines the variation of the steady wind at different flight altitudes. In addition, the character of the ABL varies diurnally and seasonally. Further details of the ABL character can be found in many resources, including Refs. 1–4.

Traditional rotorcraft typically fly above the roughness sublayer, so that an estimate of the wind measurements and the clear air turbulence (CAT) using the von Karman or Dryden approximations are reasonable engineering assumptions (Ref. 5). While these can provide estimates for early assessments at altitudes $\leq$ 609 m (2000 ft), more accurate approximations can arguably be obtained by the Mann (Ref. 6) or Kaimal (Ref. 7) turbulence models. However, flight within the roughness sublayer (RS) or the urban boundary layer (UBL) will be dominated by the unsteady turbulent wake from the obstacles, such as buildings and trees. Atmospheric turbulence intensities of up to 40% may occur in urban flows, rather than the maximum 10% encountered in the free atmosphere (Ref. 8).

The AAM/UAV community has recently become more focused on this aspect of flight within the UBL and RS, and its importance to safety, for example, Ref. 10 and the recent downwash/outwash workshop (Ref. 11). The importance of the prediction of the outcomes of flights near these buildings was clearly demonstrated by several recent exploratory studies. Branco et al. (Ref. 12) resolved the turbulent wind environment of a vertiport for a generic hospital after a fatality at a hospital in England using steady Reynolds-averaged Navier-Stokes (RANS) simulations. Their predictions indicated that under some environmental conditions, helicopter outwash/downwash interactions with the building wakes did result in dangerous wind levels for pedestrians. Studies using modeling and simulation (M & S) tools, such as the industry-standard FLIGHTLAB, have been applied to further understand the character of the outwash/downwash (Ref. 13). These M & S tools have also indicated that emergencies, such as engine failures, in the turbulent wake of building structures can jeopardize safety (Ref. 14). Beyond the terminal area, these turbulent conditions could result in unacceptable ride quality for paying air taxi passengers or victims in air taxis (Ref. 15). The impact of these complex environments on UAV flight and control authority is also a topic of recent research (Refs. 16, 17).

Several sub-scale wind tunnel experimental campaigns have

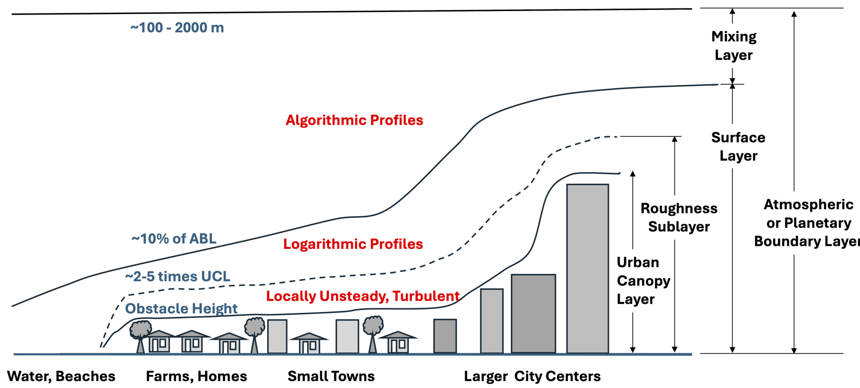

Figure 1: Vertical structure and scale of the ABL in an urban/suburban environment. From Ref. 9

been completed, using notional buildings (Ref. 18) or representative portions of cities (Ref. 19). In the latter case, results from improved delayed detached eddy simulations (IDDES) indicate that the wall-modeled large eddy simulations (WMLES) correlate well with the experimental results. Ship airwakes, which encounter similar turbulent wakes, indicate that WMLES and lattice-Boltzmann methods (LBM) also correlate well with experimental wake particle image velocimetry (PIV), when best practices are employed (Refs. 20, 21).

As these AAM vehicles will operate in complex unsteady, turbulent conditions associated with these urban environments, the ability to accurately model them is needed. Very high fidelity WMLES methods can provide the most accurate details of these complex wake flowfields (Refs. 19,20,22), but require high-performance computing over many days or even weeks, depending on the extent of the unsteady domain that is modeled. Mid-fidelity methods such as LBM offer several orders of magnitude reduction in these costs, with good correlation, especially for Reynolds-independent wakes (Refs. 20, 21). These computational approaches provide further understanding of detailed interactional aerodynamics, but may not be appropriate for all use cases.

Rapid and accurate characterizations of these complex environments are necessary for many use cases, including, but not limited to, flight planning, safety assessments, and pilot training. These unsteady environments will need to be coupled to design, flight control, M & S, and flight planning solvers, requiring near real-time or real-time unsteady wind metrics. There are two different approaches to achieve these goals: computationally resolving a set of fluid equations of motion or reduced-order models. The most accurate approach in the long term is the use of LBM. Ashok and Rauleder demonstrated that for frequencies between 0.2-2 Hz, LBM could provide accurate real-time unsteady turbulent winds for ship airwake problems with a single-node 8-GPU (graphics processor unit) computer (Ref. 23). Expansion to real-time predictions of urban landscapes has yet to be explored. Another approach is the use of computational mass conservation methods, coupled with inflow turbulence generators, that were originally developed for wind farms (Ref. 24). This approach shows promise, but initial results indicate further development is required to improve the accuracy of the wind predictions. Both the LBM and mass conservation approaches cited here included two-way coupling with the vehicle through the rotor(s), but these methods also require meshing of the entire computational domain.

The reduced-order modeling approaches rely heavily on modern machine learning (ML) algorithms. They also do not require meshing of the three-dimensional domain of interest. The trade-off is the more intensive development of these methods and the need to ensure that the ML algorithms can capture different building configurations and wind orientations. Folk et al. (Ref. 18) are exploring deep-learning Neural Network (NN) techniques where Light Detection and Ranging (LIDAR) data are combined with weather information to reproduce urban geographical areas. The process uses groups of buildings to train the model. These data are then applied to determine routes through the flowfield for autonomous flight via a flight controller.

A second reduced-order modeling approach is the subject of this paper. This approach, known as the Representative Environment Model (REM), has been developed specifically as a standalone input module for real-time applications. The efficacy of the model was demonstrated by Waanders et al. (Ref. 9), using a linear Proper Orthogonal Decomposition (POD) model as a baseline machine learning tool. A secondary focus of Ref. 9 was to explore the requirements of the machine learning model to accurately capture the flowfield using LBM and experiment as validation. While the features were accurately captured, improvements in the training requirements, speed, and applicability to new configurations were identified. This paper discusses these improvements.

REPRESENTATIVE ENVIRONMENT MODEL (REM)

The Representative Environment Model (REM) was developed as a tool that can be applied to many different use cases, from rapid design estimates to high-fidelity WMLES computations. It was also designed so that different contributions

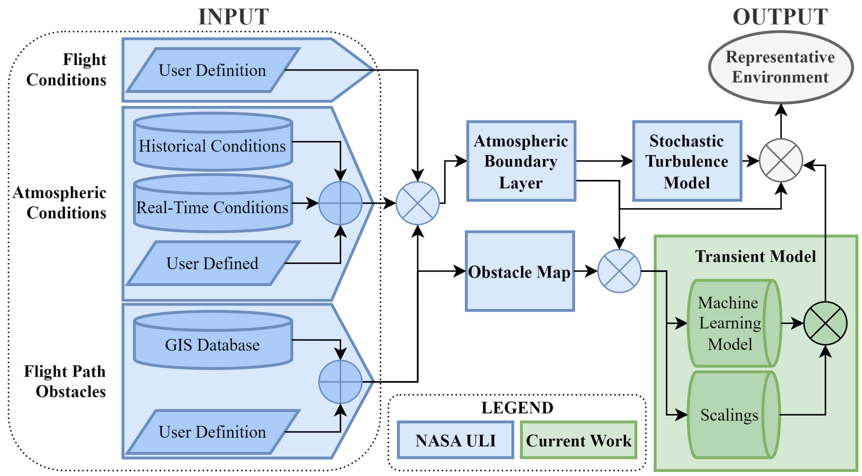

Figure 2: Overview of the Representative Environment Method (REM). From Ref. 9

of the unsteady (wind) environment on the application could be evaluated, as illustrated in the schematic of Fig. 2. This modular design combines the local or historical weather (temperature, steady wind) to first determine the character of the ABL (stable, unstable) so that the local winds at the relevant flight altitude can be computed. A stochastic or synthetic estimate of the air turbulence, without the influence of the local terrain, can then be added to characterize calm to extreme atmospheric conditions. As this turbulence is stochastic, the actual wind variations are defined with random number generators, so that each individual REM prediction will have small variations, comparable to the actual atmospheric conditions. The local terrain (natural or man-made) is included through user-defined or Geographic Information Systems (GIS) to generate accurate representations of the flight environment, without meshing. As the vehicle flies through the GIS- or user-developed flight route, the changing velocities are generated by REM and applied to the flight model, such as GTABB (Ref. 25) or FLIGHTLAB (Refs. 14, 15).

The machine learning model in Fig. 2 applies wake transients around buildings, as needed, by applying reduced-order models of nominal canonical shapes developed through machine learning from a design of experiments using the Lattice-Boltzmann Method (LBM). The machine learning model provides the ability to extend the unsteady wake information from a baseline building shape to multiple shapes and wind orientations.

For operational use cases, the REM removes the need for volume or surface meshing of the terrain and can generate real-time data on a personal computer without the need for special processors or extra memory. For design and analysis, the REM permits rapid change of the aleatory conditions to effect optimization and sensitivity analysis, and it can update the unsteady vector wind components with minimal cost, as it does not require recomputation of the flowfield.

MACHINE LEARNING FOR UNSTEADY THREE-DIMENSIONAL FLOWFIELDS

The integration of machine learning techniques into fluid dynamics has transformed how researchers model, interpret, and predict turbulent flows. There are many different techniques and algorithms, including linear and nonlinear resolutions that hold promise for the salient applications of interest in this effort. As noted earlier, linear POD models have been demonstrated to provide unsteady airwake estimates (Ref. 9), but do encounter accuracy issues in capturing some features and are not optimal for real-time applications. The use of neural networks (NN) to inform airwakes in clusters of buildings also is of interest (Ref. 18), but a detailed assessment of the prediction accuracy for use cases of interest here is needed.

Deterministic autoencoders have previously been used to predict the steady-state laminar flowfield over various obstacles from their fluid mask (Refs. 26, 27), but they may not be well suited to preserve the underlying stochastic structure of highly turbulent flowfields. Long short-term memory (LSTM) networks have also been applied to capture temporal dynamics in unsteady flows (Refs. 28, 29). More recently, transformer architectures, originally developed for natural language processing, have been adapted to learn long-range spatial and temporal dependencies in fluid flows with remarkable accuracy (Refs. 30, 31).

Generative models such as variational autoencoders (VAEs) have emerged as powerful tools in flow prediction tasks. These models are capable of learning a compressed latent representation of the high-dimensional flowfield, enabling both efficient generation of new samples and interpolation between flow states (Refs. 32, 33).

MACHINE LEARNING METHODS

Among these techniques, VAE has emerged as a powerful tool for learning low-dimensional representations of

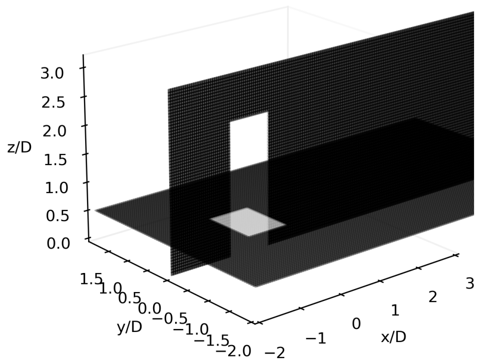

Figure 3: Binary fluid mask. 0 is colored black and 1 white.

high-dimensional flowfields while preserving their inherent stochastic characteristics. The proposed approach extends previous work by predicting the time-accurate three-dimensional (3D) flowfield using VAE, which is a probabilistic model and is more suited to predict the stochastic nature of turbulent flowfields.

Furthermore, the ROM is capable of generating time-accurate three-dimensional (3D) airwakes at Reynolds numbers of $\mathscr{O}(10^7)$, utilizing a fluid mask and a timestamp as input. The fluid mask is a binary mask of the building, one for the foreground (i.e., the building) and zero for the background (Fig. 3).

Variational autoencoder

VAE is a type of deep generative model that simultaneously learns a probabilistic encoder and decoder (Refs. 34, 35). The encoder maps high-dimensional inputs, $\mathbf{x}$ (e.g., velocity flowfields), to a low-dimensional latent space, parameterized as a multivariate Gaussian distribution. The encoder outputs a mean (μ_i) and a log-variance ($\ln(\sigma_i^2)$) for each latent variable z_i, where $i = 1, 2, \ldots, d$ and $d = 128$ is the size of the latent vector used in this study. Latent samples are generated via the re-parameterisation trick (Ref. 34), which facilitates gradient backpropagation, enabling end-to-end training via gradient descent. The decoder then maps $\mathbf{z}$ to the reconstructed flowfields $\tilde{\mathbf{x}}$.

Training is conducted by optimizing the evidence lower bound (ELBO), which balances reconstruction accuracy with the divergence of the latent distribution from a Gaussian distribution. Thus, VAE's training is performed with a compound loss function $\mathscr{L}$,

$$\mathscr{L}_{VAE} = \frac{1}{N}\sum_{i=1}^{N}(\mathbf{x}_i - \tilde{\mathbf{x}}_i)^2 - \frac{1}{2}\sum_{i=1}^{d}(1 + \ln(\sigma_i^2) - \mu_i^2 - \sigma_i^2), \quad (1)$$

where N represents the number of spatial points in $\mathbf{x}$. The first term in Eq. 1 is the mean squared error (MSE) reconstruction loss ($\mathscr{L}_{rec}$), which encourages accurate reconstruction of the inputs. The second term in Eq. 1 is the Kullback-Leibler (KL) divergence loss ($\mathscr{L}_{KL}$) (Ref. 36), which regularizes the latent distribution toward the prior, ensuring smoothness and continuity in the latent space (Ref. 37), so the model can smoothly interpolate between flows.

The balance between reconstruction accuracy and latent-space disentanglement can be fine-tuned by extending the VAE to a β-VAE (Refs. 38, 39). A β-VAE modifies a VAE's loss function (Eq. 1) by incorporating a scalar hyperparameter $\beta > 0$ into its loss function,

$$\mathscr{L}_{\beta-VAE} = \mathscr{L}_{rec} - \beta \mathscr{L}_{KL}. \quad (2)$$

The classic VAE is recovered when $\beta = 1$. A higher value of β promotes a more disentangled representation but may reduce reconstruction accuracy, while a lower value of β encourages more accurate reconstructions but a less disentangled latent space.

A Convolutional Neural Network (CNN) (Ref. 40) in the encoder and decoder captures the spatial relationships in the 3D fields. As illustrated in Fig. 4, the encoder is designed with four convolution layers that halve the spatial dimensions after each layer. All convolution layers use a $3 \times 3 \times 3$ kernel size, and the number of filters is quadrupled in each layer, from an initial value of 32, to preserve the flow of information while compressing the spatial dimension. Furthermore, each convolution layer is followed by the leaky rectified linear unit (LeakyRelu) (Ref. 41) activation function. Subsequently, the spatial and channel dimensions are flattened into a vector and fed into a fully connected layer with LeakyRelu activation to aggregate the information. Finally, the mean vector $\mu \in \mathbb{R}^{d_z}$ and the log-variance vector $\ln(\sigma^2) \in \mathbb{R}^{d_z}$ of the latent distribution are obtained through two fully connected parallel layers with linear activation.

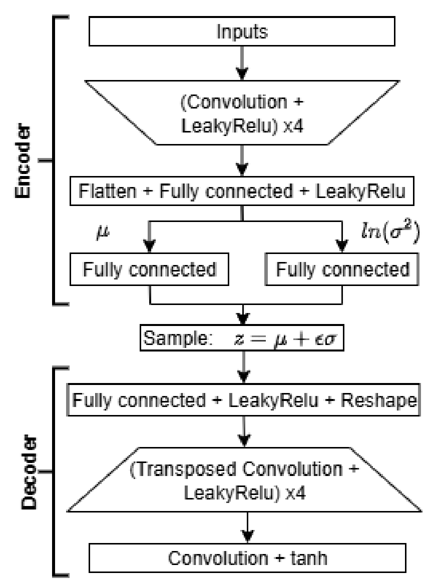

Figure 4: The VAE architecture.

The decoder network mirrors the encoder model, excluding the output layers. The latent vector, generated by sampling the encoder's output mean and log-variance, is fed into a fully connected layer to recover the size of the last convolution layer in the encoder and reshaped to match the shape. Four transposed convolution layers with a decreasing number of filters, each followed by a LeakyReLU activation layer , increase the spatial dimension to recover the encoder's input spatial dimension. A final convolution layer, with three filters and a hyperbolic tangent activation function (tanh) (Ref. 42), finishes the reconstruction of the encoder input. Because the data used in this work are normalized between -1 and 1, the final layer uses the hyperbolic tangent activation function to bound the output within this range.

Transformer model

The transformer model predicts the temporal evolution of turbulent flows in a reduced-order latent space. The model is designed to operate on sequences of latent vectors, where each input vector represents the sum of a compressed fluid mask ($\mathbf{z}_{fm}$) and a temporal positional encoding, and the sequence dimension captures the time evolution. The transformer model (Ref. 43) learns to map these input sequences to an output sequence of equal length, representing the latent vectors of the corresponding flowfields ($\mathbf{z}_{ff}$).

The architecture of the transformer model is shown in Fig. 5. The transformer model is implemented with a dimension $d_{model} = 1,024$, a sequence length $T = 100$, and 8 heads for the multi-head self-attention mechanism. The transformer model is also trained with a compound loss function, consisting of the MSE and cosine similarity loss. Using this compound loss function, the transformer model captures both the shape and scale of the flowfields' latent vectors.

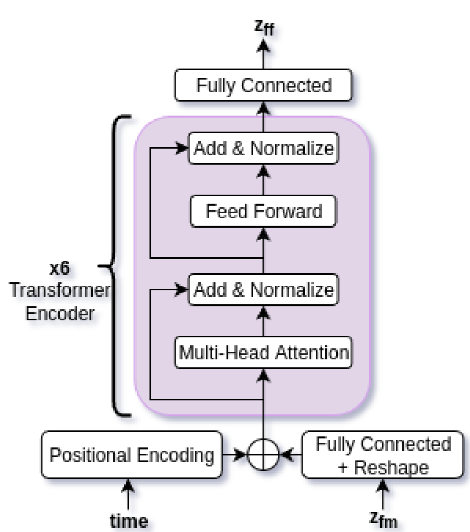

Figure 5: The transformer architecture. $\mathbf{z_{fm}}$ and $\mathbf{z_{ff}}$ represent the fluid mask and flowfield latent vector, respectively, and $\oplus$ indicates addition.

Diffusion Model

A conditional denoising diffusion probabilistic model (cDDPM) is used to predict residuals that enhance blurry flowfields. The model focuses on reconstructing high-frequency features lost by the ffVAE and ROM to obtain turbulent airwakes with realistic frequencies. The residual $r_0 = X_{LBM} - X_{ffVAE}$, where X_{LBM} and X_{ffVAE} represent the high-resolution solution from the LBM computations and the predicted flowfields from the ffVAE, respectively. The objective is to model a conditional distribution using cDDPM (Ref. 44), where the model learns to denoise a corrupted residual conditioned on the blurry input.

In the forward diffusion process, Gaussian noise is progressively added over T timesteps via a predefined noise schedule. Unlike the original cDDPM formulation, which predicts the added noise, the v-parameterization (Ref. 45) is employed, where the model predicts V_t, a weighted combination of the added noise and r_0. The authors in Ref. 45 showed that this formulation improves numerical stability and generation quality.

To guide the generation process, the model is conditioned on the blurry input (X_{ffVAE}). The training objective is a simple MSE between the true and predicted V_t. At inference time, the residual r_0 is predicted via the reverse diffusion process initialized from Gaussian noise. Finally, flowfields with high-frequency details are obtained by adding r_0 to X_{ffVAE}. The U-net architecture in (Ref. 44) was modified to perform 3D operations in the implemented cDDPM.

MACHINE-LEARNING-BASED ROMS

The baseline ROM is obtained by combining two variational autoencoders (VAE) and a transformer model, which have been individually trained (Fig. 6, offline training).

The first VAE is trained only on the fluid masks (fmVAE). The fmVAE is trained to compress the fluid mask into a latent representation using the encoder, and the decoder reconstructs the fluid mask from this compressed representation. The second VAE is trained on full-scale flowfields (ffVAE) obtained from the Lattice-Boltzmann (LBM) computational fluid dynamics (CFD) results, used as the truth data for this development. Similarly to the fmVAE, the trained ffVAE compresses the flowfield into a latent representation and reconstructs the flowfield from this compressed representation.

The transformer model takes as input the fluid mask's compressed representation from the trained fmVAE at each instance in time and predicts the resulting flowfield compressed representation. This is then used as input into the decoder-only part of the trained ffVAE to reconstruct the flowfield. in real time. By combining these models, the ROM generates time-accurate 3D flowfield over an object from the object's fluid mask (Fig. 6, online predictions).

Enhanced ROM (e-ROM)

After initial development, it was observed that the ROM yields blurred flowfields which fail to capture the high-

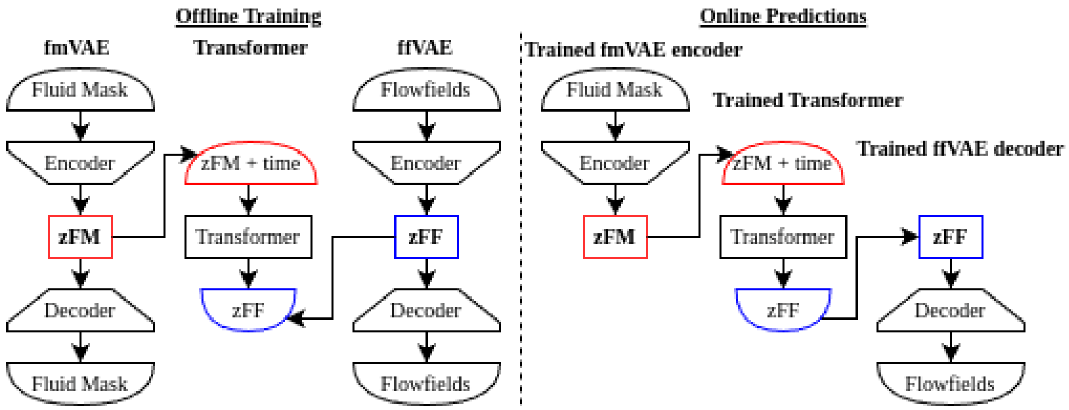

Figure 6: Machine learning ROM framework.

frequency, fine-scale turbulence. To address this limitation, the predictions of the ROM were enhanced using the diffusion model, cDDPM, which was not part of the original ROM. This approach, where the cDDPM is conditioned on the blurry outputs of the ROM, facilitates the reconstruction of the absent fine-scale turbulence. Consequently, the integration of the cDDPM output with the ROM output results in the creation of an enhanced-ROM (e-ROM).

Truth Datasets

The flowfields used as truth data were obtained using a Lattice-Boltzmann Method (LBM) solver (Ref. 9) for a canonical building. The building is a square cylinder with height $H = 2.4D$, where $D = 50$ m is the building width. A nominal freestream velocity of $U_\infty = 10$ m/s and Reynolds number $Re_D \approx 34 \times 10^6$ formed the baseline assessment. 1,050 3D solutions are recorded at a frequency of 1 Hz during each computation. The ffVAE is trained on four cases where the building yaw angle is varied $\theta = 0°, 20°, 30°$, and $45°$, resulting in a total of 4,200 3D flowfields for training. Three θ angles, $\theta = 15°, 22.5°$, and $37.5°$, are used to test the models.

The LBM simulations are performed on Cartesian grids with a spatial resolution of 0.5 m, containing $901 \times 401 \times 351$ points. The LBM solutions are downsampled to a spatial resolution of 2.5 m, containing $176 \times 80 \times 64$ points to reduce the GPU requirements for training. The maximum nondimensional velocity for all cases was approximately 2.7. For the ROM training, all cases were normalized by a factor of three so that they lie between -1 and 1.

Binary fluid masks (E.g., Fig. 3) for buildings at $\theta = -45°$ to $45°$ in an interval of $2.5°$ (i.e., 37 fluid masks) were implemented to train the fmVAE. Data augmentation was applied to this set of fluid masks to increase the number of samples by a factor of seven. The six new sets of fluid masks are obtained by reversing the order of elements along the following axes: x-axis, y-axis, z-axis, x-y-axis, x-z-axis, and y-z-axis, resulting in 259 samples for training.

ROM Training

The models and training pipeline were implemented using TensorFlow (Ref. 46). All models were trained using the Adam algorithm (Ref. 34) with its default setting and performed on a single NVIDIA GeForce RTX 4090 GPU. The training dataset was randomly shuffled, where the first 80% was used for training and the rest for validation.

The fmVAE was first trained using an initial learning rate of 1×10^{-3}, $\beta = 1$, and a batch size of 4. This relatively small batch size was deliberately chosen due to the limited availability of only 259 fluid mask samples. The learning rate was systematically reduced by a factor of 0.4 whenever the loss of the validation dataset did not improve after 25 consecutive complete passes through the dataset (epoch). Training was stopped when the loss of the validation dataset did not improve after 200 consecutive epochs, which required approximately 300 epochs over 15 minutes. The encoder of the trained fmVAE is used to create the fluid mask latent vector (z_{fm}), which is then used as input to the transformer model. The encoder required less than 30 seconds to resolve the four training cases.

Next, the ffVAE was trained using an initial learning rate of 5×10^{-4}, $\beta = 1 \times 10^{-5}$, alongside a batch size of 16, which represents the maximum capacity that the GPU memory can accommodate. The learning rate decreased by a factor of 0.4 when the loss of the validation dataset did not exhibit improvement after 25 consecutive epochs. Training was stopped when the loss of the validation dataset did not improve after 200 consecutive epochs. Training stopped after approximately 1,500 epochs over around 20 hours. The encoder of the trained ffVAE is used to create the flowfields latent vectors (z_{ff}), which are then used as the training target of the transformer model. The encoder takes around 95 seconds to predict the 1,050 latent vectors of each training case.

Finally, using the fmVAE and ffVAE outputs, the transformer model was trained. Training is carried out with an initial learning rate of 1×10^{-4} and a batch size of 64. The learning rate decreased by a factor of 0.4 when the loss of the validation dataset did not improve after 51 consecutive epochs. Training was stopped when the loss of the validation dataset did not improve after 500 consecutive epochs. Training stopped after approximately 7,000 epochs over around 9 hours.

Training of the cDDPM is carried out with a fixed learning rate of 5×10^{-5} and a batch size of 8. Training was performed for 250 epochs and took 36 hours.

RESULTS

To assess the efficacy of the ROM model for these urban airwakes, the predictions of the nominal building are evaluated and correlated with the original LBM computations. The results marked as 'ffVAE' are obtained using the ffVAE alone; that is, the flowfields of the test cases are fed as input and the reconstruction ability of the ffVAE is evaluated. ROM represents the results obtained directly following the process described previously (Fig. 6, online predictions). Lastly, e-ROM represents the results obtained by the enhanced ROM approach as described in the enhanced ROM subsection. The average results illustrated herein are time-averaged over all of the 1,050 seconds of prediction.

As the data were trained on the cases where the building yaw angle was $\theta = 0°, 20°, 30°$, and $45°$, three different building yaw angles of $\theta = 15°, 22.5°$, and $37.5°$ are used to test the models. To permit a deeper evaluation of the model, the building yaw angle of 22.5° is employed, but all three cases are discussed.

The primary features of interest about a rectangular cross-section building are visualized in Fig. 7. Of interest is the ability of the original and enhanced ROMs and ffVAE to capture the top and side shear layers and the recirculation behind the building.

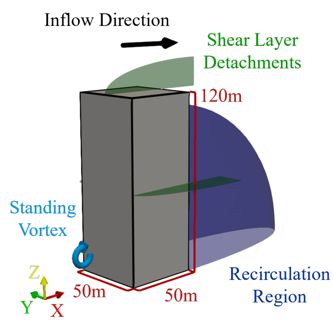

Figure 7: Visualization of the primary aerodynamic flow features for a building with a head-on wind ($\theta = 0°$). From Ref.9

A vertical plane of time-averaged velocity magnitude in the building wake is first examined for the three validation building yaw angles in Fig. 8. The plane is located two building widths ($X/D = 2$) downstream of the building center, so that it is well within the recirculation region, and it extends beyond each side and top of the building to capture the shear layers. As the yaw angle of the building increases, both the front and the right sides of the building create the top and side shear layers. Because the front side dominates the right side until $\theta = 45°$, an asymmetric wake is formed. Apart from some magnitude differences, the models effectively predict these asymmetric wake shapes at various test yaw angles. The ffVAE indicates smaller velocity magnitudes in the recirculation zone than observed by the averaged LBM data, but these are primarily recovered in the final ROM and e-ROM predictions.

To further explore these predictions, the time-averaged streamwise, spanwise, and normal wind components, nondimensionalized by the freestream wind, are analyzed in Figs. 9 – 11, respectively, at several vertical planes in the streamwise direction. In these figures, the building is located on the right, so that the wind also enters from the right and moves downstream to the left. The predicted flow topology from the models is similar to those observed in the LBM computation at those planes. Above the building at $X/D = 0$, the shear layer emanating from the front of the building overshadows the one from the right side of the building. At $X/D = 2$, the models capture the asymmetric shape of the wake and the recirculation zone characterized by the reversed flow region in the streamwise velocity. Additionally, counterrotating vortices are convected downstream as indicated by the pair of positive and negative spanwise velocity and the centered column of negative normal velocity bounded by positive normal velocity. At $X/D = 4$, the wake has reduced in size while the asymmetric shape persists.

To quantify the difference between the model predictions and the LBM results, the errors between the two are presented for a streamwise, centerline plane ($Y/D = 0$) and a horizontal plane through the middle of the building ($Z/D = 1.2$) in Fig. 12. The training portion of the ffVAE indicates variations along the surface of the building, which are not of particular concern in the flight use cases of these data. These variations propagate to the ROM and e-ROM, as expected. They are most apparent in the streamwise velocity and velocity magnitude plots. The top and side shear layers from the building indicate the most significant difference, as the angle of the shear layers and their extent are slightly offset between the LBM and the two ROMs. This indicates that the training data in their regions may need to be further refined so that these high gradients may be further resolved.

Overall, the velocities in each direction, other than the shear layer, are accurately captured to within ± 10% of the steady wind. For a moderate wind speed of 10 m/s, this means that the prediction accuracy is within 1 m/s.

To further quantify the accuracy of the different model predictions to the LBM solutions, the time-averaged streamwise velocity and velocity magnitude profiles at several X/D locations along the center-plane $Y/D = 0$ are evaluated in Fig. 13. The ffVAE predictions are accurate outside of the recirculation region, but underpredict the reverse flow velocity and the velocity magnitude in the wake. Interestingly, the ROM predictions are more accurate than those of the ffVAE in the recirculation region. However, the ROM does not capture the increase in velocity when the clockwise recirculating flow

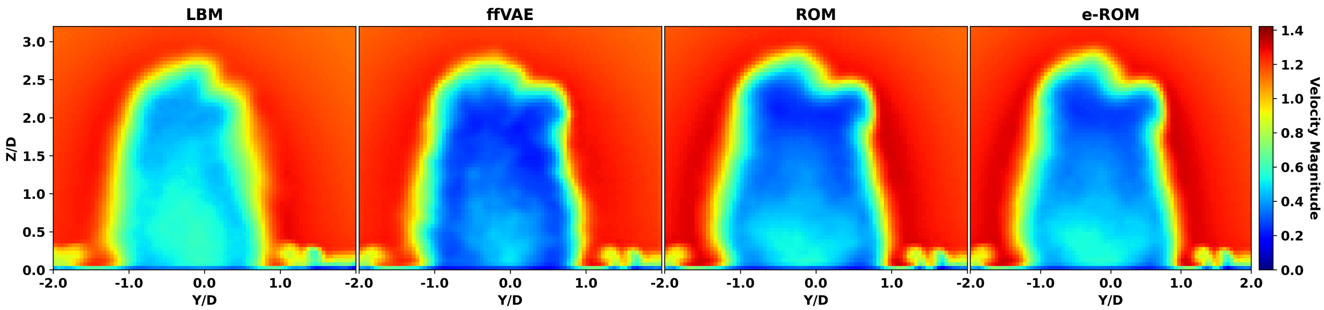

(a) Test case $\theta = 15°$

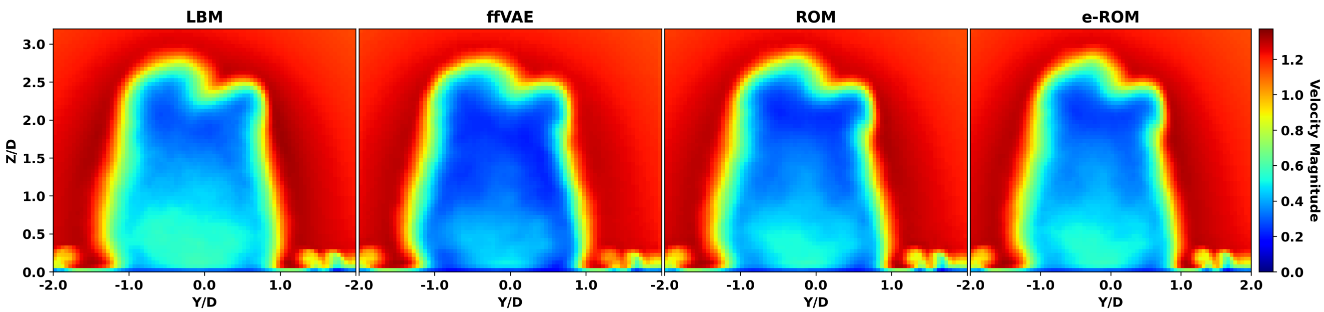

(b) Test case $\theta = 22.5°$

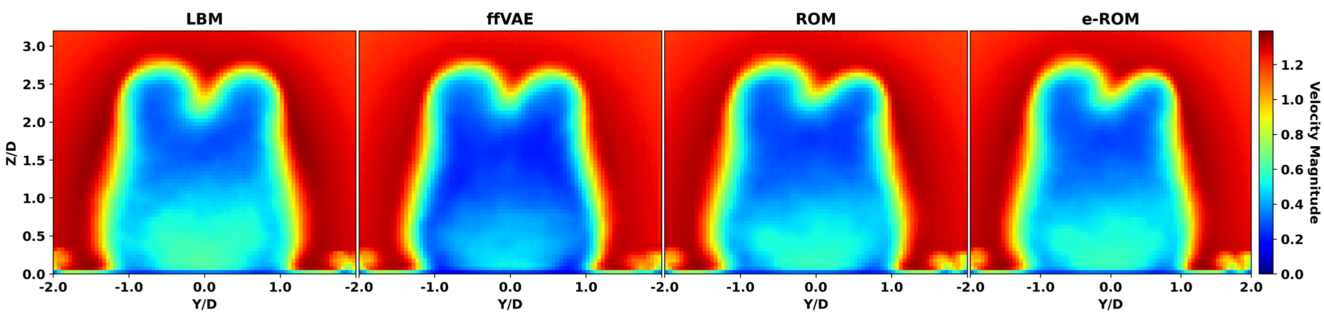

(c) Test case $\theta = 37.5°$

Figure 8: Average velocity magnitude at the plane $X/D = 2$ for the three test case.

merges with the region of separated flow, caused by the top shear layer, above the building at $X/D = 1$ (just above the dotted line, best observed at the $X/D = 0 - 2$ planes). By X/D = 3 – 4, these vortices have dissipated. There are only small changes between the ROM and e-ROM. However, the e-ROM shows an improvement from the ROM predictions for the velocity magnitude profiles. These differences primarily lie in the recirculation zone from Z/D = 0–2, and these are due to the modification of the original ROM formulation to more accurately capture the reverse flow velocities, as described earlier. Aft of the recirculation region at $X/D = 4 - 5$, both of the ROM models' predictions are comparable.

The unsteadiness of the wake is evaluated through the standard deviation of the velocity magnitude ($\sigma_{V_{mag}}$) in Fig. 14. While the time-averaged velocities showed good agreement between the model predictions and the LBM results, some differences can be observed in the unsteady content of the wake. For all three streamwise wake planes, the ffVAE shows significantly reduced unsteady content compared to LBM, ROM, and e-ROM results, although the character or shape of the wake features is captured. Above the building (X/D = 0), both the ROM and e-ROM capture the asymmetric unsteady content, though the e-ROM approach is closer to the LBM magnitudes. At X/D = 2, the magnitude and extent of the flow unsteadiness increase near the ground plane, indicating its influence on airwakes at low altitudes, potentially influencing terminal operations. The magnitude and location of the strongest unsteady flow are along the side edges of the building, so that gaps between buildings will play a significant role in flight quality, as demonstrated through FLIGHTLAB flight simulations by Goericke et al. (Ref. 15). The prediction of the ground plane and side shear layers is most accurate in the e-ROM prediction, but the extent and magnitude of the unsteadiness are still less than the LBM predictions. Farther downstream at $X/D = 4$, the recirculation zone is collapsing, and strong unsteadiness is observed throughout the near-elliptical zone. Again, the character and strength of the unsteadiness in this area are best captured by the e-ROM prediction, though it is less than the LBM unsteadiness.

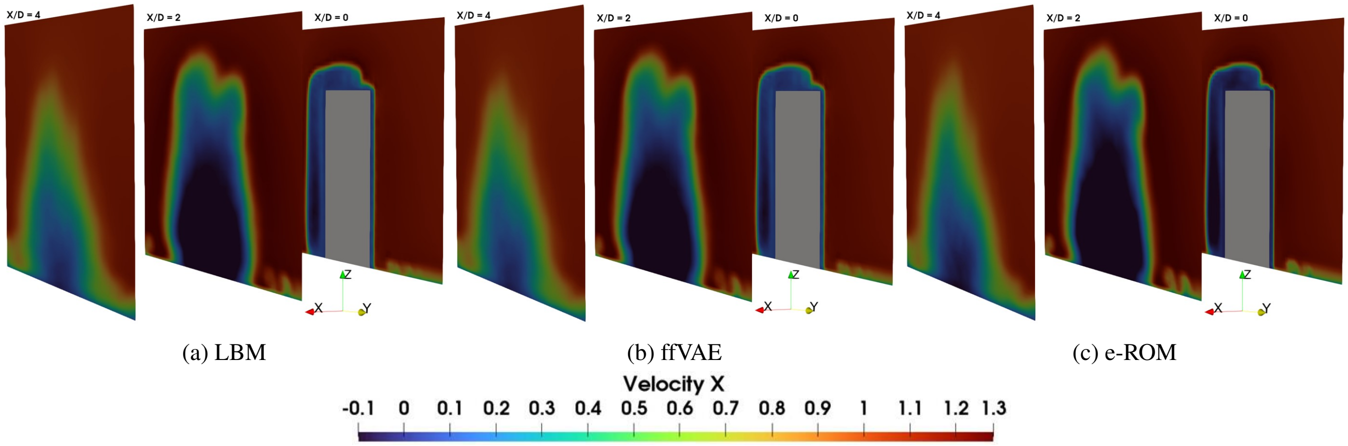

(a) LBM (b) ffVAE (c) e-ROM

Figure 9: Average streamwise (x) velocity at planes $X/D = 0, 2$, and 4, right to left, respectively, for a building yaw angle $\theta = 22.5°$.

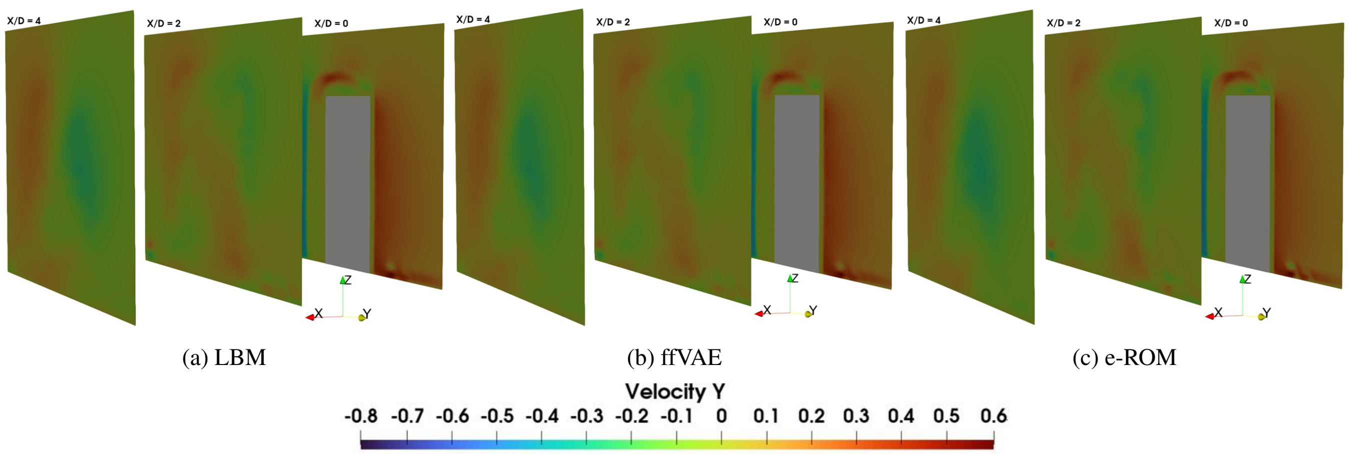

(a) LBM (b) ffVAE (c) e-ROM

Figure 10: Average spanwise (y) velocity at planes $X/D = 0, 2$, and 4, right to left, respectively, for a building yaw angle $\theta = 22.5°$.

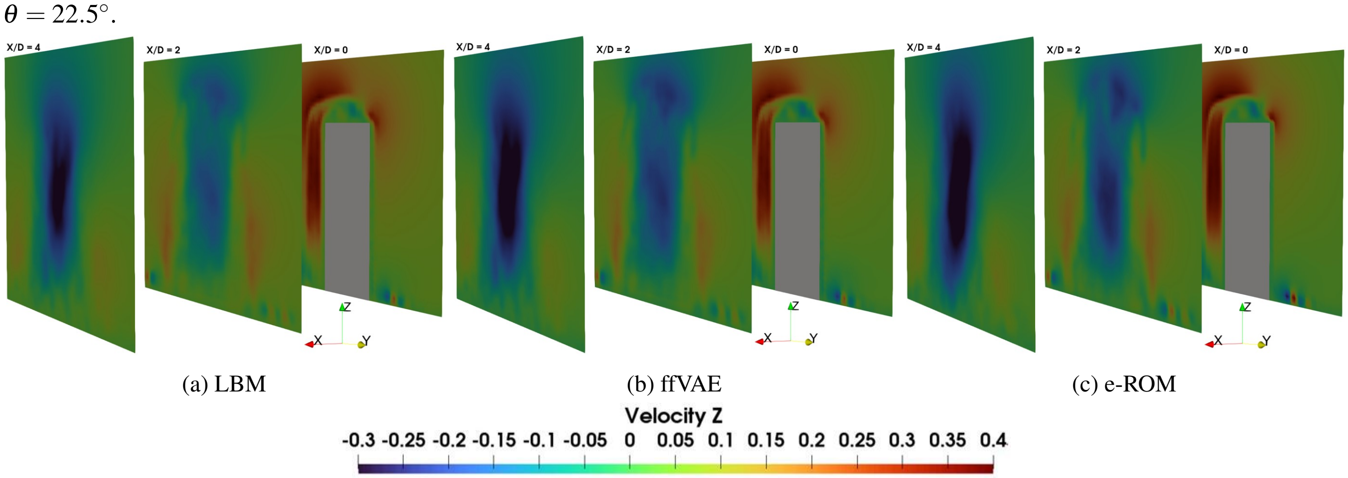

(a) LBM (b) ffVAE (c) e-ROM

Figure 11: Average normal (z velocity at planes $X/D = 0, 2$, and 4, right to left, respectively, for a building yaw angle $\theta = 22.5°$.

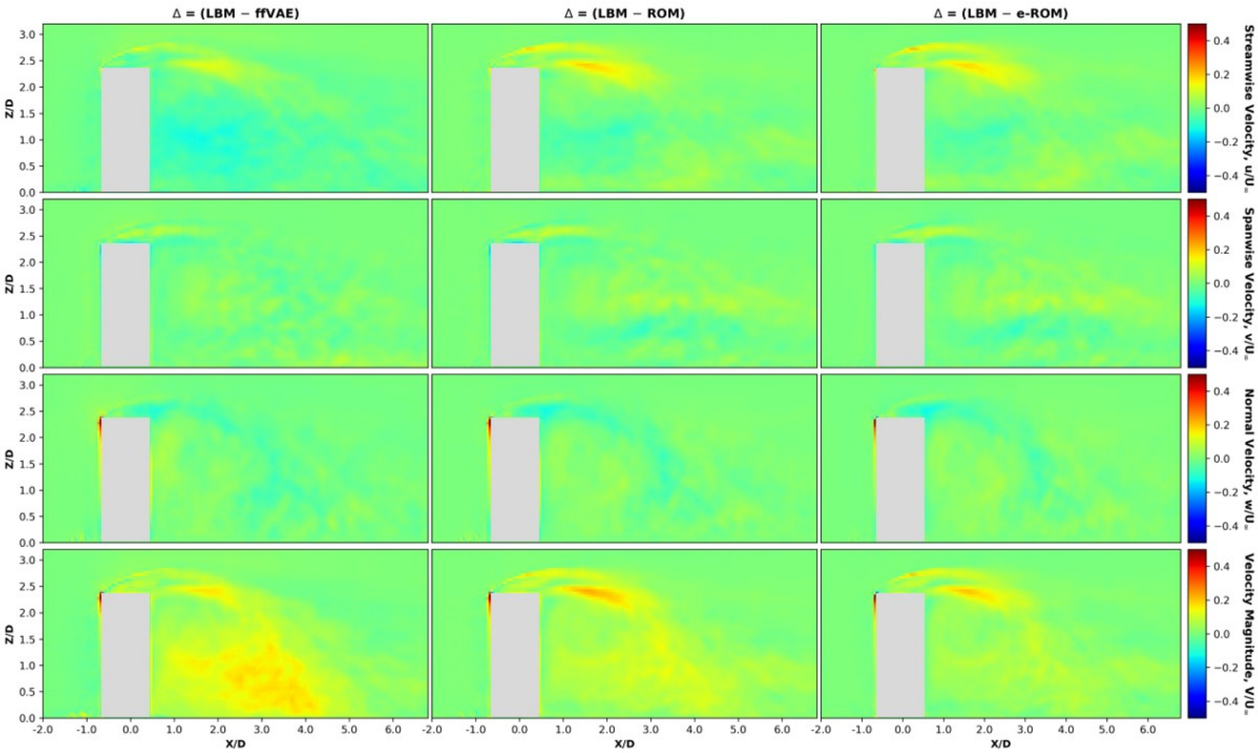

(a) Plane at $Y/D = 0$

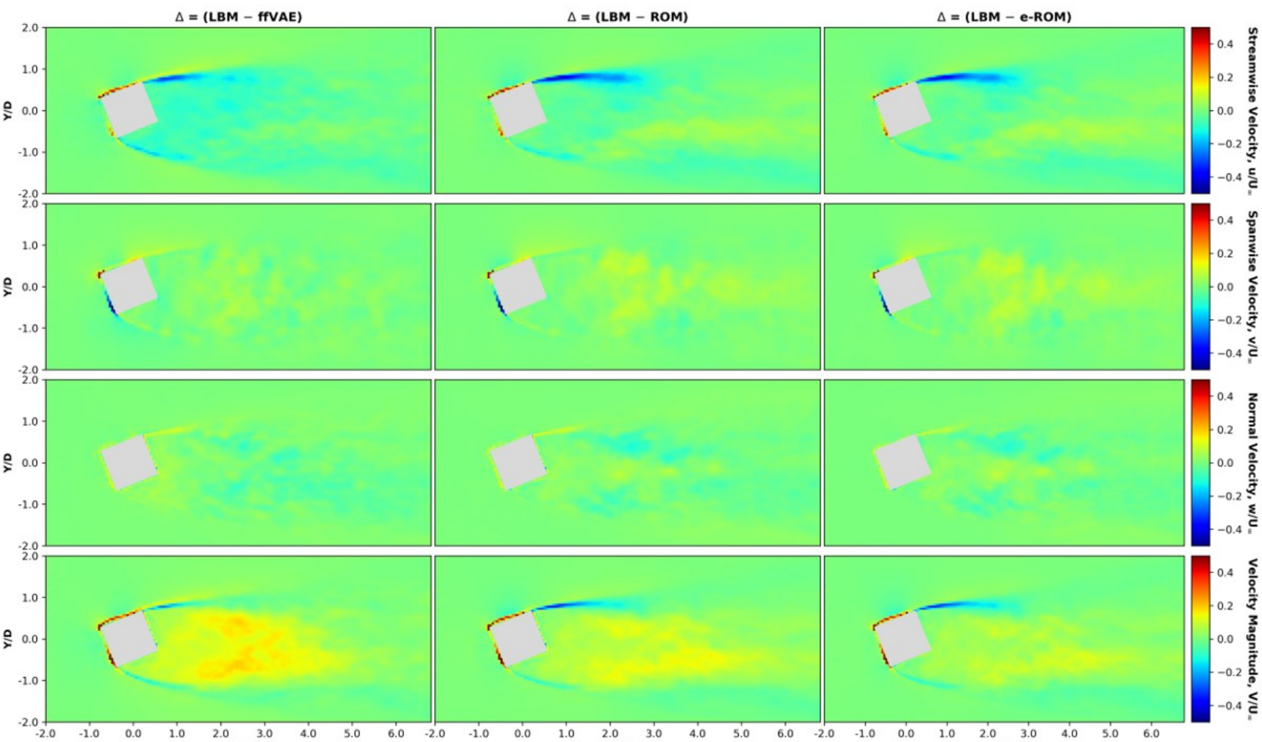

(b) Plane at $Z/D = 1.2$

Figure 12: Time-averaged velocity error for the $\theta = 22.5^\circ$ test case on the planes $y/D = 0$ and $Z/D = 1.2$. From top to bottom, row $1-4$ shows the streamwise, spanwise, normal, and magnitude velocity, respectively. From left to right, the prediction error of the ffVAE, ROM, and e-ROM compared to the LBM solution.

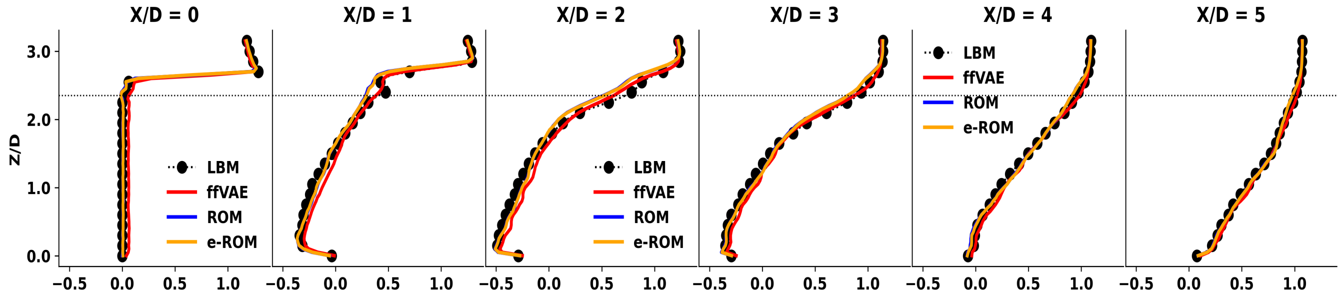

(a) Streamwise velocity, $\theta = 22.5°$

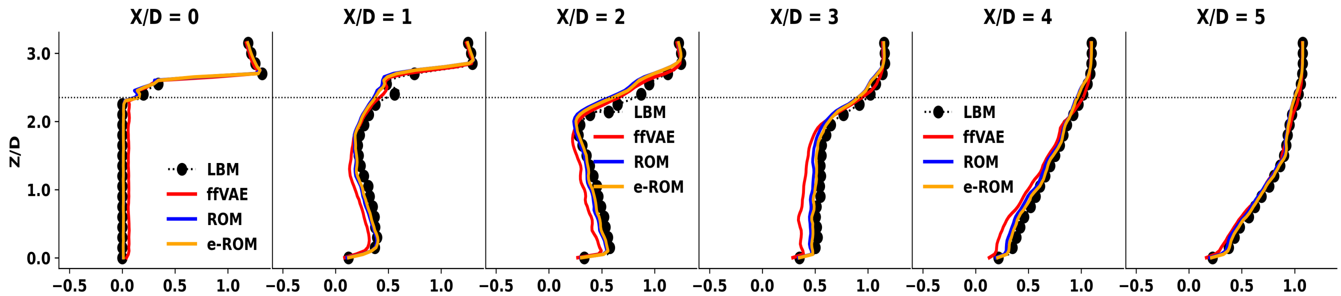

(b) Velocity magnitude, $\theta = 22.5°$

Figure 13: Averaged velocity profile for the $\theta = 22.5°$ test case along the centerline $y/D = 0$ at various distances downstream of the building center. The horizontal dotted line represents the building height.

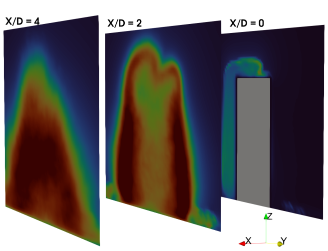

(a) LBM

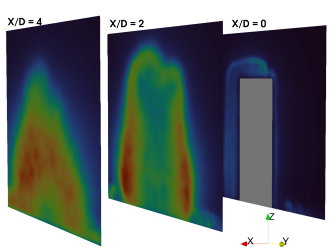

(b) ffVAE

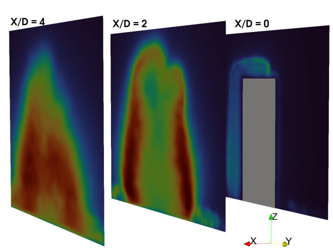

(c) ROM

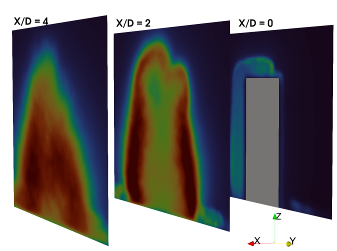

(d) e-ROM

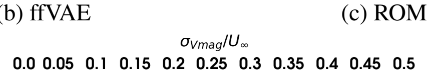

Figure 14: Standard deviation of velocity magnitude at planes $X/D = 0, 2$, and 4 for the test case $\theta = 22.5°$.

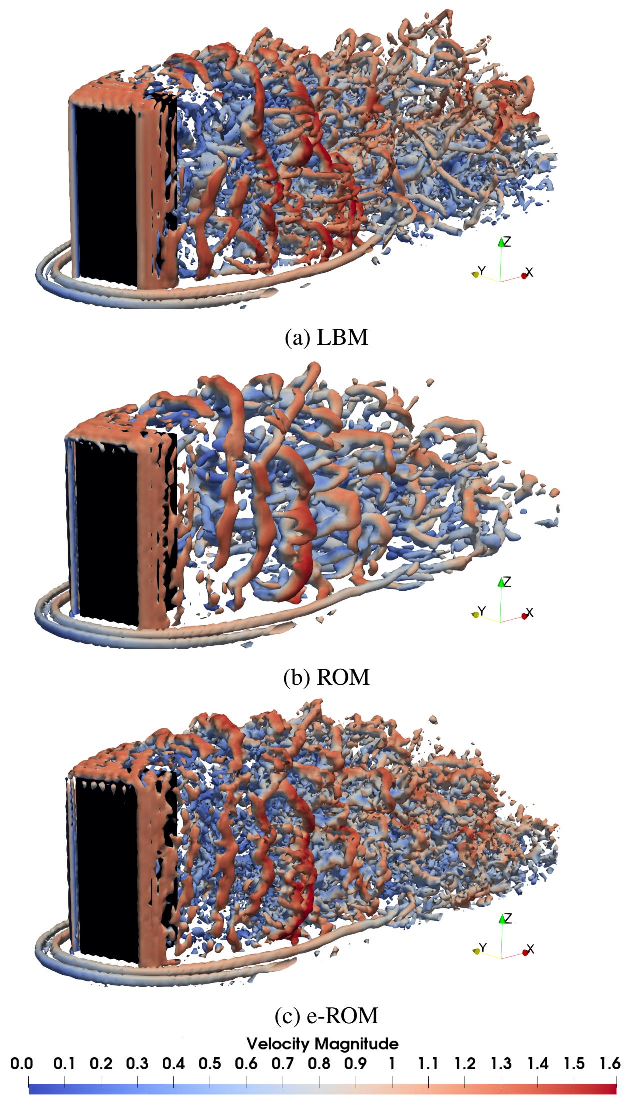

Figure 15: Isosurfaces of Q-criterion colored with velocity magnitude for the test case $\theta = 22.5°$. The building is colored black.

The complex character of the wake can be visualized using isocontours of Q-criterion, illustrated in Fig. 15. The complex wake predicted by the LBM simulations (Fig. 15a) is manifested by a large number of rollers of different velocity magnitudes. The shear layers (red) and the recirculation zone (blue) are clearly observed. This visualization is the best demonstration of the enhancements made to the original ROM. The original ROM wake in Fig. 15b predicts a much sparser wake structure that is also significantly dissipated compared to the LBM results. The enhanced ROM (e-ROM) wake in Fig. 15c recovers the complex structure of the wake for both the shear layers and recirculation zone. In addition, comparison of the velocity magnitudes indicates that the e-ROM restores the larger velocity variations observed in the LBM solution that are missed in the ROM predictions.

Overall, the results show good correlation with the LBM simulation and experimental flow visualization of a surface-mounted obstacle (Ref. 47). Importantly, the ROM successfully captures principal vortical structures, including the horseshoe vortex that originates ahead of the building and advects downstream along the lateral facades. Additionally, the ROM accurately identifies the roof and lateral vortices, which contribute to the formation of arch vortices downstream of the building. Furthermore, the e-ROM method effectively reintegrates the missing fine-scale turbulence into the ROM, thereby enhancing the realism of the predictions.

The frequency content in the flowfields is analyzed via a spatially-averaged power spectral density (PSD) computed using Welch's overlapped segment averaging estimator (Ref. 48). PSD is calculated to evaluate turbulence characteristics in the frequency domain. Because the frequency of disturbances affects handling qualities and pilot workload, predicting turbulent airwakes with realistic frequencies is crucial. The process consists of calculating the PSD at each point in the volume and then averaging these spectral densities to yield the average frequency content throughout the entire region of interest. Fig. 16 illustrates the continuous PSD distribution for the three velocity components and the velocity magnitude.

The LBM indicated that the primary frequencies of interest lie within 0-0.5 Hz, and a simple linear scale is applied here to demonstrate the differences of the models at the higher frequencies of interest. The ffVAE predicts the general trend, but the magnitudes are weaker than the LBM predictions at all frequencies of interest. The magnitude predicted by the ROM is stronger at the lower frequencies, but also misses the higher frequency content of the LBM data. The pronounced decrement in the PSD at approximately 0.2 Hz mirrors findings obtained through modal decomposition techniques, such as POD, which was the baseline approach for the REM demonstrated by Waanders et al. (Ref. 9). In this context, the most energetic modes are employed, whereas the high-frequency modes with lower energy contributions are excluded in the formulation of the ROM. The enhancements added to the ROM accurately capture the trends and magnitudes of the LBM data. This is directly related to the e-ROM's ability to restore the LBM fine-scale turbulence lost by the ROM predictions.

Finally, the e-ROM's ability to incorporate the higher-frequency fine-scale turbulence can be confirmed by comparing the instantaneous velocities in Fig. 17. When applying these representative airwakes for time-accurate use cases, it is important that the fine-scale details provide an accurate representation of the airwake on the order of the vehicle reference lengths and smaller. The two planes evaluated previously in Fig. 12 are presented. The ROM (middle column) consistently washes out the fine scales of the instantaneous flowfield, as was observed for the averaged flowfield values. The e-ROM enhancements indicate that it more consistently replicates the instantaneous character in all three constitutive directions.

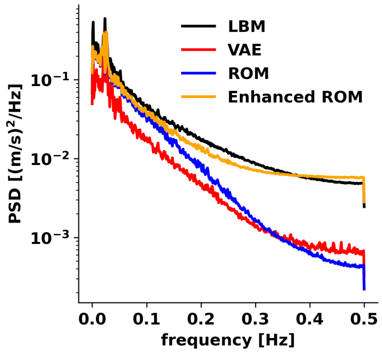

(a) Streamwise velocity

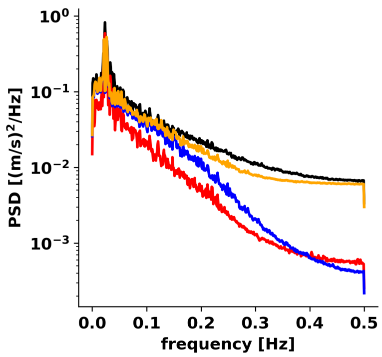

(b) Spanwise velocity

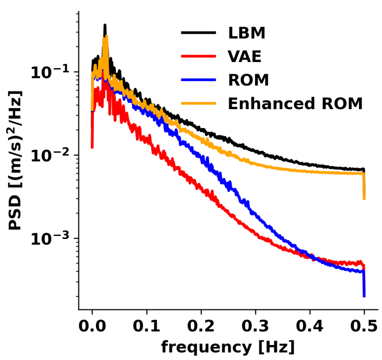

(c) Normal velocity

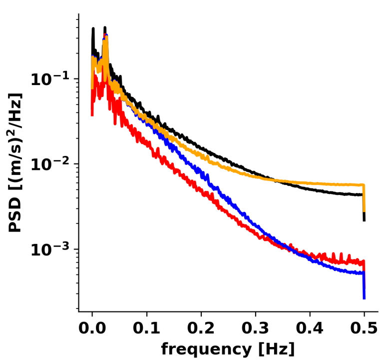

(d) Velocity magnitude

Figure 16: Spatially-averaged PSD for the test case $\theta = 22.5°$.

CONCLUSIONS

After assessment of several different reduced-order models (ROMs) that have been applied to flowfields around bluff bodies, the Variational Autoencoder (VAE) was determined to have the greatest potential for application in airwake predictions. The development of the ROM required several modifications to capture sufficient details needed for an airwake analysis. The original ROM formulation, based on prior models applied to unsteady flowfields, encountered difficulties accurately predicting the reverse flows in the leeward side of the bluff bodies (rectangular cross-sectional cylinder or prism). This was corrected by modifying the transformer model so that it directly output the flowfield latent vectors rather than a mean and log-variance. The original formulation was also observed to be biased towards the low-frequency content, similar to POD, while not capturing the higher-frequency content. This was corrected by including the conditional denoising diffusion probabilistic model (cDDPM) to generate fine-scale turbulence to more accurately predict the higher frequencies.

The enhanced ROM (e-ROM) was demonstrated to accurately capture the unsteady, turbulent airwakes within $\pm 10\%$ of the nominal wind. For most flights, this translates to errors well below 1 m/s. In addition, the location of the major airwake features was more accurately captured than the previous linear ROM model (POD) from Waanders et al. (Ref. 9). This capability was demonstrated in both time-averaged flowfields, as well as instantaneous flowfields, compared with the lattice-Boltzmann model (LBM) simulations that served as truth data. Further, the training requires much less information and the online predictions (predictions during application) are sufficiently rapid to permit real-time assessments.

Further exploration of the training data needed to capture the strong gradients associated with the top and side shear layers is underway. While the e-ROM captures these well, some improvements may be observed with an optimal spatio-temporal refinement in these areas.

The application of the e-ROM into the representative environment model (REM) is underway to assess the impact of these improvements for various use cases, such as flight safety, power requirements, and vibratory loads (ride quality).

Author contact:
Jeremy Seeyave jseeyave3@gatech.edu
Marilyn J. Smith ms55@gatech.edu

ACKNOWLEDGMENTS

This effort has been funded in part by the U.S. Government under Cooperative Agreement No. W911W6-21-2-0001 (Technical Monitor: Mahendra Bhagwat) and the NASA University Leadership Initiative (ULI), Cooperative Agreement 80NSSC21M0113 (Technical Monitor: Mia Siochi). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies of position, either expressed or implied, of the National Aeronautics and Space Administration (NASA) or the U.S. Government.

The authors would like to thank Georgia Tech graduate students Walter Chong, Aaron Crawford, Sheldon Salins, and

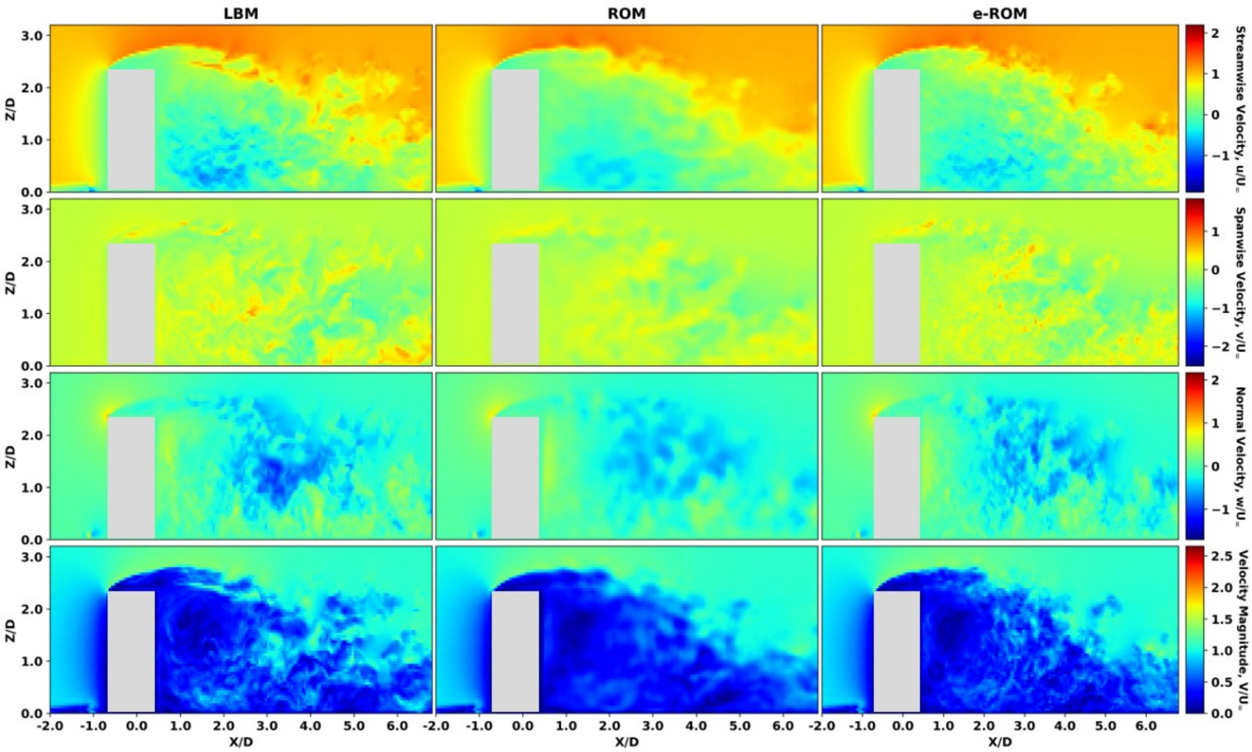

(a) Plane at $Y/D = 0$

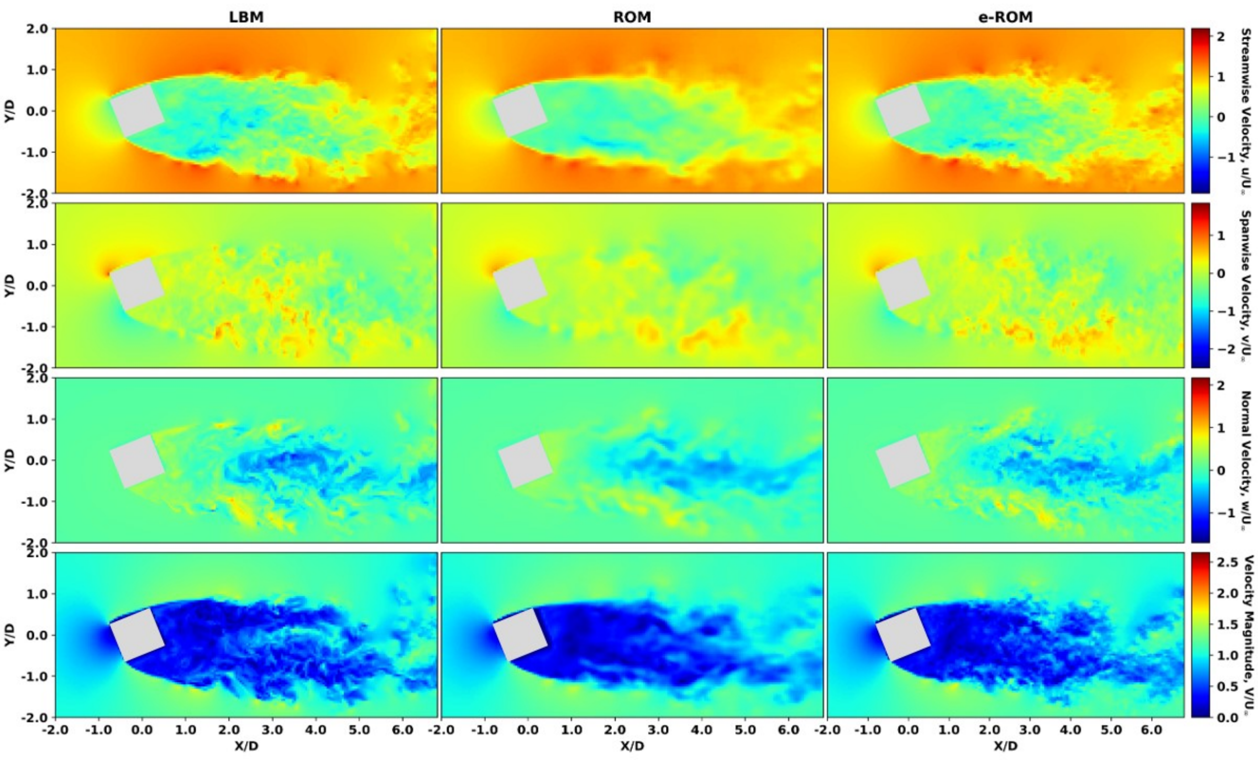

(b) Plane at $Z/D = 1.2$

Figure 17: Instantaneous velocity for the $\theta = 22.5^{\circ}$ test case on the planes $y/D = 0$ and $Z/D = 1.2$, demonstrating e-ROM ability to add high-frequency fine-scale turbulence. From top to bottom, row $1-4$ shows the streamwise, spanwise, normal, and magnitude velocity, respectively. From left to right, the LBM solution, the ROM, and e-ROM predictions.

Duncan Waanders (advised by Dr. Juergen Rauleder) who provided LBM simulations used for training and ML correlations.

REFERENCES

1. Stull, R., *Practical Meteorology: An Algebra-based Survey of Atmospheric Science -version 1.02b*, Dept. of Earth, Ocean & Atmospheric Sciences, University of British Columbia, `https://www.eoas.ubc.ca/books/Practical_Meteorology/`, 2017.

2. Oke, T. R., Mills, G., Christen, A., and Voogt, J. A., *Urban Climates*, Cambridge University Press, 2017.

3. Cantero, E., Sanz, J., Borbón, F., Paredes, D., and García, A., "On the Measurement of Stability Parameter over Complex Mountainous Terrain," *Wind Energy Science*, Vol. 7, (1), 2022, pp. 221–235. DOI: 10.5194/wes-7-221-2022.

4. Barlow, J. F., "Progress in Observing and Modelling the Urban Boundary Layer," *Urban Climate*, Vol. 10, ICUC8: The 8th International Conference on Urban Climate and the 10th Symposium on the Urban Environment, 2014, pp. 216–240. DOI: 10.1016/j.uclim.2014.03.011.

5. Beal, T., "Digital Simulation of Atmospheric Turbulence for Dryden and von Karman Models," *Journal of Guidance, Control, and Dynamics*, Vol. 16, (1), 1993, pp. 132–138. DOI: https://doi.org/10.2514/3.11437.

6. Mann, J., "Wind Field Simulation," *Probabilistic Engineering Mechanics*, Vol. 13, (4), 1998, pp. 269–282. DOI: 10.1016/S0266-8920(97)00036-2.

7. Kaimal, J. C., Wyngaard, J. C., Izumi, Y., and Coté, O. R., "Spectral Characteristics of Surface-Layer Turbulence," *Quarterly Journal of the Royal Meteorological Society*, Vol. 13, (4), 1972, pp. 563–589. DOI: 10.1016/S0266-8920(97)00036-2.

8. Barber, H., and Wall, A., "Urban Flow: what drone pilots need to know," Technical Report LTR-AL-2020-0075, 2021-01-29, National Research Council of Canada, Aerospace Research Centre, January 2021. DOI: 10.4224/40002000.

9. Waanders, D., Salins, S., Rauleder, J., and Smith, M. J., "A Real-Time Reduced-Order Model for the Atmospheric Boundary Layer Including Roughness Sublayer," Proceedings of the Vertical Flight Society 81st Annual Forum, Virginia Beach, VA, May 20–22 2025. DOI: 10.4050/F-0081-2025-0325.

10. Schajnoha, S., Barber, H., Larose, G., Al-Labbad, M., and Wall, A., "The Safety of Advanced Air Mobility and The Effects of Wind in the Urban Canyon," Proceedings of the Vertical Flight Society 78th Annual Forum, may 2022. DOI: 10.4050/f-0078-2022-17611.

11. Vertical Flight Society, "Downwash & Outwash Workshop," URL=https://vtol.org/annual-forum/forum-81/downwash-and-outwash-workshop.

12. Branco, D. S. and Stephens, I. and White, M. D. and Watson, N. A., "Assessing the Interaction of Helicopter Rotor Downwash and Turbulent Airwakes near Hospital Landing Sites," Proceedings of the 50th European Rotorcraft Forum, Marseille, France, September 10–12 2024.

13. Goericke, J., Brenner, F., Lee, D., and Habana, Z., "Investigation of eVTOL Operations in Turbulent Airwake Using Comprehensive Rotorcraft Analysis and Physics-Based Electric Propulsion System Models," Proceedings of the VFS 6th Decennial Aeromechanics Specialists' Conference, Santa Clara, CA, February 6–8 2024.

14. Goericke, J., Brenner, F., and Lee, D., "Simulation and Analysis of Electric Motor Failure During eVtol Aircraft Operations in Turbulent Airwake," Proceedings of the 50th European Rotorcraft Forum, Marseille, France, September 10–12 2024.

15. Goericke, J., Smith, M. J., and Lee, D., "Simulation of AAM Aircraft Operations in Urban Turbulent Airwake Environments," Proceedings of the 51st European Rotorcraft Forum, Venice, Italy, September 10–12 2024.

16. Stutz, C. M., Hrynuk, J. T., and Bohl, D. G., "Effects of Vehicle Dynamics on Small UAS-Gust Encounters," *Aerospace Science and Technology*, Vol. 142, October 2023. DOI: 10.1016/j.ast.2023.108679.

17. Wall, A., Comeau, P., Barber, H., and Lebrasseur, J., "Toward the Development of Processes for Assessing the Response of RPAS and UAS to Complex Flow Fields," Proceedings of the Vertical Flight Society 81st Annual Forum, may 2025. DOI: 10.4050/F-0081-2025-0130.

18. Folk, S., Melton, J., Margolis, B., Yim, M., and Kumar, V., "Real-Time Safe and Efficient Navigation in Windy Urban Environments," https://ntrs.nasa.gov/citations/20250003994, 2025.

19. Larose, G., Labbad, M., Schajnoha, S., and Chen, J., "A Comparative Study of the Use of Experimental and Computational Techniques to Assess Turbulence Above the FATO to Support the Design and Operation of Vertiports in the Built Environment," Proceedings of the Vertical Flight Society 81st Annual Forum, may 2025. DOI: 10.4050/F-0081-2025-0186.

20. Kurban, E., Oates, B., Smith, M. J., and Rauleder, J., "Correlation of High- and Mid-Fidelity Ship Airwake Analysis to Experiment," *Journal of Aircraft*, Vol. In press, 2025. DOI: 10.2514/1.C038087.

21. Ashok, S. G., and Rauleder, J., "Mid-fidelity NATO Generic Destroyer Static and Moving-ship Airwake Simulations using the Lattice-Boltzmann Method Compared with Experimental Data," *Ocean Engineering*, Vol. 298, 2024, pp. 117264. DOI: 10.1016/j.oceaneng.2024.117264.

22. Crawford, A., *Fidelity Assessments in Computational Prediction of Unsteady Loads for Multi-Rotor Vehicles*, Ph.D. thesis, Georgia Institute of Technology, Atlanta, Georgia, 2025.

23. Ashok, S. G., and Rauleder, J., "Real-Time Coupled Ship–Rotorcraft Interactional Simulations Using GPU-Accelerated Lattice-Boltzmann Method," *Journal of the American Helicopter Society*, Vol. 70, (4), 2025, pp. 22005–25. DOI: 10.4050/JAHS.70.022005.

24. Nithya, D. S. and Rylko, A. and Muscarello, V. and Liang, M. and Quaranta, G., "Effects of Microscale Wind Disturbance on eVTOL Aircraft Performance during Landing," Proceedings of the 50th European Rotorcraft Forum, Marseille, France, September 10–12 2024.

25. Koukpaizan, N., Wilks, A., Afman, P., and Smith, M. J., "Rapid Vehicle Aerodynamic Modeling for Use in Early Design," Proceedings of the AHS Aeromechanics Design for Transformative Vertical Flight Conference, San Francisco, CA, January 16–19 2018.

26. Guo, X., Li, W., and Iorio, F., "Convolutional Neural Networks for Steady Flow Approximation," Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 2016.

27. Ribeiro, M. D., Rehman, A., Ahmed, S., and Dengel, A., "DeepCFD: Efficient Steady-State Laminar Flow Approximation with Deep Convolutional Neural Networks," *arXiv preprint arXiv:2004.08826*, 2020.

28. Hasegawa, K., Fukami, K., Murata, T., and Fukagata, K., "Machine-Learning-Based Reduced-Order Modeling for Unsteady Flows Around Bluff Bodies of Various Shapes," *Theoretical and Computational Fluid Dynamics*, Vol. 34, 2020, pp. 367–383.

29. Yousif, M. Z., and Lim, H.-C., "Reduced-Order Modeling for Turbulent Wake of a Finite Wall-Mounted Square Cylinder Based on Artificial Neural Network," *Physics of Fluids*, Vol. 34, (1), 2022.

30. Geneva, N., and Zabaras, N., "Transformers for Modeling Physical Systems," *Neural Networks*, Vol. 146, 2022, pp. 272–289.

31. Yousif, M. Z., Zhang, M., Yu, L., Vinuesa, R., and Lim, H., "A Transformer-Based Synthetic-Inflow Generator for Spatially Developing Turbulent Boundary Layers," *Journal of Fluid Mechanics*, Vol. 957, 2023, pp. A6.

32. Kim, B., Azevedo, V. C., Thuerey, N., Kim, T., Gross, M., and Solenthaler, B., "Deep Fluids: A Generative Network for Parameterized Fluid Simulations," Computer Graphics Forum, Vol. 38, 2019.

33. Solera-Rico, A., Sanmiguel Vila, C., Gómez-López, M., Wang, Y., Almashjary, A., Dawson, S. T., and Vinuesa, R., "β-Variational Autoencoders and Transformers for Reduced-Order Modelling of Fluid Flows," *Nature Communications*, Vol. 15, (1), 2024, pp. 1361.

34. Kingma, D. P., and Welling, M., "Auto-Encoding Variational Bayes," , 2013.

35. Rezende, D. J., Mohamed, S., and Wierstra, D., "Stochastic Backpropagation and Approximate Inference in Deep Generative Models," International conference on machine learning, 2014.

36. Kullback, S., and Leibler, R. A., "On Information and Sufficiency," *The Annals of Mathematical Statistics*, Vol. 22, (1), 1951, pp. 79–86.

37. Doersch, C., "Tutorial on Variational Autoencoders," *arXiv preprint arXiv:1606.05908*, 2016.

38. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A., "β-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework," International conference on learning representations, 2017.

39. Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., and Lerchner, A., "Understanding Disentangling in β-VAE," *arXiv preprint arXiv:1804.03599*, 2018.

40. O'shea, K., and Nash, R., "An Introduction to Convolutional Neural Networks," *arXiv preprint arXiv:1511.08458*, 2015.

41. Maas, A. L., Hannun, A. Y., Ng, A. Y., *et al.*, "Rectifier Nonlinearities Improve Neural Network Acoustic Models," Proc. icml, Vol. 30, 2013.

42. LeCun, Y., Bottou, L., Orr, G. B., and Müller, K.-R., "Efficient Backprop," *Neural networks: Tricks of the trade*, Springer, 2002, pp. 9–50.

43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I., "Attention is All you Need," Advances in Neural Information Processing Systems, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Vol. 30, 2017.

44. Ho, J., Jain, A., and Abbeel, P., "Denoising Diffusion Probabilistic Models," *Advances in Neural Information Processing Systems*, Vol. 33, 2020, pp. 6840–6851.

45. Salimans, T., and Ho, J., "Progressive Distillation for Fast Sampling of Diffusion Models," *arXiv preprint arXiv:2202.00512*, 2022.

46. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems," Software available from tensorflow.org, 2015.

47. Hunt, J. C., Abell, C. J., Peterka, J. A., and Woo, H., "Kinematical Studies of the Flows Around Free or Surface-Mounted Obstacles; Applying Topology to Flow Visualization," *Journal of Fluid Mechanics*, Vol. 86, (1), 1978, pp. 179–200.

48. Welch, P. D., "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms," *IEEE Transactions on Audio and Electroacoustics*, Vol. 15, (2), 1967, pp. 70–73.