


APPLICATION PAPER

MLP-mixer-based deep learning network for pedestrian-level wind assessment

Adam Clarke¹ , Knut Erik Teigen Giljarhus^{2,3} , Luca Oggiano³, Alistair Saddington¹  and Karthik Depuru-Mohan¹ 

¹Centre for Defence Engineering, Cranfield University, Defence Academy of the UK, Shrivenham, United Kingdom

²Department of Mechanical and Structural Engineering and Materials Science, University of Stavanger, Stavanger, Norway

³Nablaflow AS, Stavanger, Norway

Corresponding author: Adam Clarke; Email: adam.p.clarke@cranfield.ac.uk

Received: 24 June 2024; **Accepted:** 21 November 2024

Keywords: deep learning; pedestrian-level wind; surrogate modelling; wind comfort

Abstract

This article addresses the challenges of assessing pedestrian-level wind conditions in urban environments using a deep learning approach. The influence of large buildings on urban wind patterns has significant implications for thermal comfort, pollutant transport, pedestrian safety, and energy usage. Traditional methods, such as wind tunnel testing, are time-consuming and costly, leading to a growing interest in computational methods like computational fluid dynamics (CFD) simulations. However, CFD still requires a significant time investment for such studies, limiting the available time for design modification prior to lockdown. This study proposes a deep learning surrogate model based on a MLP-mixer architecture to predict mean flow conditions for complex arrays of buildings. The model is trained on a diverse dataset of synthetic geometries and corresponding CFD simulations, demonstrating its effectiveness in capturing intricate wind dynamics. The article discusses the model architecture and data preparation and evaluates its performance qualitatively and quantitatively. Results show promising capabilities in replicating key wind features with a mean error of 0.3 m/s and rarely exceeding 0.75 m/s, making the proposed model a valuable tool for early-stage urban wind modelling.

Impact Statement

This research introduces a deep learning surrogate model for pedestrian-level wind assessment in urban environments, offering a novel approach to address critical challenges in urban planning and design. By leveraging relatively simple machine learning techniques, the proposed model provides a rapid and cost-effective alternative to traditional wind assessment methods. The impact of this research has the potential to span across various sectors, including city planning, public health, and energy generation, where considerations for thermal comfort, pollutant transport, and pedestrian safety are paramount. The model's ability to accurately predict wind conditions in complex urban configurations can augment the early stages of urban development, enabling designers and planners to optimise building layouts for enhanced sustainability and human well-being. Moreover, the study contributes to the growing field of climate informatics, showcasing the potential of artificial intelligence in understanding and mitigating the impact of urban structures on microscale wind patterns. Overall, the research lays the foundation for a deep-learning model with the potential to shape more resilient and liveable urban spaces in the face of evolving climate challenges.

1. Introduction

Urban climates are significantly influenced by the morphology of the built environment and the presence of tall buildings impacting local wind patterns and posing challenges for residents and urban planners. Architects, engineers, and city planners must consider factors such as thermal comfort, pollutant transport, and pedestrian safety when proposing new urban developments (Moonen et al., 2012). High-wind conditions can render public spaces and infrastructure unusable and unsafe, affecting economic activities (Afe, 1970). To address these concerns, some city authorities mandate assessments of wind effects before issuing permits for new developments (City of London, 2019; Leeds City Council, 2021). Moreover, the phenomenon of the urban heat island effect plays an important role in overall energy consumption amidst trends in global climate as inhabitants strive to maintain comfortable living conditions. Ambient temperatures in cities can increase 4 °C above the surrounding area on average with peaks exceeding a 10 °C rise (Santamouris, 2015; Santamouris, 2016). Wind plays multiple roles in easing the rise in temperature through natural ventilation, convecting heat energy away from the area (Oke et al., 1991) while also significantly influencing individuals perception of thermal comfort (Nikolopoulou, 2004). Pedestrian-level wind (PLW) assessment is the quantification of resulting microclimate as drag effects create turbulent flow, buildings block or deflect flows, and street canyons artificially increase velocities. Wind tunnel testing using physical models has been used in the past to conduct such studies; however, owing to the time and financial costs incurred by this type of study, they are often delayed until designs are locked down in the later stages, favouring computational methods that can be used to rapidly estimate wind conditions and iterate through designs toward an optimal solution (Mittal et al., 2018). Effective PLW assessment requires meteorological data, aerodynamic information, and defined assessment criteria (Blocken et al., 2016). Many such wind comfort criteria exist (Isyumov and Davenport, 1977; Lawson, 1978; Melbourne, 1978; Willemsen and Wisse, 2007), and each stipulates comfort bands based on a threshold velocity and probability of exceedance. To ensure a comprehensive assessment of the area, studies usually incorporate data from up to 32 wind directions across multiple seasons and times of day, as the nature of the wind changes with respect to these variables (Hågbo and Giljarhus, 2022; Hågbo and Teigen Giljarhus, 2023). The requirement for many situations to be accounted for nudges the industry toward lower cost and rapid solutions such as computational fluid dynamics (CFD). Large Eddy simulations (LES) and Reynolds Averaged Navier–Stokes (RANS) are typically used with RANS being the preferred method for conducting PLW assessments (Blocken, 2018). RANS offers a time-averaged solution, providing the assessor with mean flow conditions. In contrast, LES directly solves turbulent flow down to the scale of the spatial discretization, relying on models to solve the finer details. It is known that RANS performs poorly in describing wind conditions in the low-velocity wake region of large buildings (Blocken et al., 2016). However, as low wind conditions do not pose threats to pedestrian safety or comfort, this shortcoming is often overlooked in favour of the reduced time and computational cost in the early stages of design assessment. There is a clear tolerance within the industry to trade some level of accuracy for increased cadence through the design iterations. This has motivated the exploration of surrogate models (Vasan and Saneinejad, 2023). Machine learning methods have become ubiquitous across many areas of science and engineering, owing to their versatility and powerful ability to fit models to data that are intractable to humans. Several studies have attempted to develop surrogate models with the ability to predict urban wind environments (Benmoshe et al., 2023; Hoiness et al., 2022; Mokhtar et al., 2020; Weerasuriya et al., 2021). In these studies, geometric information was provided to the models, while CFD results were used as a ground truth. Notably, the urban configurations in these studies are comparatively smaller than those presented here.

In this study, we propose a learned surrogate model that leverages global information communication methods to learn a mapping between a given set of boundary conditions and the expected mean flow conditions for complex arrays of bluff-shaped buildings. A novel extension to the model architecture that aids the learning process is introduced, and a number of model iterations are assessed for their efficacy in generating accurate flow fields.

2. Methodology

2.1. Problem formulation

The problem at hand involves the determination of the pedestrian-level flow field, framed as an image-to-image translation task. Specifically, the task is approached by leveraging a set of preprocessed geometries denoted as X , represented as grey-scale images. These images are input into a model f with the objective of learning a mapping to a pre-computed flow field Y , derived from a RANS simulation. Each pixel value in an individual image $x_i \in X$ encodes the heights of the buildings within the spatial domain. In contrast, the corresponding output image y_i is a 3-channel RGB image, where each channel represents the x, y, z velocity components at a 2 m height. It is crucial to note that the boundary conditions for the CFD solutions are uniform across all data points in Y and are implicitly incorporated into the model. The learning process, encapsulated in the mapping from $X \rightarrow Y$, is achieved by adjusting the model's free parameters denoted as θ . This adjustment is facilitated through backpropagation, employing a defined set of loss function equations:

$$f_{\theta} : x_i \rightarrow y_i \quad x_i, y_i \in X, Y \quad (1)$$

$$\min_{\theta} \left[\frac{1}{2} (\|y_i - f(x, \theta)\|_1 + \|\hat{y}_i - f(\hat{x}, \theta)\|_1) \right] \quad (2)$$

The loss function employed in this study aims to quantify the disparity between the model's output and the precomputed flow fields, utilising the mean absolute error (MAE). To further emphasise the significance of the central section in the evaluation process, an additional term has been introduced. This supplementary term accounts for the central portion of the image, \hat{x} , acknowledging its heightened importance in capturing nuanced details critical to the accuracy of the flow field prediction. The loss function is designed as the average of two components: the MAE computed across the entire image and the MAE calculated for a cropped central section. This approach ensures a balanced evaluation, promoting both overall accuracy and the model's ability to capture fine details in the central region of the flow field.

2.2. Geometries

The training dataset used in this study consists of 163 distinct synthetic geometries, each serving as the foundation for generating eight individual CFD scenarios. These scenarios account for wind flow from each of the cardinal and ordinal directions. The choice of a circular formation for each geometry introduces a deliberate uniformity among simulations, particularly in terms of wind interaction around the outer edge of the urban configuration.

The buildings within each geometry are characterised by simple, smooth prisms, featuring either sharp corners or fillet angles reminiscent of architectural styles found in early-stage design. Each building maintains a consistent cross-sectional area along its height axis. Notably, the architectural design intentionally omits intricate features such as bridges, skywalks, balconies, or masts. Moreover, it is crucial to highlight that the buildings in the synthetic geometries are entirely impermeable, contributing to a simplified yet representative simulation environment. Each geometry within the dataset exhibits variability in terms of density, building height, and the extent of open areas, providing a diverse range of scenarios for training the model.

2.3. CFD simulations

The CFD solutions in this study were generated using the open-source OpenFOAM v2206 software, utilising the simpleFOAM steady-state solver tailored for incompressible turbulent flow. To model turbulence closure, a $k-\epsilon$ model was employed, with specific coefficients, $C_{\mu}, C_{\epsilon 1}, C_{\epsilon 2}, \sigma_k$ set to 0.09, 1.44, 1.92, 1.11, respectively, as per the work of Hargreaves et al (Hargreaves and Wright, 2007).

The synthetic geometries were positioned at the centre of a cylindrical domain with a diameter of 3000 m and a height of 300 m. An internal orthogonal grid was used to refine the mesh within 75 m from the extremities of the geometry. This refinement not only enhanced cell quality but also provided a higher spatial resolution crucial for capturing the nuanced physics of the flow. The mesh resolution was adaptively adjusted, progressively increasing cell volumes with distance from the ground where a finer resolution was deemed unnecessary, thereby mitigating computational load. The mesh generation around the building geometries was facilitated by the in-built `snappyHexMesh` routine (OpenFOAM, n.d.).

The inflow conditions were characterised by a logarithmic profile with a reference velocity of 5 m/s at a height of 10 m. A slip boundary condition was applied at the top surface of the domain. A total of 1304 unique results were obtained by simulating wind flow from eight directions for each geometry, considering both cardinal and ordinal directions. The solver was run for 1000 iterations with a time step size Δt of 1 s chosen as a tradeoff between accuracy, convergence, and computational efficiency to generate a sufficiently robust training dataset.

2.4. Data preparation

For each of the 163 unique geometries, a transformation into a 2-dimensional grey-scale image is performed using the `PyVista` Python package (Bane Sullivan and Alexander Kaszynski, 2019), yielding images with dimensions $H \times W = 1024 \times 1024 \text{px}$. Each pixel represents an area of approximately 1m^2 and has a value within the range of $[0, 1]$ representing the scaled height of the building at that specific point in space.

Similarly, the precomputed flow fields undergo preparation also using the `PyVista` package. a slice of the XY plane is taken, capturing the velocity components of the flow field in separate channels as RGB pixel values. Specifically, the red, green, and blue channels encode the x, y, z velocity components, resulting in a tensor of size $H \times W \times C = 1024 \times 1024 \times 3 \text{px}$. A colour mapping is applied to the velocity components, with $[-6, 6] \rightarrow [0, 1]$ for x, y direction and $[-2, 2] \rightarrow [0, 1]$ for the z component. A sample training pair is shown in Figure 1.

To ensure consistency, each of the 1304 image representations of the velocity fields is rotated such that the wind inflow is from the top of the image. The relevant components are transformed using a rotation matrix congruent with their shift in the frame of reference. Augmentation is further applied by mirroring each rotational orientation about the Y -axis, resulting in the full 16 octagonal symmetries for each geometry. Consequently, the final training set comprises 2608 geometry-flow field pairs, providing a robust and diverse dataset for model training. To describe the velocity flow field in our analysis, we employ a Cartesian coordinate system with axes U_x , U_y , and U_z , representing the horizontal, vertical and

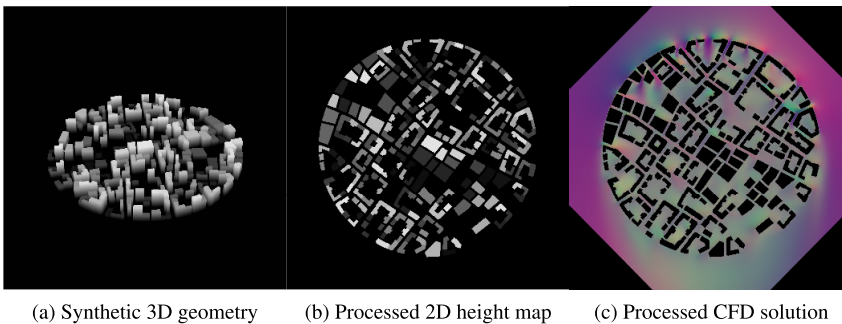


Figure 1. Training Data Overview: (a) 3D model of a synthetic geometry used to generate training data for the deep learning model. (b) Processed geometry, a 2D representation of the synthetic geometry after preprocessing, where pixel colour corresponds to the height of the buildings. (c) Postprocessed CFD Data where the RGB channels encode the velocity components in the x, y , and z directions. This image provides a visual representation of the ground truth used for training and evaluating the deep learning model.

depth components. Positive values indicate motion toward the right and top of the image for the U_x and U_y components and emerging from the page for the U_z component.

Given that the geometries undergo rotations at angles such as 45, 135, 225 degrees, etc., a notable consequence is the absence of data in the corners of the resulting images. While such data gaps do not pose an issue when rotations occur in multiples of 90 degrees, the irregular angles necessitate a corrective measure. To address this, a corner mask is systematically applied to all images. This mask effectively removes the data in the corners, ensuring uniformity across all orientations and mitigating confusion in the model, thereby enhancing the overall consistency and reliability of the dataset.

2.5. Model architecture

The model architecture draws inspiration from the image-to-image multilayer perceptron (MLP) mixer model introduced by Mansour et al. (Mansour et al., 2023). This supervised learning model, comprising solely of linear transformations, nonlinear activations, and data transpositions, demonstrated state-of-the-art performance in computer vision tasks such as reconstructing noisy images. Information transfer in this architecture is facilitated through MLP layers acting on linear transformations of input images across all spatial dimensions and token channels. Unlike a convolutional neural network, which inherently captures local spatial relationships due to their convolutional nature and receptive fields, giving them a strong inductive bias toward translation invariance while allowing them to learn hierarchical spatial representations, MLP mixers rely more on learning global interactions across image patches resulting in a lower inductive bias toward spatial relationships. Crucially, the number of model parameters scales linearly with the size of the input dimension, rendering it well-suited for handling higher-resolution images without the need for prior compression via auto-encoders or similar techniques.

Each input image undergoes a discretization process, where it is divided into discrete patches of size $P \times P$. Each patch is then transformed into an embedding vector with an arbitrary number of channels C . These latent vectors are amalgamated to form a tensor with dimensions $\frac{H}{P} \times \frac{W}{P} \times C$, preserving their relative position in the image. This tensor undergoes mixing operations in each of its three dimensions using a shared MLP-mixer block, consisting of two MLP layers in series with a GELU activation function separating them. The size of the single hidden layer in each MLP layer is proportional to the size of the layer input and is determined by a hyperparameter f .

Multiple mixer blocks can be stacked, performing repeated mixing operations as the model attends to different areas in each layer, increasing the network's learning power. After n such mixing layers, the latent tensor is transformed back to the desired image dimension. The image reconstruction is managed in the patch expansion step, where each transformed latent vector is expanded into a flattened patch of size CP^2 via a shared linear transformation. The grouped vectors are reshaped to form a tensor, restoring the height and width dimensions of the image. Finally, a 1×1 convolution layer collapses the channel dimension into the original three colour channels.

2.5.1. Architectural enhancements

Unique to this model formulation is an additional mixing step that incorporates information from immediate neighbours within a predefined area, enhancing the model's capacity to consider local context during the learning process. This is achieved using two 2-dimensional convolutional layers in series separated by a GELU activation. The size of the latent dimension between the two layers is governed by a hyperparameter f_c . Similar to the original mixer block, the modified mixer block maintains the size of the input tensor through its operation. A schematic of the modified model is shown in Figure 2.

2.6. Compute resource

The training and hyperparameter tuning processes were conducted on a Linux machine, leveraging the computational power of a single A100 GPU, accompanied by 16 CPU cores and 64GB RAM. The model underwent training for a total of 20 epochs, taking approximately 8 hours.

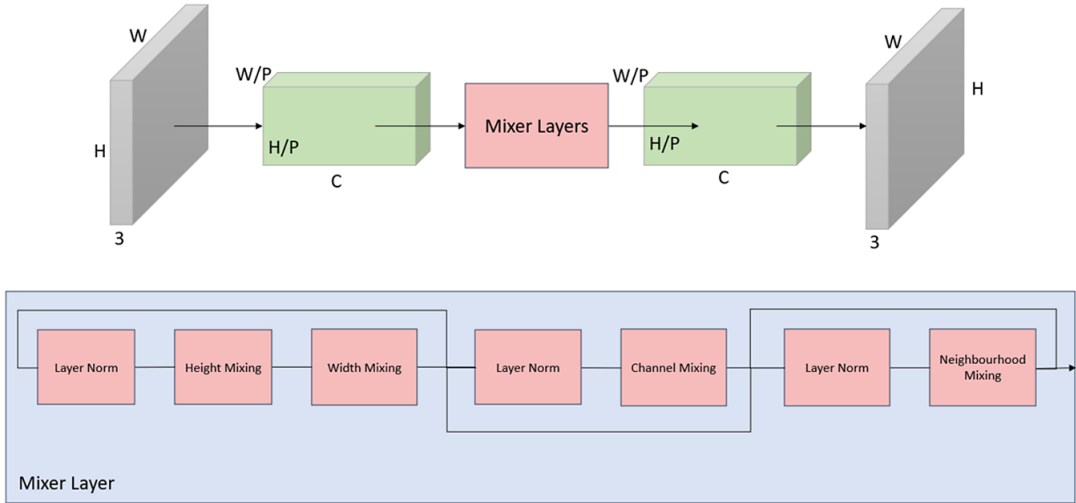


Figure 2. A schematic representation of the proposed modified image-to-image mixer model adapted from (Mansour et al., 2023). Details of the mixing layers are shown underneath.

3. Discussion of results

To gauge the effectiveness of the model, we conducted evaluations by comparing generated flow fields to their pre-computed counterparts for a reserved set of geometries that remained concealed from the model during the training phase. We begin the section with a qualitative comparison. Subsequently, quantitative measures are defined to systematically evaluate performance and suitability for the specific task of generating accurate flow fields at pedestrian-level height (2 m). Following the qualitative and quantitative analyses, we examine of the benefits of the neighbourhood mixing modification by using a standard, unmodified mixer model as per Mansour et al.’s work (Mansour et al., 2023) as a baseline for the assessment.

3.1. Qualitative comparison to CFD

From a qualitative perspective, we compared the model-generated flow patterns to those produced by CFD simulations. This evaluation aimed to assess the model’s capability to capture the expected physical behavior of the wind, particularly in areas characterised by high wind velocity amplification, often associated with increased discomfort and risk to pedestrians. This qualitative assessment involved a visual inspection of model predictions alongside the corresponding ground truth for each velocity component and magnitude. An illustrative example of this comparison is presented in Figure 3. Notably, the model exhibits proficiency in replicating key features around the geometry, including small areas of stagnation at the leading edge of the geometry and the wake region on the leeward side. As the wind circulates around the outer perimeter of the geometry, the model accurately captures regions of high velocity, particularly in gaps between buildings. Additionally, the model successfully reproduces amplified flows penetrating from the outer edge through canyons formed between buildings. In areas where courtyards exist, characterised by open spaces surrounded by buildings, the model appropriately reflects the increase in wind flow. Toward the centre of the geometry, the model demonstrates its ability to predict heightened flow velocities between buildings, even far from the free stream, after undergoing complex interactions.

Upon examining individual velocity components, the model effectively captures the majority of significant flow features for both the U_x and U_y directions. Particularly noteworthy is the model’s accurate representation of the reversal of flow direction in the U_y component, observed in the wake region and on the windward side of the geometry. An analysis of the U_z component reveals the model’s success in reproducing downwash on the windward side of the buildings throughout the entire geometry.

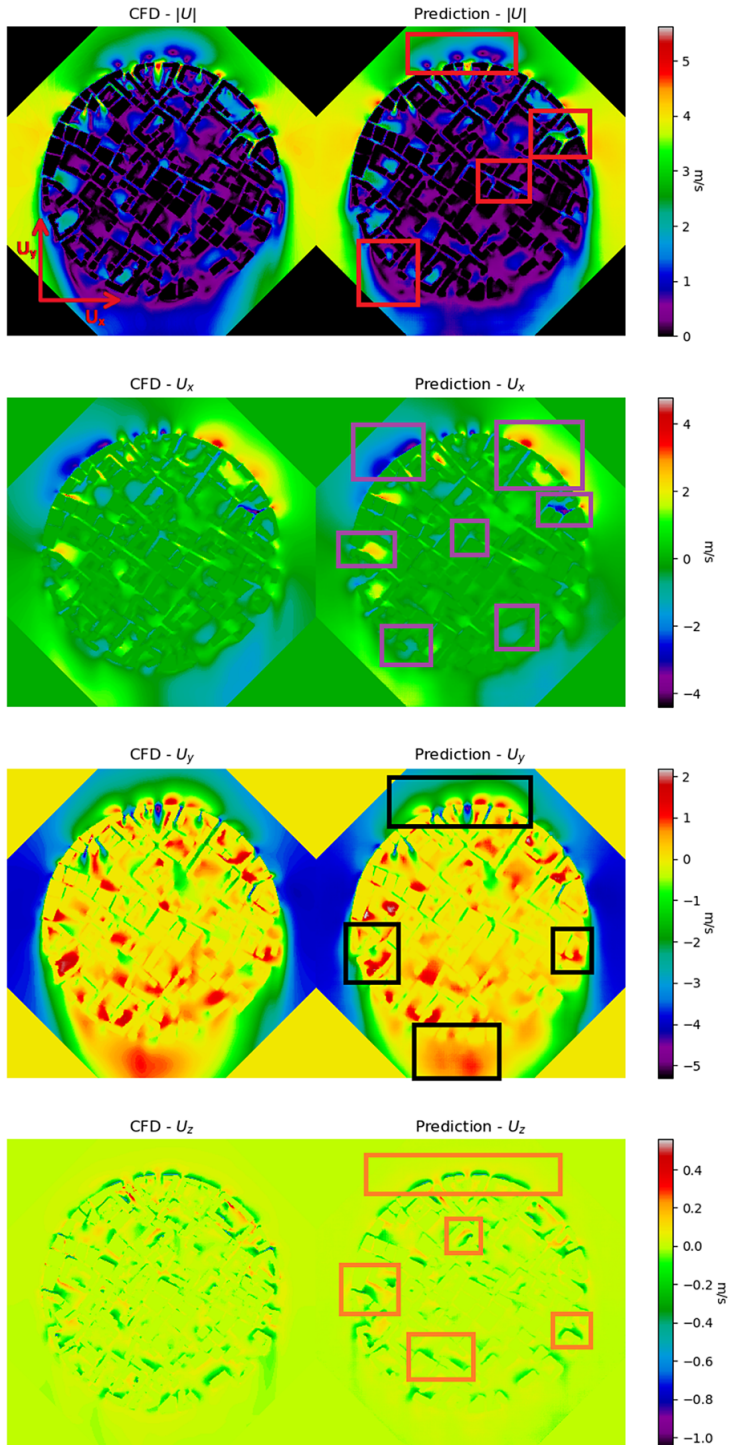


Figure 3. A qualitative comparison between the predicted flow field generated by the deep learning model and the corresponding CFD simulation. Highlighted areas pinpoint instances where the model successfully replicates essential features of the wind flow, providing valuable insights into its performance.

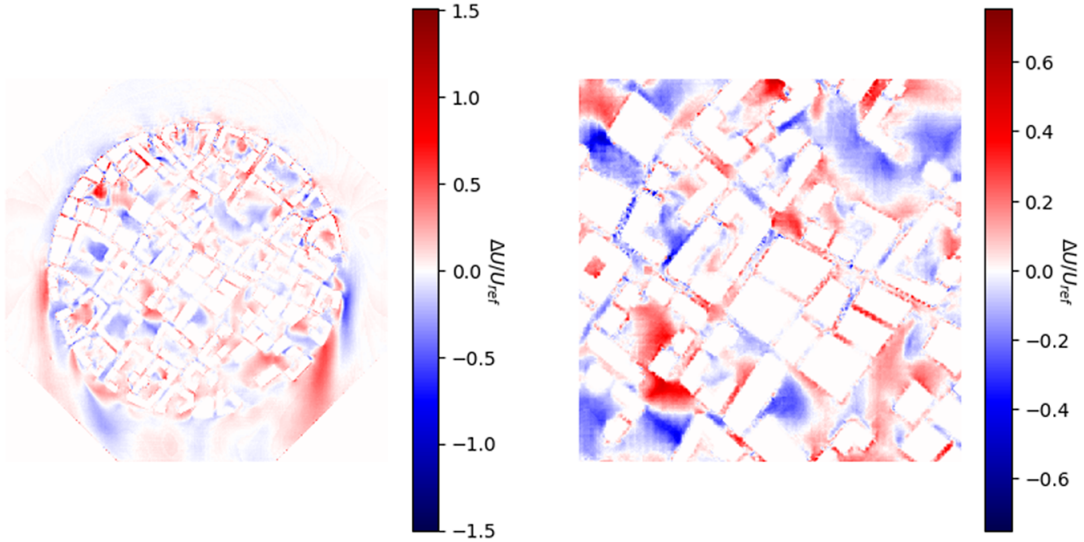


Figure 4. Error plots depicting the difference between the predicted magnitude generated by the deep learning model and the simulated magnitude from CFD for the entire image and the centre section. The magnitude difference is normalised by the reference velocity at a 2 m height.

Table 1. Performance comparison between the standard MLP mixer and the modified version measured on the test set. Mean absolute errors and the 90th percentile absolute errors for the entire image and the centre section, excluding building or masked corner pixels, provide insights into the accuracy and robustness of each model in predicting pedestrian-level wind conditions. Lower errors indicate superior performance.

Model	Full image		Cropped section	
	Mean absolute error	90 th % Absolute error	Mean absolute error	90 th % Absolute error
Standard	0.43	0.97	0.46	0.91
Neighbourhood mixing	0.32	0.75	0.35	0.71

The agreement between the model and the ground truth is reasonably high for the z-components, which aligns with expectations due to the naturally lower velocities associated with this component.

3.2. Quantitative comparison to CFD

To quantify the model’s performance, we initially plot the error $\|y_i\| - \|f(x_i, \theta)\|$ normalised by the reference velocity at a height of 2 m shown in Figure 4. In general, the model demonstrates commendable performance across the domain, with low errors typically falling within the range of ± 0.625 m/s. However, it becomes evident that while the majority of flow features are effectively captured, the precise shapes of these features and the predicted associated wind velocities can exhibit errors in the region of 2 m/s with rare occurrences exceeding 3.5 m/s. The level of under/over-prediction by the model is not consistent throughout the domain, displaying no clear pattern.

The findings from the error plots are further supported by descriptive statistics provided in Table 1.

3.3. Effect of neighbourhood mixing

Comparison between a model equipped with a neighbourhood mixing layer to the standard unmodified architecture reveals significant enhancements across all error metrics and in the representation of flow patterns. The quality of the comparison is particularly marked by a reduction in pixel-wise errors and effective mitigation of discontinuities in the generated flow fields. Moreover, the standard mixer exhibits a tendency to introduce artifacts into the generated image, implying valuable information can be captured in the immediate local area. Omitting this information leads to confusion and the presentation of behaviors that contradict physical laws such as discontinuities in the flow field. These improvements signify the efficacy of the introduced architectural modification in refining the model's ability to generate more accurate and coherent representations of wind behaviour. Consequently, this enhancement contributes to an overall improvement in the model's performance and predictive capabilities, which is clearly depicted in Figure 5.

3.4. Effect of training set size

Deep learning models are inherently shaped by the data they are trained on, making the volume and quality of the training dataset pivotal factors in determining model performance. In our study, we explore the impact of training set size by comparing models trained on varying proportions of the total dataset: 100%, 80%, 60%, and 40%. The mean absolute error (MAE) for each trained model is documented in Table 2. Notably, the largest decrease in loss is observed between the models trained on 40% and 60% of the dataset, indicating the importance of sufficient data for model effectiveness. Furthermore, as the proportion of training samples increases, there is a consistent linear decrease in loss, underscoring the direct correlation between data volume and model performance.

Despite these improvements, our analysis suggests that there remains potential for further enhancement, particularly through the expansion of the training dataset. Increasing the size of the training set could offer additional opportunities for refining model accuracy and generalisation, thereby optimising performance in PLW assessment tasks.

3.5. Effect of model size

We conducted training iterations of the model using different numbers of layers: 4, 6, 8, and 10. Interestingly, models with 4 and 6 layers exhibited comparable performance, with minimal disparity in

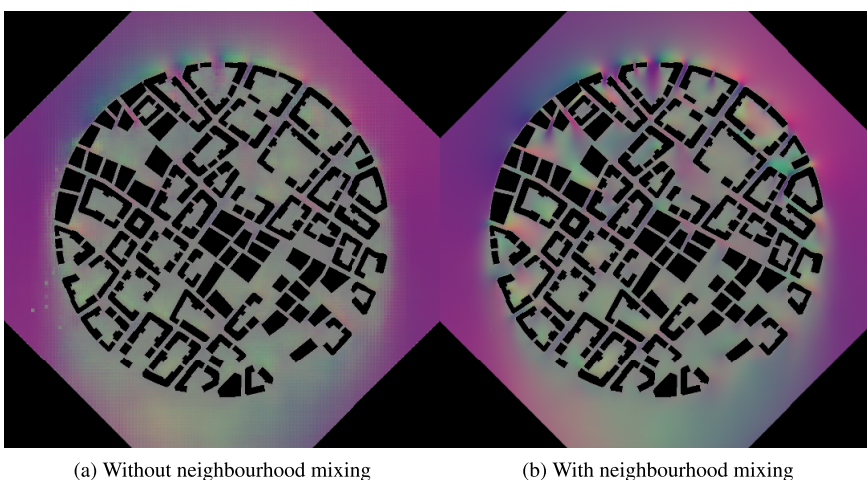


Figure 5. A side-by-side comparison of the predicted flow fields produced by the standard mixer model (a) and the proposed modified version (b).

Table 2. Top: Influence of training set size on model performance. Bottom: Influence of model size on performance

Training samples	Parameters (Millions)	Training loss (MAE)
1036	78.98	0.0902
1555	78.98	0.0791
2073	78.98	0.0759
2592	78.98	0.0717
Model layers		
4	34.119	0.0863
6	49.073	0.0860
8	64.026	0.0807
10	78.98	0.0717

MAE loss. However, as additional layers were added, the loss demonstrated a further decrease, suggesting the potential for continued improvement with larger models.

It is noteworthy, however, that increasing the number of layers introduces a risk of overfitting, particularly when working with a relatively small dataset. As the model parameters expand, so does the susceptibility to overfitting, where the model may excel in learning from the training data but struggle to generalise effectively to unseen data.

3.6. Inference time

CFD studies are known to be time-intensive. The data generated for model training and testing took 80–100 minutes to resolve a single training example. In contrast, our developed deep learning model achieves remarkable efficiency, with an average inference time of just 0.008 seconds per forward pass. This significant speed enhancement is particularly advantageous considering the multitude of wind directions that must be simulated for each design iteration in a PLW assessment. Consequently, our model holds considerable promise as a rapid and effective tool for preliminary design assessments, offering substantial time savings over traditional CFD approaches.

4. Conclusion

This article introduced a modified multilayer perceptron machine learning network designed to serve as a surrogate for the generation of accurate wind flow fields in the specific context of PLW assessment. The presented model demonstrates its capability to produce detailed, physically consistent, flow fields with a mean error of approximately 0.3 m/s for complex urban configurations. A notable feature of the model is its inherent capacity to capture long-range dependencies and positional information, crucial for making accurate inferences in the intricate urban wind environment. Given the broad ranges of the PLW assessment criteria, it could be feasible to use such a model as a viable tool for early-stage urban wind modelling, providing a rapid inference time while requiring a reasonably small set of training examples.

Open peer review. To view the open peer review materials for this article, please visit <http://doi.org/10.1017/eds.2024.44>.

Acknowledgments. The authors would like to express sincere gratitude to Franz Forsberg (<https://www.spacio.ai/>) and Nablaflow (<https://nablaflow.io/>) for generously providing the valuable geometry data and CFD solutions respectively, a significant contribution to the success of this project.

Mr Clarke is pleased to acknowledge the contribution of the IMechE Whitworth Senior Scholarship Award in supporting this research.

Author contribution. Adam Clarke: Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing—Original Draft, Visualisation, Project administration. Knut Erik Giljarhus: Conceptualisation, Software, Investigation, Writing—Review & Editing, Supervision, Funding acquisition. Luca Oggiano: Conceptualisation, Writing—Review & Editing, Funding

acquisition. Alistair Saddington: Writing—Review & Editing, Supervision. Karthik Depuru-Mohan: Conceptualisation, Writing—Review & Editing, Supervision, Project administration, Funding acquisition.

Competing interest. None.

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

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Funding statement. This research was supported by funding provided by Cranfield University and Nablaflow AS (<https://nablaflow.io/>).

Provenance. This article was accepted into the Climate Informatics 2024 (CI2024) Conference. It has been published in Environmental Data Science on the strength of the CI2024 review process.

References

- Afe W (1970) Wind effects due to groups of buildings. *Royal Society Symposium Architectural Aerodynamics*, 26–27.
- Bane S and Alexander K (2019) PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *The Open Journal* 4(37), 1450. <https://doi.org/10.21105/joss.01450>
- Benmoshe N, Fattal E, Leitl B and Arav Y (2023) using machine learning to predict wind flow in urban areas. *Atmosphere* 14, 990. <https://doi.org/10.3390/atmos14060990>
- Blocken B, Stathopoulos T and van Beeck JPAJ (2016) Pedestrian-level wind conditions around buildings: Review of wind-tunnel and CFD techniques and their accuracy for wind comfort assessment. *Building and Environment* 100, 50–81. <https://doi.org/10.1016/j.buildenv.2016.02.004>
- Blocken B (2018) LES over RANS in building simulation for outdoor and indoor applications: a foregone conclusion? *Building Simulation* 11(5), 821–870. <https://doi.org/10.1007/s12273-018-0459-3>
- City of London (2019, August) *Wind Microclimate Guidelines for Developments in the City of London* (tech. rep.). Available at <https://www.cityoflondon.gov.uk/assets/Services-Environment/wind-microclimate-guidelines.pdf> (accessed 19 December 2022)
- Hågbo T-O and Giljarhus KET (2022) Pedestrian wind comfort assessment using computational fluid dynamics simulations with varying number of wind directions. *Frontiers in Built Environment* 8, 858067. <https://doi.org/10.3389/fbuil.2022.858067>
- Hågbo T-O and Teigen Giljarhus KE (2024, September) Sensitivity of urban morphology and the number of CFD simulated wind directions on pedestrian wind comfort and safety assessments. *Building and Environment*, 253, 111310, ISSN 0360-1323, <https://doi.org/10.1016/j.buildenv.2024.111310>.
- Hargreaves DM and Wright NG (2007) On the use of the k-ε model in commercial CFD software to model the neutral atmospheric boundary layer. *Journal of Wind Engineering and Industrial Aerodynamics* 95(5), 355–369. <https://doi.org/10.1016/j.jweia.2006.08.002>
- Hoeiness H, Gjerde K, Oggiano L, Giljarhus KET and Ruocco M (2022, January) Positional Encoding Augmented GAN for the Assessment of Wind Flow for Pedestrian Comfort in Urban Areas [arXiv:2112.08447 [cs]]. Available at <http://arxiv.org/abs/2112.08447> (accessed 21 October 2022)
- Isyumov I and Davenport AG (1977) The ground level wind environment in built-up areas. In *Proceedings of the Fourth International Conference on Wind Effects on Buildings and Structures: Heathrow 1975*. Cambridge University Press.
- Lawson TV (1978) The wind content of the built environment. *Journal of Wind Engineering and Industrial Aerodynamics* 3(2), 93–105. [https://doi.org/10.1016/0167-6105\(78\)90002-8](https://doi.org/10.1016/0167-6105(78)90002-8)
- Leeds City Council (2021) Draft Wind and Microclimate Toolkit. Available at <https://www.leeds.gov.uk/docs/Draft%20wind%20and%20microclimate%20toolkit.pdf> (accessed 15 January 2024)
- Mansour Y, Lin K and Heckel R (2023) Image-to-Image MLP-Mixer for Image Reconstruction. Available at <https://openreview.net/forum?id=wsuQ2h6KZXQ> (accessed 7 April 2023)
- Melbourne WH (1978) Criteria for environmental wind conditions. *Journal of Wind Engineering and Industrial Aerodynamics* 3(2), 241–249. [https://doi.org/10.1016/0167-6105\(78\)90013-2](https://doi.org/10.1016/0167-6105(78)90013-2)
- Mittal H, Sharma A and Gairola A (2018) A review on the study of urban wind at the pedestrian level around buildings. *Journal of Building Engineering* 18, 154–163. <https://doi.org/10.1016/j.jobe.2018.03.006>
- Mokhtar S, Sojka A and Davila CC (2020) Conditional generative adversarial networks for pedestrian wind flow approximation. *Society for Modeling & Simulation International (SCS), SimAUD* 2020, 8.
- Moonen P, Defraeye T, Dorer V, Blocken B and Carmeliet J (2012) Urban physics: effect of the micro-climate on comfort, health and energy demand. *Frontiers of Architectural Research* 1(3), 197–228. <https://doi.org/10.1016/j.foar.2012.05.002>
- Nikolopoulou M (2004) *Designing Open Spaces in the Urban Environment: A Bioclimatic Approach*. Centre for Renewable Energy Sources, EESD, FP5.

- Oke TR, Johnson G, Steyn D and Watson I** (1991) Simulation of surface urban heat islands under ‘ideal’ conditions at night part 2: diagnosis of causation. *Boundary-Layer Meteorology* 56, 339–358.
- OpenFOAM** (n.d.) OpenFOAM: User Guide: Snappyhexmesh. Available at <https://www.openfoam.com/documentation/guides/latest/doc/guide-meshing-snappyhexmesh.html> (accessed 24 January 2024)
- Santamouris M** (2015) Analyzing the heat island magnitude and characteristics in one hundred asian and australian cities and regions. *Science of The Total Environment* 512–513, 582–598. <https://doi.org/10.1016/j.scitotenv.2015.01.060>
- Santamouris M** (2016) Innovating to zero the building sector in europe: minimising the energy consumption, eradication of the energy poverty and mitigating the local climate change [Special issue: progress in Solar Energy]. *Solar Energy* 128, 61–94. <https://doi.org/10.1016/j.solener.2016.01.021>
- Vasan N and Saneinejad S** (2023) State of the art in wind consulting: a perspective on the application of current advances in CFD and AI. In *16th International Conference on Wind Engineering*, Florence, Italy: Springer Nature.
- Weerasuriya AU, Zhang X, Lu B, Tse KT and Liu CH** (2021) A gaussian process-based emulator for modeling pedestrian-level wind field. *Building and Environment* 188, 107500. <https://doi.org/10.1016/j.buildenv.2020.107500>
- Willemsen E and Wisse JA** (2007). Design for wind comfort in the netherlands: procedures, criteria and open research issues. *Journal of Wind Engineering and Industrial Aerodynamics* 95(9).