

Diffusion models for optimal sensor placement and sparse reconstruction for simplified urban flows

By A. Vishwasrao[†], S. Gutha[‡], A. Patil, K. Wijk[‡], B. J. McKeon, C. Gorle,
H. Azizpour[‡] AND R. Vinuesa[†]

Wind patterns within urban canopies profoundly impact pollution dispersion, drone navigation, heat distribution, and overall urban microclimates, making understanding and monitoring these flows crucial. Traditionally, this has been achieved through deploying sensors, the optimized placement of which is essential due to the high costs and practical constraints involved and then performing flow-field reconstruction based on these sparse measurements. This paper explores advanced methodologies for sparse reconstruction based on denoising diffusion probabilistic models (DDPMs). DDPMs offer a stable and effective approach to state-of-the-art generative modelling. Despite their success in natural image generation, their application to turbulent flows has been limited. We extend the use of DDPMs to flow modelling by applying wrapper models such as denoising diffusion restoration models (DDRM) and pseudoinverse-guided diffusion models (IIGDM) to the sparse reconstruction of flow fields. Additionally, we propose a novel framework for optimal sensor placement utilizing Shapley additive explanations (SHAP). These methodologies have shown great potential by significantly outperforming conventional linear methods while maintaining high fidelity as compared with direct numerical simulation (DNS) data.

1. Introduction

The rapid urbanization of cities has introduced significant challenges in infrastructure development, particularly in managing wind patterns and coherent structures within the urban canopy. Studies highlight their importance in public health, with air pollution in cities contributing to 7 million premature deaths annually worldwide (WHO 2024). Hence, understanding urban wind patterns is essential for optimizing urban planning, improving public health, and enhancing the efficiency of emerging technologies like autonomous drones (Torres et al. 2021).

To gain insights into complex urban wind fields, sensor deployment within the urban canopy has long been a standard approach, providing the empirical data needed to model and predict wind behaviour. Fluid flows, governed by the Navier–Stokes equations, are best analyzed using numerical methods. While large-eddy and direct numerical simulations (LES and DNS, respectively) can provide accurate data, their computational demands hinder real-time monitoring. Data assimilation-based sparse-reconstruction methods help by estimating flow fields from limited sensor data, but as these methods are linear in nature, they have two major limitations: accurately capturing turbulent structures and optimizing sensor placement (OSP) within urban constraints. Solving these is key to improving wind predictions and sensor efficiency.

[†] FLOW-SCI, KTH Royal Institute of Technology, Sweden

[‡] RPL-EECS, KTH Royal Institute of Technology, Sweden

Denoising diffusion probabilistic models (DDPMs) are probabilistic models that generate data by reversing a diffusion process. This results in high-quality samples, comparable to those of generative adversarial networks (GANs, see Güemes et al. (2021)), but with more stable training grounded in likelihood estimation (Ho et al. 2020). While diffusion models have excelled in image generation, their application to turbulent flow data remains underexplored. Most studies focus on generating spatiotemporal flow data using conditional diffusion models (Shu et al. 2023; Du et al. 2024; Gao et al. 2024; Zhuang et al. 2024). While the existing literature highlights the potential of diffusion models for fluid flow reconstruction, the field is still in its early stages, with few studies addressing the complexities of turbulent flows.

This work explores diffusion models for sparse-reconstruction and OSP in turbulent flows. We focus on models like denoising diffusion restoration models (DDRMs) (Kawar et al. 2022) and pseudoinverse-guided diffusion models (PIGDMs) (Song et al. 2022) for sparse reconstruction. For sensor placement, we propose a coupled framework based on Shapley values (Shapley 1953) and the above-mentioned diffusion models. These methods operate as wrappers on pretrained models, reducing computational costs while maintaining efficiency.

2. Methodology

This section details the three major constituents of the proposed framework, namely training a prior model, sparse reconstruction and optimal sensor placement.

2.1. Diffusion models for unconditional data generation

We use DDPMs as generative priors for turbulent flows. In DDPM, Gaussian noise is iteratively added to a sample $\mathbf{x}_0 \in \mathbb{R}^n$ in the forward process $\mathbf{x}_0 \rightarrow \mathbf{x}_1 \rightarrow \dots \rightarrow \mathbf{x}_T$ after which the model is tasked with learning the reverse process $\mathbf{x}_0 \leftarrow \mathbf{x}_1 \leftarrow \dots \leftarrow \mathbf{x}_T$ that denoises \mathbf{x}_T . For certain transition probabilities in the forward process (see, Eq. (2.1b)), $p(\mathbf{x}_T)$ is a standard Gaussian distribution. Consequently, passing the noisy samples $\mathbf{x}_T \sim p(\mathbf{x}_T)$ through the learned reverse process creates clean data samples \mathbf{x}_0 . Note that $(\beta_0, \beta_1, \beta_2, \dots, \beta_T)$ denotes a monotonically increasing variance schedule with $\beta_0 = 0$, $0 < \beta_i < 1, \forall i \in \{1..T\}$. This is also referred to as the variance-preserving formulation. The forward process follows a Markov chain, with its transition kernel:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad (2.1a)$$

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}). \quad (2.1b)$$

Similarly, the reverse process is a Markov chain with the reverse transition kernels approximated to be Gaussian:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t), \quad (2.2a)$$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)), \quad p(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2.2b)$$

where, p_θ denotes a parameterized distribution with parameters θ . The model is trained using a variational objective that simplifies to an \mathcal{L}_2 loss between the original standard normal noise ϵ_t and the model's prediction of this noise $\hat{\epsilon}_t = \epsilon_\theta(\mathbf{x}_t, t)$ from a noisy sample $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon_t$, where \mathbf{x}_0 is a clean data sample and $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$. We train the DDPM by treating each snapshot as an independent and identically distributed

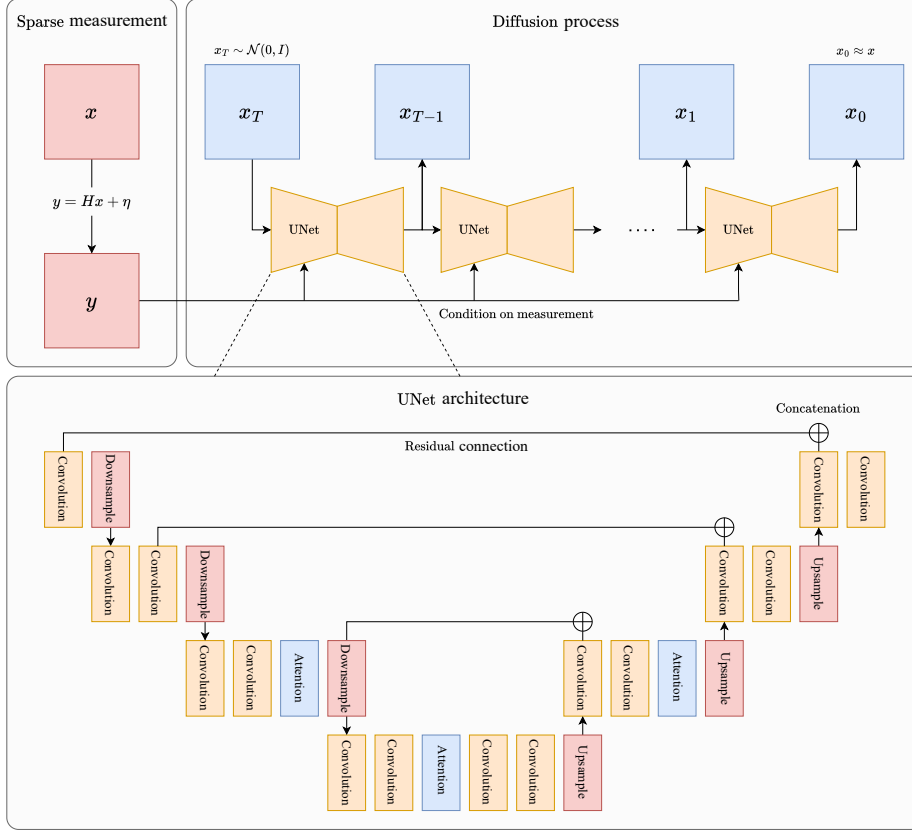


FIGURE 1. A schematic of the sparse reconstruction process, combining the diffusion process and UNet architecture. At each diffusion timestep, the UNet receives a noisy image as input and outputs the denoised image.

(IID) sample of the underlying distribution of snapshots represented by the training set. More details regarding the simulation dataset are stated in Section 3.1.

2.2. Diffusion models for sparse reconstruction

We turn our attention to the reconstruction problem in Eq. (2.3) in which the goal is to reconstruct an instantaneous flow field $\mathbf{x} \in \mathbb{R}^n$ from partial measurements $\mathbf{y} \in \mathbb{R}^m$ obtained through the forward operator $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. In this work, we focus on noiseless inpainting, or sparse reconstruction, with the forward operator $\mathcal{A}(\mathbf{x}) = \mathbf{H}\mathbf{x}$, where \mathbf{H} is a binary measurement matrix $\mathbf{H} \in \{0, 1\}^{m \times n}$:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \boldsymbol{\eta}. \quad (2.3)$$

We solve the sparse-reconstruction problem using a pretrained diffusion model as the prior $P(x)$ and generate x given y . This approach requires no additional training of the model. To this end, we explore two different methods: (1) denoising diffusion restoration models and (2) pseudoinverse-guided diffusion models. Note that typically $\boldsymbol{\eta} \sim \mathcal{N}(0, \sigma_y^2 \mathbf{I})$.

2.2.1. Denoising diffusion restoration models

DDRMs use the variance exploding (VE) formulation of DDPM to design the conditional forward process in the spectral space of \mathbf{H} . A typical state transition kernel of a VE

DDPM is described in Eq. (2.4). The mapping from the pixel space to the spectral space of \mathbf{H} is given by $\bar{x} = \mathbf{V}^T x$, and $\bar{y} = \mathbf{\Sigma}^\dagger \mathbf{U}^T y$, using the singular value decomposition (SVD) of $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; x_{t-1}, (\sigma_t^2 - \sigma_{t-1}^2)\mathbf{I}), \quad (2.4a)$$

$$q(x_t|x_0) = \mathcal{N}(x_t; x_0, \sigma_t^2\mathbf{I}). \quad (2.4b)$$

2.2.2. Pseudoinverse-guided diffusion models

Song et al. (2021) introduced the continuous generalization of diffusion models using stochastic differential equations (SDEs). The forward process, which is denoted by an SDE, gradually corrupts a clean data sample to a completely noisy sample. The reverse process, which is also represented as an SDE is then solved using SDE solvers to generate new samples from the data distribution. Conditional generation involves sampling from the posterior $p(x|y)$, which translates to solving the modified reverse SDE in Eq. (2.5). From Eq. (2.6), $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$ is learned by the unconditional diffusion model; however, the term $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$ is intractable (Eq. 2.7) since $p(\mathbf{y}|\mathbf{x}_0)$ is tractable (Eq. 2.3) but $p(\mathbf{x}_0|x_t)$ is intractable:

$$dx = [\mathbf{f}(x, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}|\mathbf{y})]dt + g(t)d\bar{w}, \quad (2.5)$$

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t), \quad (2.6)$$

$$p(\mathbf{y}|\mathbf{x}_t) = \int_{\mathbf{x}_0} p(\mathbf{y}|\mathbf{x}_0, \mathbf{x}_t) p(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0 = \int_{\mathbf{x}_0} p(\mathbf{y}|\mathbf{x}_0) p(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0. \quad (2.7)$$

PIGDM makes a Gaussian approximation such as $p(\mathbf{x}_0|\mathbf{x}_t) = \mathcal{N}(\mu_t, r_t^2\mathbb{I})$, where $\mu_t = \mathbb{E}(\mathbf{x}_0|\mathbf{x}_t) = \mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) \approx \mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} S(\mathbf{x}_t, t)$, where $S(\cdot, \cdot)$ denotes the score function learned by the pretrained unconditional diffusion model. The standard deviation r_t is set proportional to σ_t . This results in $p(\mathbf{y}|\mathbf{x}_t)$ being Gaussian, hence tractable. In our experiments, we set $r_t = \sigma_t$. We refer to Song et al. (2022) for additional details regarding the sampling and using DDPM checkpoints with PIGDMs.

2.3. SHAP values for optimal sensor placement

Shapley Additive Explanations (SHAP, see Lundberg and Lee 2017) is a framework for interpreting machine-learning models by utilizing Shapley values from game theory (Shapley 1953) to quantify feature contributions. It has been successfully applied to various fields (Jeyakumar et al. 2020) and also in fluid mechanics, to identify coherent structures in wall-bounded turbulent flows using a UNet architecture (Cremades et al. 2024). In this work, we employ SHAP to generate a sensor-importance matrix for OSP in turbulent flow reconstruction, representing the model’s estimated prediction error as a linear combination of sensor contributions through an additive-feature attribution method. Detailed explanations regarding Shapley value computations can be found in Cremades et al. (2024) and Lundberg and Lee (2017). Here, shapley values are computed across 1000 randomly selected instantaneous fields for 756 plausible sensor locations (see, Section 3.4.1), combinations of which are utilised by reduced order model (ROM) based on DDRMs or PIGDMs (see, Section 2.2) as inputs for accurate sparse reconstruction. Practically, the SHAP framework learns a kernel function for optimized sensor locations based on the iterative computation of estimated prediction error for masks within the range from a full mask (maximum sensor selection region) to a zero mask (no sensors). This approach allows flexible sensor-placement strategies, enabling better consideration of physical and economic constraints.




FIGURE 2. Instantaneous visualization of the flow around a wall-mounted square obstacle (Martínez-Sánchez et al. 2023) and extraction of 2D plane from the 3D domain at $z/h = 0$.

3. Results and discussion

3.1. Numerical simulation and flow description

We use DNS data of flow around a wall-mounted square cylinder, representing a simplified urban environment and simulating flow around a building. The DNS dataset is generated using the open-source code Nek5000, which solves the incompressible Navier–Stokes equations with the spectral element method. The square obstacle has a width-to-height ratio (b/h) of 0.25, and the simulation is performed at a Reynolds number of 2000, based on obstacle height h . The spectral element mesh consists of 21.8 million grid points, and the computational domain is extracted near the obstacle, covering $-1 \leq x/h \leq 5$, $0 \leq y/h \leq 2$, and $-1.5 \leq z/h \leq 1.5$. Data is spectrally interpolated onto a uniform mesh with $(N_x, N_y, N_z) = (300, 100, 150)$. Temporal parameters are expressed in convective time units, with a time step $\Delta t_s = 0.005$, providing adequate resolution for capturing flow phenomena. The dataset comprises 30,000 snapshots, equivalent to 150 convective time units after the removal of transient time steps, and is split into a 5:1 ratio for training and testing, preserving temporal order to simulate real-world scenarios where future flow states are predicted from past observations. More details regarding the simulation setup can be found in Martínez-Sánchez et al. (2023). Mean fields are subtracted from the velocity fields, and the DDPM exclusively focuses on predicting turbulent fluctuations, increasing generative model accuracy compared to conventional methods. While the original simulation covers a 3D domain, this study focuses on a 2D section at $z/h = 0$, balancing computational efficiency and model complexity. The goal is to eventually extend these methods to 3D flow fields.

3.2. Statistical evaluation of stochastically generated samples

The primary aim of this study is to develop a deep generative model that serves as a probabilistic prior, enabling the generation of synthetic turbulent flow data by learning the underlying distribution from a simulation. For this purpose, we utilize a denoising diffusion probabilistic model. To assess the performance of this prior, we randomly sample 25,000 snapshots—equivalent to the training dataset size—using the pretrained diffusion model without any conditioning. We then compare key turbulent statistical quantities between the generated samples and the ground truth DNS data to evaluate the accuracy of the generative model.

Figure 3 provides a qualitative comparison between a randomly selected spatiotem-

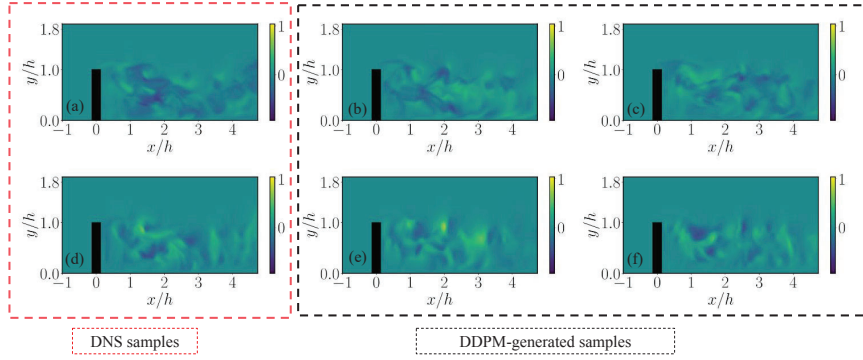


FIGURE 3. Instantaneous visualizations of randomly selected DNS samples (a) and (d) compared with DDPM-generated samples (b),(c) and (e), (f), streamwise and wall-normal velocity fluctuations, respectively.




FIGURE 4. Comparison of Reynolds stresses on the entire 2D plane: colored contours represent stress values computed based on DNS data, while the black outlines represent contour based on DDPM-generated samples.

poral DNS sample and probabilistically generated DDPM samples. Figure 4 compares the Reynolds stresses across the entire 2D plane. Additionally, Figure 5 presents joint probability density functions (JPDFs) of velocity fluctuations along the streamwise direction at $y/h = 0.5$. The diffusion-based approach for simulating complex turbulent flows demonstrates robust flexibility and high accuracy in reproducing essential turbulent characteristics, particularly with an impressive statistical agreement in the wake region behind the obstacle, where only minor discrepancies exist in the streamwise direction, and we expect these deviations to further diminish as we integrate additional timestep-specific measurements in Section 3.3.

3.3. Diffusion-based sparse reconstruction

After establishing the diffusion model as a reliable prior for generating synthetic turbulent flow data, we move on to perform conditional generation based on masked sensor inputs. For this, we design a mask that considers only m per cent of the flow-field information. The rationale behind the mask design is explained in detail in Section 3.4.1. We evaluate two wrapper techniques based on diffusion models: DDRMs and IIGDMs. For testing, we randomly select 2000 samples from the test dataset, representing unseen future time steps. Both methods use different conditioning strategies for the masked inputs but maintain the same core diffusion model. We reconstruct 2000 snapshots conditioned on these designed masks. We then compute the relative mean-squared error (MSE), *i.e.*, the MSE normalized by square of maximum velocity fluctuation magnitude:




FIGURE 5. JPDFs between velocity fluctuations: streamwise (a) and wall-normal (b), with streamwise direction (x/h) at $y/h = 0.5$. Colored lines show JPDFs computed on DNS data, while dotted lines represent those corresponding to DDPM-generated data.




FIGURE 6. MSE (ε) averaged over 2000 test samples using IIGDM for streamwise (a) and wall-normal (b) velocity fluctuations with $m = 30\%$

$$\varepsilon_i = \frac{1}{N} \sum_{i=1}^N \frac{(u_{i,\text{DNS}} - u_{i,\text{Pred}})^2}{\max(u_{i,\text{DNS}})^2}. \quad (3.1)$$

Here, u_1 and u_2 represent streamwise and wall-normal velocity fluctuations, respectively. Figure 6 shows MSE results for IIGDM on the streamwise and wall-normal velocity fluctuations. We noted that IIGDM outperforms DDRM, achieving $\varepsilon \sim 1.5\%$ compared to DDRM's $\varepsilon \sim 3\%$, owing to its more advanced conditioning. Both methods perform better on streamwise components, likely since streamwise direction dominates the momentum transfer. Furthermore, Figure 7 shows a temporal instance generated by IIGDM with $m = 30\%$. While the prediction closely matches DNS data ($\varepsilon < 0.6\%$), IIGDM struggles to maintain accuracy across all instances, likely due to the probabilistic nature of the diffusion models. Future work could explore advanced conditioning, physics-based constraints, or reducing inference timesteps to improve consistency across sequential predictions and address temporal coherence issues.

3.4. Optimal sensor placement based on Shapley values

3.4.1. Design and segmentation of the mask

A key motivation for this work is to develop a sensor-selection process that is both effective and flexible enough to handle real-world challenges in sensor placement. To address these complexities, we focus on a modular mask that ensures high-fidelity predictions while maintaining flexibility for practical deployment. Traditional linear sensor placement methods, such as QR-pivoting (Brunton and Kutz 2022), prioritize sensors in




FIGURE 7. IIGDM prediction (b) compared with ground truth (a) for randomly selected instance, mask $m = 30\%$.




FIGURE 8. The mask (a) and segmentation (b) tensors, used for sparse reconstruction and OSP. Different colors in the bottom-left panel indicate pixel groupings equivalent to sensor locations. Sensor importance (c) based on SHAP values.

high-energy regions, often within the wake of urban canopies. However, deploying sensors in these regions can disturb the flow, making them impractical for real-world use. Without these sensors, linear methods perform suboptimally, as they depend on high-energy measurements that are inaccessible in such scenarios. To demonstrate this, we design a mask that limits sensor placement to practical regions, such as near the ground and walls (see, Figure 8). This modular-mask approach allows flexibility in sensor placement by subdividing the region into 756 smaller areas, each corresponding to a potential sensor with an area of $0.0024h^2$ and six pixels. This granularity enables the grouping of pixel regions into single sensors, accommodating real-world constraints and affording scalability to 3D applications in future.

3.4.2. The SHAP values

In this study, we performed SHAP analysis on the segmented mask, generating a matrix of SHAP values for 756 sensor locations to represent individual contributions to the flow-field reconstruction. These values were normalized, and masks were created based on cumulative sensor addition, with sensor count decreasing as normalized SHAP values increased from 0.4 to 0.9. Higher SHAP values indicate greater sensor importance, *i.e.*, the sensor with the highest SHAP value will be considered in all six masks, as visualized in Figure 8. Using six masks corresponding to SHAP values between 0.4 (118 sensors) and 0.9 (10 sensors), we analyzed 1000 temporal snapshots from the test dataset with both DDRMs and IIGDMs. Figure 9 illustrates the relative MSE (ε) for sensor and pixel counts, comparing results from SHAP-derived masks and QR-pivoting for the same snapshots. QR-pivoting tends to scatter sensors, failing to identify crucial regions in the flow. The results show that SHAP-based masks achieve comparable accuracy to QR-pivoting while using roughly 100 fewer sensors. Notably, QR-pivoting was applied on the DNS dataset, while SHAP analysis used data from IIGDMs, potentially favoring QR-pivoting. A direct comparison would likely reveal an even greater advantage for SHAP-based methods, as the diffusion model already ensures high reconstruction accuracy, whereas linear

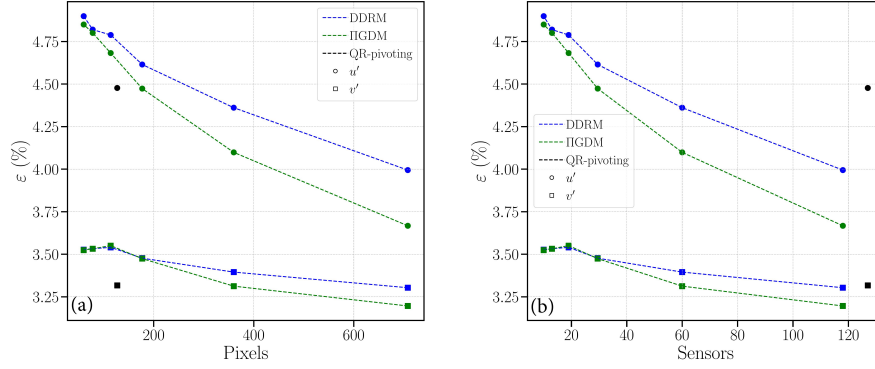


FIGURE 9. Comparative evolution of ϵ for DDRM and IIGDM predictions based on pixel (a) and sensor (b) counts. Six masks were obtained based on SHAP analysis and one from QR-pivoting.




FIGURE 10. Comparison of relative MSE (ϵ) for stochastically generated DDPM snapshots (zero sensors), QR-pivoting mask with 127 sensors, and SHAP masks with a range of sensors.




FIGURE 11. DNS data (a), IIGDM prediction (b) and corresponding MSE (ϵ) for a randomly selected instance (c) for SHAP mask with 118 sensors.

QR-pivoting-based reconstruction fails to output meaningful results. Figure 10 confirms this, comparing MSEs for stochastically sampled snapshots using DDPM against QR-pivoting and SHAP predictions. Finally, Figure 11 shows a randomly selected snapshot for u' predicted using a SHAP mask with 118 sensors and an IIGDM, demonstrating the accuracy of the SHAP-derived sensor placement.

4. Conclusion

This study introduces a novel framework for reconstructing turbulent flow fields from sparse sensor data using diffusion models and SHAP-based OSP strategies, particularly

suited for urban environments where traditional methods like QR-pivoting fall short due to accessibility issues in high-energy regions. We found that pretrained unconditional diffusion models, like DDRMs and IIGDMs, can generate synthetic turbulent data efficiently, with IIGDMs outperforming DDRMs in accuracy, especially for streamwise velocity, due to superior conditioning. SHAP-based sensor placement not only requires fewer sensors but also surpasses QR-pivoting by offering structured, flexible placement that enhances reconstruction accuracy. Future work will focus on refining diffusion models, adding physics-informed constraints, and extending the approach to 3D domains to improve robustness for real-world applications.

Acknowledgements

The authors thank the CTR program participants for their valuable discussions and insights, as well as Andres Cremades for his input on this work. A.V. and R.V. would like to acknowledge the support from Marie Skłodowska-Curie Actions project MODE-LAIR, funded by the European Union’s HORIZON Research and Innovation Programme through grant agreement no. 101072559. R.V. would like to acknowledge the European Union’s HORIZON Research and Innovation Programme’s, project RefMap, under grant agreement number 101096698.

REFERENCES

- Brunton, S. L., & Kutz, J. N. 2022. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (2nd ed.) Cambridge University Press.
- Cremades, A., Hoyas, S., Deshpande, R., Quintero, P., Lellep, M., Lee, W. J., Monty, J. P., Hutchins, N., Linkmann, M., Marusic, I., & Vinuesa, R. 2024 Identifying regions of importance in wall-bounded turbulence through explainable deep learning. *Nat. Commun.* **15**, 3864.
- Du, P., Parikh, M. H., Fan, X., Liu, X.-Y., & Wang, J.-X. 2024 CoNFILD: conditional neural field latent diffusion model generating spatiotemporal turbulence. *arXiv Preprint arXiv:2403.05940*.
- Gao, H., Han, X., Fan, X., Sun, L., Liu, L.-P., Duan, L., & Wang, J.-X. 2024. Bayesian conditional diffusion models for versatile spatiotemporal turbulence generation. *Comput. Method. Appl. M.* **427**, 117023.
- Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H., & Vinuesa, R. 2021 From coarse wall measurements to turbulent velocity fields through deep learning. *Phys. Fluids* **33**, 075121.
- Ho, J., Jain, A., & Abbeel, P. 2020 Denoising Diffusion Probabilistic Models. In H. Larochelle, M.Ranzato, R. Handsell, M.F. Balcon & H.Lin (eds.), *Advances in Neural Information Processing Systems 33*. NeurIPS.
- Jeyakumar, J. V., Noor, J., Cheng, Y.-H., Garcia, L., & Srivastava, M. 2020 How can I explain this to you? An empirical study of deep neural network explanation methods. In H. Larochelle, M.Ranzato, R. Handsell, M.F. Balcon & H.Lin (eds.), *Advances in Neural Information Processing Systems 33*. NeurIPS.
- Kawar, B., Elad, M., Ermon, S., & Song, J. 2022 Denoising Diffusion Restoration Models. In S. Koyejo, S.Mohammed, A.Agarwal, D.Belgrave, K.Cho & A.Oh (eds.), *Advances in Neural Information Processing Systems 35*. NeurIPS.
- Lundberg, S. M., & Lee, S.-I. 2017 A unified approach to interpreting model predictions.

- In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan & R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*. NeurIPS.
- Martínez-Sánchez, A., López, E., Le Clainche, S., Lozano-Durán, A., Srivastava, A., & Vinuesa, R. 2023. Causality analysis of large-scale structures in the flow around a wall-mounted square cylinder. *J. Fluid Mech.* **967**, A1.
- Shapley, L. S. 1953 A value for n-person games. In H. Kuhn & A. Tucker (eds.), *Contributions to the Theory of Games II* (pp. 307-317). Princeton University Press.
- Shu, D., Li, Z., & Barati Farimani, A. 2023 A physics-informed diffusion model for high-fidelity flow field reconstruction. *J. Comput. Phys.* **478**, 111972.
- Song, J., Vahdat, A., Mardani, M., & Kautz, J. 2022 Pseudoinverse-guided diffusion models for inverse problems. [conference presentation]. The Eleventh International Conference on Learning Representations, Kigali, Rwanda, May 1-5.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. 2021. Score-based generative modeling through stochastic differential equations. [conference presentation]. The Ninth International Conference on Learning Representations (virtual only), May 3-7.
- Torres, P., Le Clainche, S., & Vinuesa, R. 2021 On the experimental, numerical, and data-driven methods to study urban flows. *Energies* **14**, 1310.
- World Health Organization (WHO) 2024 Ambient (outdoor) air pollution, [fact Sheets].
- Zhuang, Y., Cheng, S., & Duraisamy, K. 2024 Spatially-aware diffusion models with cross-attention for global field reconstruction with sparse observations. *arXiv Preprint arXiv:2409.00230*.