CWE-8
June 7-11, 2026
London, ON, Canada

A Generative Modeling Approach for Two-dimensional Spatiotemporal Flow around Elliptical Cross-Sections

Chuanting Liu [a], Genshen Fang [b], Ke Li [c], Yaojun Ge [d]

[a] *State Key Laboratory of Disaster Reduction in Civil Engineering, Tongji University, Shanghai, 200092, China, chuantingliu@tongji.edu.cn*
[b] *State Key Laboratory of Disaster Reduction in Civil Engineering, Tongji University, Shanghai, 200092, China, 2222tjfgs@tongji.edu.cn*
[c] *School of Civil Engineering, Chongqing University, Chongqing 400045, China, keli-bridge@cqu.edu.cn*
[d] *State Key Laboratory of Disaster Reduction in Civil Engineering, Tongji University, Shanghai, 200092, China, yaojunge@tongji.edu.cn*

SUMMARY

Turbulent flow simulation is computationally expensive and time-consuming, limiting the efficiency of aerodynamic performance assessment and wind-environment analysis. Recent advances in deep learning offer a data-driven alternatives to traditional numerical simulations; however, existing methods struggle to model spatiotemporal flows that vary across different aerodynamic shapes. To address this challenge, we propose a generative framework that integrates geometry-conditioned time-averaged flow prediction, latent-space flow representation, and latent sequences generation. An optimal-transport conditional flow-matching (OT-CFM) generates the time-averaged flow from geometric prompts, a β-Variational Autoencoder (β-VAE) encodes flow fields into a compact latent space, and a second OT-CFM produces arbitrary-phase temporal dynamics directly in this latent space. The framework is trained and validated on a dataset of flow around ellipses with diverse aspect ratios and angles of attack, and effectively captures both spatial patterns and temporal fluid dynamics.

Keywords: *Turbulent Flow, Aerodynamic Shape, Generative Model*

1. INTRODUCTION

Turbulent flow around buildings and structures plays a critical role in wind-resistance design and urban wind environment analysis. When boundary conditions change, particularly when aerodynamic shape varies, the flow field must be recomputed by solving the governing partial differential equations (PDEs), which is computationally expensive and time-consuming. By data-driven deep learning (DL) models, turbulent flow simulation can be accelerated. Most existing studies address either geometric variations or temporal flow, but not both. Li et al., (2025) predicted time-averaged flow fields over different shapes. Du et al., (2024) proposed the CoNFiLD model that generates temporal flow sequences in a one-shot approach in latent space. Hu et al., (2025) proposed a diffusion model, relying on phases defined by POD as prompt inputs.

To overcome these limitations, we propose a framework that directly generates spatiotemporal flow series around different aerodynamic shapes. The framework first generates a time-averaged flow field conditioned on the geometry embedding, serving as a coarse representation of the geometry-induced flow pattern. Next, the high-dimensional flow fields are compressed into a compact latent representation to reduce computational cost. The latent temporal sequence is then generated without any predefined phase before being decoded back into full spatiotemporal flow.

2. GENERATIVE MODELING APPROACH FOR SPATIOTEMPORAL FLOW

2.1. Dataset description

We construct a dataset of 2D flows around elliptical cross-sections and randomly split it into training and validation sets with a ratio of 9:1. The parameters include the aspect ratio ($AR \in [1, 10]$) and the angle of attack ($\alpha \in [-10°, 10°]$), while the major axis is fixed at $a = 0.1$ m. Halton sequence is used to quasi-randomly sample the (AR, α) space, ensuring low-discrepancy and extensible. CFD simulations are automated using Python and OpenFOAM, as illustrated in Figure 1. 160 cases are computed using the k-ω SST model at $Re = 2500$ in a uniform flow. 250 snapshots are collected for each case with $\Delta t = 0.04$ s at steady regime. Flow fields (streamwise and cross-stream velocity u, v, and pressure p) are interpolated onto a uniform grid of 64×128. The aerodynamic shape is represented by a signed distance field (SDF) and a binary mask.

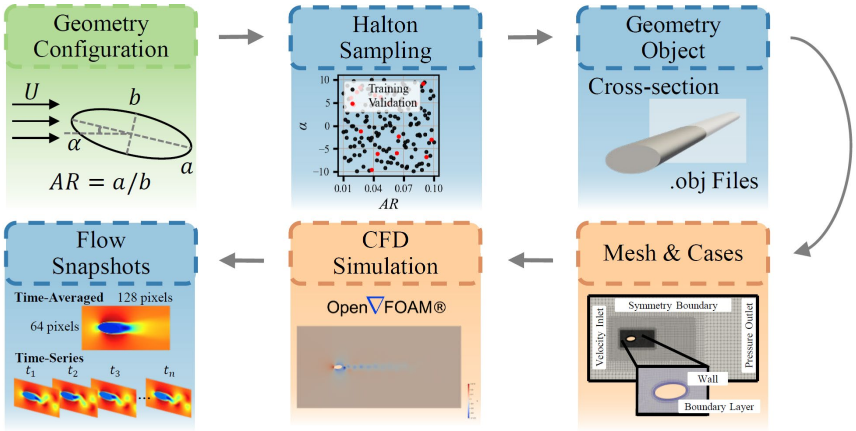

Figure 1: Process of the dataset construction.

2.2. Reduced-order modeling of spatial flow patterns

Spatiotemporal flow fields are high-dimensional, making direct sequence modeling computationally demanding. We reduce the spatial dimension and perform temporal prediction in a compact latent space, and adopt a β-VAE to extract low-dimensional representations of spatial flow patterns of time-series snapshots and time-averaged fields. The encoder maps each flow snapshot to a compact latent vector, and the decoder reconstructs the original field from the latent code. The training loss of the β-VAE is

$$\mathcal{L}(\boldsymbol{x}) = \mathcal{L}_{rec} - \beta \mathcal{L}_{KL} = \frac{1}{N_t} \sum_{i=1}^{N_t} (\boldsymbol{x} - \tilde{\boldsymbol{x}})^2 - \frac{\beta}{2} \sum_{i=1}^{d} \left(1 + \log(\sigma_i^2) - \mu_i^2 - \sigma_i^2\right) \tag{1}$$

where L_{rec} denotes the reconstruction error and L_{KL} is Kullback–Leibler (KL) divergence loss, which measures the difference between the generated probability distribution and the prior Gaussian distribution. β balances reconstruction accuracy and the latent-space disentanglement. The encoder consists of 4 convolutional layers that down sample the input to a bottleneck. A self-attention block is added at the bottleneck to better capture spatial dependencies. The encoder outputs μ_i and $\log \sigma_i^2$ in a $d = 128$ latent space. The decoder mirrors the encoder structure and reconstructs the latent samples back to full-resolution flow snapshots.

2.3. Spatiotemporal flow field generation from aerodynamic shape

The overall pipeline is shown in Figure 2. Our goal is to generate spatiotemporal flow conditioned on a given aerodynamic shape. The process consists of three parts: (i) generating the time-averaged flow from the shape; (ii) mapping between the flow field and its latent representation; (iii) generating a latent temporal sequence conditioned on the latent code of the time-averaged flow. Parts (i) and (iii) use two OT-CFM models, while the β-VAE handles part (ii). These three components are trained independently, avoiding instability caused by joint backpropagation in complex generative models.

OT-CFM is a generative method that transforms a prior distribution, into a target distribution along a learned conditional flow. At each inference step, a UNet predicts the probability flow velocity based on the current state and time. In part (i), the geometry is represented by binary mask and SDF field and used as prompt to generate the time-averaged flow from Gaussian noise. In part (iii), the latent code of the time-averaged flow serves as the condition for generating a sequence of latent codes corresponding to flow snapshots. Since the steady regime demonstrates periodic behavior, we do not impose any phase or time-stamp assumptions. This allows OT-CFM to start from random noise and produce a fixed-length sequence corresponding to any possible flow phase.

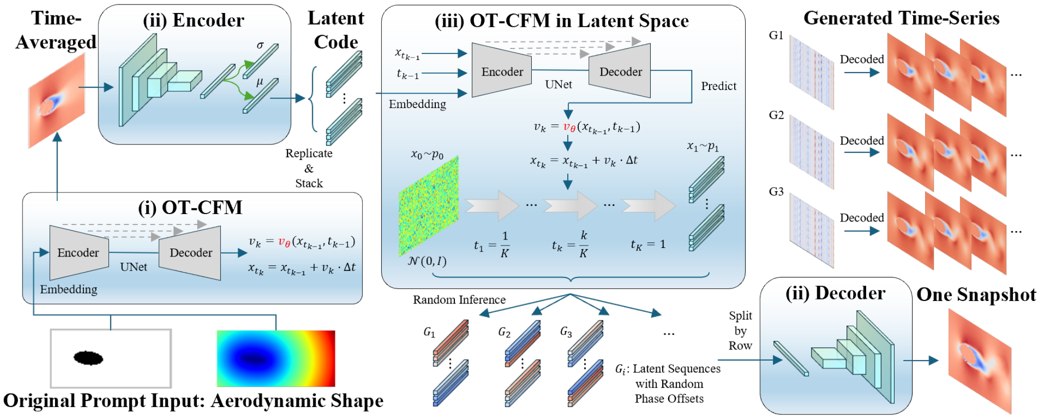

Figure 2: Overview of the pipeline—from aerodynamic shape to time-averaged flow to latent-space temporal flow.

3. RESULTS

3.1. Evaluation of Time-Averaged Flow Generation

The time-averaged flow fields characterize the primary patterns associated with different aerodynamic shapes. For brevity, we present only the generated u-component, while similar methodology applies to v and p. Figure 3 shows two examples. OT-CFM successfully captures key physical structures such as the separation point, wake width, and recirculation-region length, which are crucial for engineering applications. The spatially averaged mean absolute errors relative to CFD results on the training and validation sets are 0.0079 and 0.0160 , respectively; nevertheless, localized discrepancies remain, primarily in the boundary layer and wake region.

3.2. Evaluation of Temporal Flow Sequence Generation

The time-averaged and time-series flow fields are projected into a shared latent space using β-VAE, enabling the generation of temporal flow sequences in a significantly reduced-dimensional representation. Figure 4 presents the velocity time histories at two monitoring points for two different ellipses, each evaluated through two independent inference samples. As shown, while the phases are random, the amplitude and periodicity closely match the CFD results (i.e. ground truth), demonstrating that the model successfully captures the spatiotemporal flow dynamics.

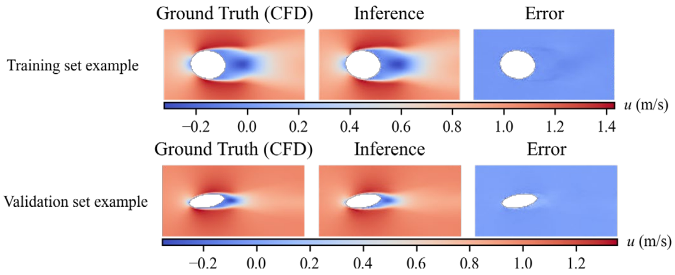

Figure 3: Generated time-averaged flow fields. From left to right: ground truth (CFD), inference and error.

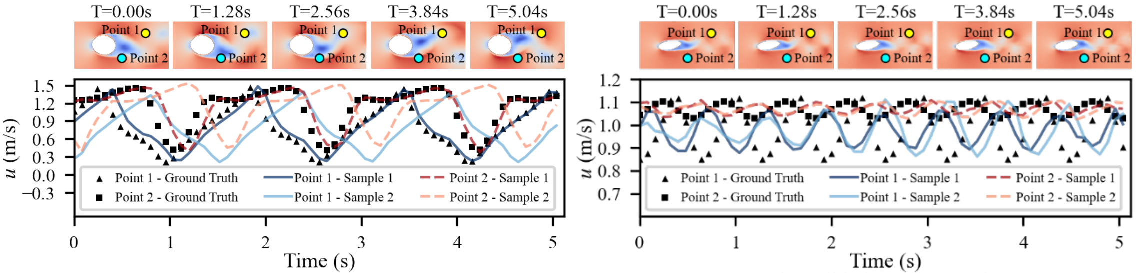

(a) Training set example

(b) Validation set example

Figure 4: Time histories of u at sampled points in the generated spatiotemporal flow.

4. CONCLUSIONS

We propose an “aerodynamic shape–time-averaged flow–latent-space temporal flow” generative framework, which significantly shortens the time required to obtain flow fields while preserving key flow characteristics. Future work will be extended to a broader range of cross-sections to provide an efficient and reliable foundation for large-scale shape exploration, rapid aerodynamic evaluation, and iterative optimization in downstream tasks.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the National Natural Science Foundation of China (52278520, 52578602, 52508578), Young Elite Scientists Sponsorship Program by CAST (2023QNRC001), Natural Science Foundation of Shanghai (25ZR1402494), the Committee of Science and Technology of Shanghai, China (24160712200).

REFERENCES

Li, K., Wang, L., Chen, Z., Zhao, W., Qin, Y., Li, S., 2025. Intelligent prediction method of flow field and aerodynamic characteristics for two-dimensional blunt body combined sections in steady wind. Acta Aerodynamica Sinica. 43(5): 1−12. https://doi.org/10.7638/kqdlxxb-2025.0036 (in Chinese)

Du, P., Parikh, M. H., Fan, X., Liu, X., Wang, J., 2024. Conditional neural field latent diffusion model for generating spatiotemporal turbulence. Nature Communications, 15(1), 10416. https://doi.org/10.1038/s41467-024-54712-1

Hu, J., Lu, Z., Yang, Y., 2025. Generative prediction of flow fields around an obstacle using the diffusion model. Physical Review Fluids, 10(9), 094903. https://doi.org/10.1103/8flg-k6s5