

# Generative reconstruction of spatiotemporal Wall-pressure in turbulent boundary layers via patchwise latent diffusion

Xiantao Fan <sup>1</sup>, Meet Hemant Parikh <sup>1</sup>, Yi Liu <sup>1,2</sup>, Xin-Yang Liu <sup>2</sup>, Junyi Guo <sup>1</sup>,  
Meng Wang <sup>2</sup>, and Jian-Xun Wang <sup>1,2,\*</sup>

<sup>1</sup>*Sibley School of Mechanical and Aerospace Engineering, [Cornell University](#), Ithaca, New York, USA*

<sup>2</sup>*Department of Aerospace and Mechanical Engineering, [University of Notre Dame](#),  
Notre Dame, Indiana, USA*



(Received 21 November 2025; accepted 22 July 2026; published 13 August 2026)

Wall-pressure fluctuations beneath turbulent boundary layers induce structural vibration and acoustic radiation, especially in underwater and aerospace systems. Accurate prediction of their wave number-frequency spectra is critical for effective noise mitigation and design, yet empirical/analytical models rely on simplifying assumptions and miss the full spatiotemporal complexity, while high-fidelity simulations are prohibitive at high Reynolds numbers. Experimental measurements, though accessible, typically provide only pointwise signals and lack the resolution to recover full spatiotemporal fields. We propose a probabilistic generative framework that couples a patchwise (domain-decomposed) conditional neural field with a latent diffusion model to synthesize spatiotemporal wall-pressure fields under varying pressure-gradient conditions. The model conditions on sparse surface-sensor measurements and a low-cost mean-pressure descriptor, supports zero-shot adaptation to new sensor layouts, and produces ensembles with calibrated uncertainty. Validation against reference data shows accurate recovery of instantaneous fields and key statistics.

DOI: [10.1103/ln5n-v7db](#)

## I. INTRODUCTION

Wall-pressure fluctuations beneath turbulent boundary layers (TBLs) are central to many engineering problems. They drive structural vibration and aeroacoustics, contributing to fatigue, noise pollution, and loss of integrity in engineering systems from aircraft and submarines to trains and automobiles [1–3]. Accurate characterization of spatiotemporal wall-pressure fluctuations and their wave number-frequency spectra is essential for providing inputs to structural and acoustic models that predict flow-induced vibration and noise [4,5].

Research over the past several decades has investigated wall-pressure fluctuations through experiments, high-fidelity simulations, and empirical/analytical models [6]. Experiments deliver high-fidelity point signals, but the measured spectra are convolved with the device transfer function, making deconvolution and multiscale recovery difficult [7]; moreover, such measurements are often sparse and do not resolve even the large-scale, spatially coherent structures of TBLs. High-fidelity numerical approaches such as direct numerical simulation (DNS) or wall-resolved large-eddy simulation (WRLES) can resolve wall-pressure dynamics and near-wall turbulence, but their computational cost increases rapidly with Reynolds number, making them impractical for design or optimal control applications [8]. As a faster alternative, analytical and empirical models directly estimate wall-pressure power spectral densities (PSDs), often via the Poisson equation in the wave number-frequency domain or by normalizing canonical spectra [9–11]. However, these

\*Contact author: [jw2837@cornell.edu](mailto:jw2837@cornell.edu)

approaches require accurate second-order turbulence statistics (e.g., Reynolds stress) as input, which are difficult to obtain in practice, and they generalize poorly beyond zero-pressure-gradient (ZPG) flows, especially under adverse/favorable pressure gradients (APG/FPG) [12,13]. Critically, PSD-based models yield spectra rather than instantaneous fields and therefore cannot provide the spatiotemporal fluctuations needed for flow control. As a result, robust reconstruction of wall-pressure fields and their spectra across nonequilibrium boundary layers remains challenging.

To tackle these challenges, machine learning (ML) offers a complementary route. ML can learn expressive nonlinear mappings, enabling flexible modeling of high-dimensional turbulence. One prominent approach develops spectrum-level surrogates: for example, gene expression programming (GEP) has been used to infer analytical formulas for pointwise frequency spectra and adapt scaling to pressure-gradient effects, especially improving low-frequency predictions in APG flows [13–18]. Lightweight neural networks likewise fit spectra in a black-box manner with higher accuracy and training efficiency than classical curve fitting [19–21]. However, these approaches are inherently narrow in scope: they are typically trained for specific conditions and degrade when extrapolated; more fundamentally, by operating on pointwise, one-dimensional spectra, they discard spatial phase and cross-spectral coherence. As a result, they provide limited insight into instantaneous physics and are inadequate for tasks that require full spatiotemporal wall-pressure reconstruction.

In response to these gaps, growing effort has focused on fusing wall-mounted sensor data with ML to reconstruct spatially and temporally resolved wall-pressure fields. Wall sensors are easy to deploy and provide time-resolved, nonintrusive measurements at scale. Leveraging such data, trained models can infer full fields and subsequently yield wave number-frequency spectra and other turbulence statistics via postprocessing. For example, Poulinakis *et al.* [22] combined a long short-term memory (LSTM) network with cubic-spline interpolation to upsample coarse pressure signals, accurately recovering low-frequency content. More broadly, ML has been widely applied to flow reconstruction, especially super-resolution of velocity fields using convolutional neural networks (CNNs) on structured grids [23–25] and graph neural networks on unstructured meshes [26]. Recent advances incorporate skip connections [27], attention mechanisms [28,29], proper orthogonal decomposition (POD)-informed models [30], and mesh-aware CNNs with spatially varying filters [31]; some works incorporate physics-informed losses to promote physical consistency during training [32,33]. Despite this progress, most ML reconstructions are fundamentally deterministic. While such models perform well for relatively simple flows [23], including laminar regimes, two-dimensional vorticity dynamics, and jet flows [34], they struggle to capture the inherent stochasticity of fully turbulent wall-pressure fields. Moreover, reconstructing full fields from sparse sensors is a severely ill-posed inverse problem: multiple plausible fields fit the same data. Deterministic ML yields a single point estimate that cannot represent posterior uncertainty, often blurs coherent structures, and underestimates extreme events. These limitations motivate the development of probabilistic generative frameworks capable of capturing the inherent uncertainty while preserving statistical stationarity in realistic wall-pressure dynamics.

Among generative modeling frameworks, diffusion models have emerged as state-of-the-art for synthesizing high-fidelity flow realizations that reproduce target turbulence statistics. Recent studies demonstrate diffusion for flow generation, reconstruction, and super-resolution [35–43]. A notable exemplar is the conditional neural field-based latent diffusion framework (CoNFILD), which synthesizes spatiotemporal (4D) wall-bounded turbulence while preserving inhomogeneity and anisotropy, and enables synthetic inflows and zero-shot sparse-sensor reconstruction [41,42]. In a Bayesian posterior-sampling view, such trained generative priors can act as foundation models that support training-free reconstruction across diverse sensor arrangements, offering a promising route to flow reconstruction conditioned on sparse wall-pressure sensors regardless of placement or count [44]. However, convection-dominated wall-pressure signals pose additional challenges: the elliptic coupling of the pressure Poisson equation induces long-range, nonlocal dependencies and complex space-time coherence. To our knowledge, no prior work reconstructs full spatiotemporal wall-pressure fields from sparse wall sensors using a generative model.

In this work, we propose a task-specific generative framework for wall-pressure reconstruction in TBLs that combines (i) a domain-decomposed conditional neural field producing compact local latents with high-resolution decoding and (ii) a patched latent diffusion model equipped with classifier-free regime guidance and diffusion posterior sampling for training-free data assimilation. The resulting model adapts at inference to arbitrary sparse sensor layouts without retraining, reconstructs spatiotemporal wall-pressure fields with calibrated uncertainty, and generalizes across a range of APG/FPG conditions. The remainder of the paper is organized as follows. Section II details the proposed framework. Section III presents reconstruction results across different flow conditions to demonstrate the effectiveness of the approach. Section IV analyzes generalization and limitations. Section V concludes and outlines future directions.

## II. METHODOLOGY

### A. Problem formulation

We consider an incompressible turbulent boundary layer (TBL) over a smooth, rigid wall. Using Reynolds decomposition  $u_i = U_i + u'_i$  and  $p = \bar{p} + p'$ , the fluctuating pressure field  $p'(\mathbf{x}, t)$  satisfies the Poisson equation obtained by taking the divergence of the momentum equations [6],

$$-\frac{\partial^2 p'}{\partial x_i^2} = 2\rho \frac{\partial U_i}{\partial x_j} \frac{\partial u'_j}{\partial x_i} + \rho \frac{\partial^2}{\partial x_i \partial x_j} (u'_i u'_j - \overline{u'_i u'_j}), \quad (1)$$

where  $\rho$  is the density;  $x_i$  denotes the Cartesian coordinates;  $U_i$  and  $u'_i$  are the mean and fluctuating velocity components, respectively;  $i, j \in \{1, 2, 3\}$  correspond to the streamwise, wall-normal, and spanwise directions; and repeated indices imply summation. The two right-hand-side terms are the classical rapid (mean-shear-turbulence) and slow (turbulence-turbulence) contributions. Although our target is the wall-pressure fluctuation field, the elliptic character of Eq. (1) implies strong nonlocality: local pressure depends on the global distribution of velocity gradients and Reynolds stresses, both of which also evolve in time.

We define the wall-pressure fluctuation  $p_w(x_1, x_3, t) \equiv p'(x_1, 0, x_3, t)$  on the wall plane. In principle, one could obtain  $p_w$  by solving Eq. (1) (e.g., via Green's functions). In practice, this would require accurate Reynolds stresses and velocity fluctuations throughout the near-wall volume, which are notoriously difficult to measure or numerically resolve at high Reynolds numbers. By contrast, wall pressure can be measured reliably and nonintrusively with flush-mounted sensors. However, such measurements are often sparse (limited mounting locations) and noisy (finite sampling rate and sensor noise), so inferring the full, densely resolved and globally coupled field from these measurements is underdetermined.

As introduced in Sec. I, most ML approaches for wall pressure target direct prediction of spectral quantities [13–15, 19, 45]. In contrast, we aim to reconstruct the full spatiotemporal wall-pressure field from sparse sensors with calibrated uncertainty. Because the inverse map is many-to-one and the dynamics are stochastic, a generative formulation is appropriate. Following the representation-dynamics factorization of CoNFILD [41], our framework (Fig. 1) first uses a domain-decomposed conditional neural field (D-CNF) [46] to encode each instantaneous wall-pressure field into a compact latent  $z_t$ , with a decoder that maps latents back to the physical domain. In the latent space, a classifier-free patched diffusion model learns the stochastic temporal evolution. Patching induces locality on the wall plane and scales to long spans and high resolutions. Two conditioning signals steer generation toward full-field reconstructions. The first is sensor consistency enforced via diffusion posterior sampling (DPS) during inference only: given sparse measurements, the reverse diffusion is augmented by the sensor likelihood gradient backpropagated through the D-CNF decoder, enabling agreement with arbitrary sensor layouts without retraining. The second is flow-regime guidance provided by a low-cost descriptor  $c$  (e.g., a mean wall-pressure profile obtainable from a RANS simulation). Using classifier-free guidance,  $c$  is applied during both training and inference, which promotes zero-shot generalization across different adverse/favorable pressure-gradient (APG/FPG)




FIG. 1. Schematic of the generative framework for reconstructing wall pressure. Top: the domain-decomposed conditional neural field (D-CNF) encodes the physical domain into latent space and decodes latents back to fields. Bottom: the latent diffusion model with two conditioning mechanisms: (I) sensor consistency via DPS (inference only) and (II) flow-regime guidance via a mean-profile descriptor used in both training and inference. Solid arrows denote training; dashed arrows denote inference.

flow regimes. With DPS and regime guidance, the diffusion model generates latent trajectories that are decoded to produce measurement-aligned spatiotemporal reconstructions of instantaneous  $p_w(x_1, x_3, t)$ , while preserving physically relevant statistics (e.g., wave number-frequency spectra and coherence) for downstream applications.

### B. Domain-decomposed conditional neural field for latent representation

The first component is a domain-decomposed conditional neural field (D-CNF) that maps each instantaneous wall-pressure field to a compact latent code and decodes latents back to the physical domain. We adopt an autoencoder design with a fully projected hypernetwork, following the CoNFILd line [42] and recent nonoverlapping CNF tiling study [46]. In an autoencoder, no amortized encoder is trained; instead, each training sample has a learnable latent code, while at inference, latents are supplied by the generative model.

Let the wall plane be  $\Omega_w \subset \mathbb{R}^2$  with in-plane coordinates  $\mathbf{x} = (x_1, x_3)$ . We partition  $\Omega_w$  into  $N_p$  nonoverlapping tiles  $\{\Omega_i\}_{i=1}^{N_p}$  whose interiors are disjoint and whose union (up to boundaries) equals  $\Omega_w$ . For each tile, an affine map  $\mathcal{T}_i : \Omega_i \rightarrow [-1, 1]^2$  normalizes coordinates to promote parameter sharing. We denote by  $t$  the discrete time index and by  $c \in \mathbb{R}^{d_c}$  a flow-regime descriptor (e.g., Clauser parameter  $\beta$ , or a mean wall-pressure/profile summary from a low-cost RANS computation). The D-CNF predicts per-tile fields  $\hat{p}_{w,i}(\mathbf{x}, t, c)$  and assembles the global reconstruction by tiling:  $\hat{p}_w(\mathbf{x}, t, c) = \hat{p}_{w,i}(\mathbf{x}, t, c)$ ,  $\mathbf{x} \in \Omega_i$ ,  $i = 1, \dots, N_p$ , so no explicit blending is required, and locality is preserved. Each tile is modeled by a sinusoidal representation network (SIREN)-based implicit field  $\mathcal{F}_\theta$  [47] acting on local coordinates  $\tilde{\mathbf{x}} = \mathcal{T}_i(\mathbf{x}) \in [-1, 1]^2$ ,

$$\hat{p}_{w,i}(\mathbf{x}, t, c) = \mathcal{F}_{\theta_{i,t,c}}(\tilde{\mathbf{x}}), \quad \theta_{i,t,c} = \zeta + \mathcal{H}_\gamma(\mathbf{s}_{i,t,c}, t, c), \quad (2)$$

where  $\mathbf{s}_{i,t,c} \in \mathbb{R}^{d_z}$  is the learnable latent for tile  $i$  at time  $t$  and descriptor  $c$ ;  $\zeta$  collects the shared base SIREN parameters; and  $\mathcal{H}_\gamma$  is a fully projected hypernetwork (parameters  $\gamma$ ) that outputs layer-wise weight/bias offsets applied to all SIREN layers. This projection is more expressive than FiLM-style modulation and has proved effective for convection-dominated wall-pressure structure [46].

Let  $\mathcal{D} = \{(t_j, c_j, p_w(\cdot, t_j, c_j))\}_{j=1}^{N_s}$  be the training set of snapshots with continuous descriptors. For each sample  $j$  and tile  $i$ , we associate a latent  $\mathbf{s}_{i,j}$  (equivalently  $\mathbf{s}_{i,t_j,c_j}$ ). Denote the restriction of  $p_w$  to  $\Omega_i$  by  $p_{w,i}(\mathbf{x}; t_j, c_j)$ . We jointly optimize latents and network parameters by minimizing an  $L_2$



reconstruction loss over points sampled in each tile,

$$(\mathbf{s}^*, \zeta^*, \gamma^*) = \arg \min_{\mathbf{s}, \zeta, \gamma} \sum_{j=1}^{N_k} \sum_{i=1}^{N_p} \|p_{w,i}(\mathbf{x}; t_j, c_j) - \mathcal{F}_{\zeta+\mathcal{H}_\gamma(\mathbf{z}_{i,j}, t_j, c_j)}(\mathcal{T}_i(\mathbf{x}))\|_{L_2(\Omega_i)}^2. \quad (3)$$

A small  $\ell_2$  regularizer on  $\mathbf{s}_{i,j}$  may be used for stability, though we did not find it necessary. The SIREN backbone uses the standard input-frequency scale  $\omega_0$  controlled by  $s_0$  and initialization  $W \sim \mathcal{U}(-r/\sqrt{n}, r/\sqrt{n})$ , where  $\mathcal{U}(\cdot)$  represents a uniform distribution with specified bounds [47]. Typical choices  $s_0 = 30$  and  $r = \sqrt{6}$  were robust in our experiments. Full architectural details are provided in Appendix A. The D-CNF therefore yields per-tile latents that (i) capture localized, convection-dominated features, (ii) provide exact gradients  $\partial \hat{p} / \partial \mathbf{s}$  for sensor-consistency in DPS, and (iii) decode to high-resolution fields via the nonoverlapping tiling.

### C. Conditional latent diffusion model with classifier-free guidance

At physical time index  $t$ , the D-CNF (Sec. II B) yields per-tile latents  $\{\mathbf{s}_{i,t,c} \in \mathbb{R}^{d_z}\}_{i=1}^{N_p}$  for a given flow regime descriptor  $c \in \mathbb{R}^{d_c}$ . For temporal modeling, we form a short window of length  $T_w$  and, for each tile  $i$ , stack the latents into a spatiotemporal tensor  $\mathbf{z}_0^{(i)} \in \mathbb{R}^{T_w \times d_z}$ . Diffusion is performed on  $\mathbf{z}_\tau^{(i)}$  with tiles processed in parallel as the batch dimension. For notational brevity below, we drop the tile index and write  $\mathbf{z}_\tau \in \mathbb{R}^{T_w \times d_z}$ .

With a variance schedule  $\{\beta_\tau\}_{\tau=1}^S$ ,  $\alpha_\tau = 1 - \beta_\tau$ , and  $\bar{\alpha}_\tau = \prod_{u=1}^\tau \alpha_u$ , the forward (noising) process is defined by

$$q(\mathbf{z}_\tau | \mathbf{z}_{\tau-1}) = \mathcal{N}(\sqrt{\alpha_\tau} \mathbf{z}_{\tau-1}, \beta_\tau \mathbf{I}), \quad q(\mathbf{z}_\tau | \mathbf{z}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_\tau} \mathbf{z}_0, (1 - \bar{\alpha}_\tau) \mathbf{I}), \quad (4)$$

where  $\tau \in \{1, \dots, S\}$  indexes diffusion steps (distinct from physical time  $t$ ).

The reverse process (generation) is achieved by learning a noise predictor  $\epsilon_\theta(\mathbf{z}_\tau, \tau, \tilde{c})$  with parameters  $\theta$ , where  $\tilde{c}$  is either the descriptor  $c$  or a null token  $\emptyset$  (for classifier-free learning). The reverse transition is parametrized as

$$p_\theta(\mathbf{z}_{\tau-1} | \mathbf{z}_\tau, c) = \mathcal{N}(\mu_\theta(\mathbf{z}_\tau, \tau, \tilde{c}), \tilde{\beta}_\tau \mathbf{I}), \quad \mu_\theta(\mathbf{z}_\tau, \tau, \tilde{c}) = \frac{1}{\sqrt{\bar{\alpha}_\tau}} \left( \mathbf{z}_\tau - \frac{\beta_\tau}{\sqrt{1 - \bar{\alpha}_\tau}} \epsilon_\theta(\mathbf{z}_\tau, \tau, \tilde{c}) \right), \quad (5)$$

with  $\tilde{\beta}_\tau$  following the usual denoising diffusion probabilistic model (DDPM) choice. The training is performed using the simplified noise-matching loss,

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{\mathbf{z}_0, \tau \sim \mathcal{U}[1, S], \epsilon \sim \mathcal{N}(0, \mathbf{I}), \tilde{c}} [\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_\tau} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_\tau} \epsilon, \tau, \tilde{c})\|_2^2]. \quad (6)$$

To generalize across a range of flow conditions with varying pressure gradients, we adopt a classifier-free guidance (CFG) strategy to handle the flow descriptor  $c$  without training a separate noisy-latent classifier and to avoid the diversity collapse often induced by classifier guidance [48]. CFG trains a single noise predictor on both conditional and unconditional inputs by randomly dropping  $c$  with probability  $p_{\text{uncond}}$ , then steers sampling with a tunable guidance strength  $w$  while preserving stochastic variability. During training, we set  $\tilde{c} = \emptyset$  with probability  $p_{\text{uncond}}$ , and at inference we combine scores as

$$\tilde{\epsilon}_\theta(\mathbf{z}_\tau, \tau, c) = (1 + w) \epsilon_\theta(\mathbf{z}_\tau, \tau, c) - w \epsilon_\theta(\mathbf{z}_\tau, \tau, \emptyset), \quad (7)$$

where  $w$  controls guidance ( $w = -1$  unconditional,  $w = 0$  conditional only,  $w > 0$  extrapolated; we use  $w = 1$  for diverse pressure-gradient scenarios unless noted). The flow regime descriptor  $c$  is embedded by a small multilayer perceptron (MLP) and fused with the timestep embedding at all U-Net blocks; we use  $p_{\text{uncond}} = 0.2$ . Specifically, a modernized U-Net is used for  $\epsilon_\theta(\cdot)$  that operates on tensors of shape  $T_w \times d_z$  (tiles in batch). The network employs residual convolutional blocks, group normalization, and multihead self-attention over time to capture long-range spatiotemporal

couplings in latent space. Diffusion step  $\tau$  is encoded via sinusoidal embeddings; the descriptor embedding for  $c$  is added to these embeddings at each block. More details of the network architecture are provided in Fig. 26 in Appendix B.

#### D. Generative reconstruction from sparse wall measurements without retraining

We aim to reconstruct the spatiotemporal wall-pressure field from sparse wall sensors under arbitrary layouts, without retraining. As sensor number, placement, and sampling rate vary across cases, to let a single pretrained model adapt zero-shot to any configuration, we assimilate sensor information exclusively at inference via DPS [39,41,43]. In DPS, the D-CNF and diffusion parameters are frozen; sensors influence sampling only through a likelihood gradient computed on the decoded field through the state-to-observable operator rather than by updating weights.

Concretely, given the flow descriptor  $c$  and sparse wall records  $\mathbf{y}$  over a temporal window of length  $T_w$ , we sample the latent “image”  $\mathbf{z}_0 \in \mathbb{R}^{T_w \times d_z}$  using the DPS-modified reverse process. Classifier-free guidance provides the conditional prior score  $s_\theta(\mathbf{z}_\tau, \tau, c; \theta^*)$ ,

$$s_\theta(\mathbf{z}_\tau, \tau, c; \theta^*) = -\frac{\epsilon_\theta(\mathbf{z}_\tau, \tau, c; \theta^*)}{\sqrt{(1 - \bar{\alpha}_\tau)}}. \quad (8)$$

The sensor measurements  $\mathbf{y}$  are assimilated by modifying this to form a posterior score,

$$s_{\text{post}}(\mathbf{z}_\tau, \tau, c; \mathbf{y}, \theta^*) = s_\theta(\mathbf{z}_\tau, \tau, c; \theta^*) + \nabla_{\mathbf{z}_\tau} \log p(\mathbf{y}|\mathbf{z}_\tau), \quad (9)$$

where the intractable marginal  $p(\mathbf{y}|\mathbf{z}_\tau)$  is approximated by evaluating a pseudolikelihood at a clean-latent estimate  $\hat{\mathbf{z}}_0(\mathbf{z}_\tau)$  [49]. Using Tweedie’s formula in the  $\epsilon$ -prediction parametrization [50,51],

$$\hat{\mathbf{z}}_0(\mathbf{z}_\tau) \approx \frac{1}{\sqrt{\bar{\alpha}_\tau}}(\mathbf{z}_\tau + (1 - \bar{\alpha}_\tau)s_\theta(\mathbf{z}_\tau, \tau, c; \theta^*)). \quad (10)$$

Assuming Gaussian measurement noise with scale  $\sigma_c(\tau)$ , the likelihood score can be approximated as

$$\nabla_{\mathbf{z}_\tau} \log p(\mathbf{y}|\mathbf{z}_\tau) \simeq \nabla_{\mathbf{z}_\tau} \log p(\mathbf{y}|\hat{\mathbf{z}}_0) = -\frac{1}{\sigma_c(\tau)^2} \nabla_{\mathbf{z}_\tau} \frac{1}{2} \|\mathcal{M}(\mathcal{G}'(\hat{\mathbf{z}}_0)) - \mathbf{y}\|_2^2, \quad (11)$$

with the gradient backpropagated through  $\mathbf{z}_\tau \rightarrow \hat{\mathbf{z}}_0(\mathbf{z}_\tau) \xrightarrow{\mathcal{G}'}$  field  $\xrightarrow{\mathcal{M}}$  sensors. Following [49], the DPS guidance step is further normalized by the current measurement residual. Therefore, using  $\sigma_c(\tau) = 1$  corresponds to adopting a unit likelihood scale in the normalized sensor space rather than imposing an unnormalized correction magnitude; an annealed schedule can emphasize sensors late in the reverse process, where  $\mathcal{G}'$  denotes the D-CNF decoder that maps the latent variable to the full pressure field, and  $\mathcal{M}$  is a spatial operator that extracts the sensor locations from the reconstructed field. Because  $\mathbf{y}$  and  $\mathcal{M}$  enter only through the likelihood, the same pretrained model adapts zero-shot to different sensor counts, layouts, and sampling rates, and no retraining or architectural changes are required. Each posterior sample  $\mathbf{z}_0$  is finally decoded by  $\mathcal{G}'$  to produce a coherent reconstruction; autoregressive temporal DPS yields long sequences [41], and ensembles provide posterior means and uncertainty.

### III. RESULTS

#### A. Case setup and training data generation

We construct a database of TBLs under varying pressure gradients using direct numerical simulations (DNS). The three-dimensional, incompressible Navier-Stokes equations are solved with a well-established finite-volume code developed at Stanford University [52]. The numerical setup follows Abe *et al.* [53,54], where an artificial boundary condition is imposed at the upper boundary (blowing-suction) to mimic pressure-gradient effects, while a no-slip condition is applied at the wall. A two-dimensional separation bubble is induced in the  $x$ - $y$  plane; periodic boundary conditions

TABLE I. Numerical settings and parameters for the DNS dataset, where  $V\cdot$  denotes the fraction of the reference amplitude.

| Case                                 | V0                                                      | V20    | V40   | V60    | V80   | V90     | V100   |
|--------------------------------------|---------------------------------------------------------|--------|-------|--------|-------|---------|--------|
| $V_{\max}$                           | 0                                                       | 0.0665 | 0.133 | 0.1995 | 0.266 | 0.29925 | 0.3325 |
| $\beta_{x_1/\theta_{\text{in}}=100}$ | 0                                                       | 2.14   | 6.92  | 22.31  | 62.31 | 90.19   | 302.55 |
| Grid                                 | $N_x \times N_y \times N_z = 512 \times 320 \times 512$ |        |       |        |       |         |        |
| Domain                               | $L_x \times L_y \times L_z = 400 \times 80 \times 160$  |        |       |        |       |         |        |
| $\Delta T$                           | 1.6                                                     |        |       |        |       |         |        |
| $Re_{\theta_{\text{in}}}$            | 300                                                     |        |       |        |       |         |        |
| $T_{\text{flow}}^{\text{train}}/N$   | 0 – 4/1000                                              |        |       |        |       |         |        |
| $T_{\text{flow}}^{\text{test}}/N$    | 8 – 12/1000                                             |        |       |        |       |         |        |

are enforced in the spanwise ( $z$ ) direction. In the streamwise direction, inflow and convective outflow boundary conditions are employed. Turbulent inflow is supplied via the rescale-and-recycle procedure of Lund *et al.* [55]. Unless noted, all quantities are normalized by outer variables (free-stream velocity  $U_\infty$  and inlet momentum thickness  $\theta_{\text{in}}$ ). The Reynolds number based on the inlet momentum thickness is  $Re_{\theta_{\text{in}}} = 300$ .

To enrich coverage beyond the three cases in [53], we vary the magnitude of the upper-boundary wall-normal blowing/suction to generate a wide range of TBL cases, including attached and separated flows, spanning zero-, adverse-, and favorable-pressure-gradient (ZPG, APG, FPG) regimes. The wall-normal velocity at the top boundary is prescribed as

$$V_{\text{top}} = V_{\max} \sqrt{2} \left( \frac{x_c - x_1}{\sigma} \right) \exp \left[ \psi - \left( \frac{x_c - x_1}{\sigma} \right)^2 \right], \quad (12)$$

with  $x_c = 195.3125$ ,  $\sigma = 89$ , and  $\psi = 0.95$ , and where  $V_{\max}$  controls the pressure-gradient strength (see Table I). At the top boundary, the streamwise velocity  $u$  is specified to maintain zero spanwise vorticity; a Neumann condition is imposed on the spanwise velocity  $w$ . We characterize pressure-gradient intensity via the Clauser parameter [54],

$$\beta = \frac{\delta^*}{\tau_w} \frac{d\bar{p}}{dx_1}, \quad (13)$$

where  $\delta^*$  is the displacement thickness and  $\tau_w$  is the wall shear stress. For all cases, the largest pressure gradient occurs near  $x_1/\theta_{\text{in}} = 100$ ; therefore, we report the corresponding  $\beta$  values at this location in Table I. Case labels  $V\cdot$  denote the fraction of the reference amplitude (e.g.,  $V20 \equiv 0.2 V_{\max}$ ;  $V100 \equiv V_{\max}$ ). The suite spans ZPG, FPG, attached APG up to incipient separation, and strongly separated APG (V80-V100;  $\beta > 39$ ) [56]. Numerical settings and  $\beta$  values are also summarized in Table I. We validate the DNS by comparing (i) inlet ZPG mean and fluctuation profiles against the reference DNS of Spalart [57], and (ii) skin friction  $C_f$  and mean pressure coefficient  $C_p$  under representative APG/FPG (V90) against Abe [53]. As shown in Fig. 2, the present results agree closely with the references: the ZPG inner-outer profiles collapse onto the benchmark data, and the APG/FPG case reproduces the separation/reattachment behavior in  $C_f$  and the streamwise variation of  $C_p$ . This agreement supports the fidelity of our simulations and their use for training and evaluation. The instantaneous streamwise velocity fields for all cases are presented in Fig. 27 to illustrate the flow dynamics.

Wall-pressure snapshots on the wall plane have resolution  $N_x \times N_z = 512 \times 512$ . Data are sampled every  $\Delta T = 1.6\theta_{\text{in}}/U_\infty$  (roughly ten numerical time steps). Training covers four flow-through times ( $T_{\text{flow}}$ ); testing uses disjoint windows eight to 12 flow-through times downstream (Table I). Only cases V0, V40, and V80 are included in training; all other cases are held out for testing to assess cross-regime generalization. To ensure statistical independence between training and testing




FIG. 2. Dataset validation: (a) Mean velocity and (b) velocity fluctuations at the inlet, compared with DNS data from [57]; (c) skin friction and (d) mean pressure coefficient for a turbulent boundary layer under strong APG/FPG ( $V90$ ,  $\beta = 90.19$ ), compared with DNS results from [53]. Reference data are shown as scatter points, while our results are plotted as lines.

data, all test samples are extracted from time intervals beyond the initial four flow-through periods. Following Sec. II B, the wall plane is partitioned into  $16 \times 16$  nonoverlapping tiles (256 total), each of size  $32 \times 32$ , and encoded by the D-CNF into a latent vector of length  $d_z = 64$  per tile/time. For diffusion modeling, we form temporal windows of length  $T_w$  and stack per-tile latents into spatiotemporal “latent images” of size  $T_w \times d_z$  (tiles in the batch dimension). The resulting latent dataset  $\{\mathbf{z}_0\}$  is used to train the diffusion model.

To test the reconstruction performance of the trained model, we consider two scenarios. *Scenario I (single-regime testing)*: the flow condition (e.g.,  $Re$  and pressure-gradient setting) is fixed, but the test sequences are new spatiotemporal realizations drawn from disjoint time windows and/or new initial conditions; inference uses only sparse sensors (Condition I) via DPS, with no parameter conditioning. *Scenario II (cross-regime generalization)*: the model reconstructs fields at previously unseen pressure-gradient strengths (ZPG/APG/FPG). Training includes Condition II via classifier-free guidance with a continuous descriptor (e.g., mean  $C_p$ ), and testing uses both sparse sensors (Condition I) and the descriptor (Condition II) to enable zero-shot transfer across regimes.

### B. Generative wall-pressure reconstruction (single-regime testing)

We first evaluate reconstruction in a single-regime testing setting: the flow condition is fixed, but test sequences are new spatiotemporal realizations drawn from new initial conditions. Three representative regimes are considered, including ZPG (V0), moderate APG/FPG (V40), and strong APG/FPG with separation (V80). For each flow regime, a model is trained once and tested on new




FIG. 3. Reconstructed wall-pressure fluctuations for Case V0 (ZPG,  $\beta = 0$ ), conditioned on sensor arrays at unseen time steps.

realizations. During inference, only Condition I is used: wall-pressure measurements from a  $16 \times 16$  sensor array collocated with the D-CNF tiles are assimilated via DPS.

Figure 3 shows instantaneous reconstructions for V0 over one flow-through time. Even without an imposed pressure gradient,  $p_w$  exhibits rich small-scale, convecting structure; the reconstructions closely track the DNS both frame-by-frame and across the window. As the APG/FPG strength increases to V40 ( $\beta = 6.92$ ; Fig. 4), larger coherent footprints emerge downstream, marking the onset toward separation; these features are faithfully reproduced by the model. Under strong APG/FPG (V80,  $\beta = 62.31$ ; Fig. 5), the DNS shows wall-pressure signatures concentrated near separation and reattachment; the reconstructions recover these large-scale footprints and their temporal evolution while maintaining the embedded small-scale fluctuations. Across all three regimes, DPS-conditioned samples remain spatiotemporally consistent with DNS, indicating that




FIG. 4. Reconstructed wall-pressure fluctuations for Case V40 (APG/FPG,  $\beta = 6.92$ ), conditioned on sensor arrays at unseen time steps.






FIG. 5. Reconstructed wall-pressure fluctuations for Case V80 (APG/FPG,  $\beta = 62.31$ ), conditioned on sensor arrays at unseen time steps.

the latent diffusion prior and D-CNF decoder together can accurately reconstruct the instantaneous wall-pressure fluctuation fields.

At unobserved (i.e., unsensed) locations, the reconstructions are essentially indistinguishable from DNS in both amplitude and phase; the three independent posterior samples lie tightly on the DNS trace with only a very thin scatter band (Fig. 6), indicating highly accurate recovery away from




FIG. 6. Time series of wall-pressure fluctuations at three unobserved locations (no sensors) for Case V80 (APG/FPG,  $\beta = 62.31$ ), including three different sample realizations.






FIG. 7. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for Case V80 (APG/FPG,  $\beta = 62.31$ ): (a) station locations; (b) PDF near the separation point  $x_1/\theta_{\text{in}} = 100$ ; (c) PDF near the bubble center  $x_1/\theta_{\text{in}} = 200$ ; (d) PDF near the reattachment point  $x_1/\theta_{\text{in}} = 300$ .

sensors. Although the  $16 \times 16$  array imposes strong conditioning and can reduce sample diversity, the generator still retains stochastic flexibility: small but physically consistent variations persist across posterior draws, reflecting calibrated uncertainty rather than model bias. The same behavior is observed for V0 and V40 (Figs. 28 and 30), confirming that within a fixed flow regime the method produces DNS-quality reconstructions at unsensed points while still expressing modest, plausible variability.

We next assess statistical fidelity at three representative streamwise stations for V80: the separation region ( $x_1/\theta_{\text{in}} = 100$ ), the center of the recirculation bubble ( $x_1/\theta_{\text{in}} = 200$ ), and the reattachment region ( $x_1/\theta_{\text{in}} = 300$ ). Figure 7 compares the probability density functions (PDFs) of wall-pressure fluctuations at these locations. The reconstructed distributions closely match DNS in mean and variance and capture the tail behavior across all stations. Small discrepancies are limited to a mild underestimation of the most negative tail events near separation, while the bubble center and reattachment stations are nearly indistinguishable from DNS. Consistent agreement is observed for V0 and V40 (Figs. 29 and 31). To further examine the zero-shot capability of the proposed framework, we also consider two additional sensor layouts for Case V80, including a shifted layout and a coarser layout. As shown in Fig. 32, when the sensor arrangement is shifted by four grid points in both the  $x_1$  and  $x_3$  directions, the reconstructed field remains in good agreement with the ground truth, which provides clear evidence of the zero-shot capability. When the number of sensors is reduced to  $12 \times 12$ , such that not every patch contains a sensor although the layout is still relatively dense, the reconstruction accuracy decreases. While the overall flow pattern is still captured reasonably well, noticeable degradation appears in the turbulence statistics, particularly in the  $C_{p,\text{rms}}$  and frequency spectra.




FIG. 8. Turbulence statistics for all training ZPG and APG/FPG regimes on unseen time windows: (a) mean pressure coefficient; (b) root-mean-square of wall-pressure fluctuations; (c) wall-pressure frequency spectra at three representative streamwise locations.

A more comprehensive statistical comparison across the three training pressure-gradient conditions is shown in Fig. 8. The mean wall-pressure coefficient  $C_p$  profile is reproduced with high fidelity [panel (a)], including the growth and downstream decay of the APG-induced hump in V40 and V80. The root-mean-square (rms) pressure fluctuations  $C_{p,rms}$  are also captured reasonably well [panel (b)]: the model recovers the single peak associated with the incipient-separation region in V40 and the two distinct peaks near separation and reattachment in V80. Frequency spectra at three streamwise stations [panel (c)] agree closely with DNS, preserving both the low-frequency energy associated with large scales and the correct high-frequency roll-off. Small discrepancies appear near the bubble center ( $x_1/\theta_{in} = 200$ ) in V80, where the flow is most sensitive to local variations, leading to a mild mismatch around intermediate frequencies; elsewhere the agreement is excellent.

In convection-dominated TBLs, coherent structures are organized by a local convection velocity  $U_c$ : at a fixed point the wall pressure signal decorrelates rapidly in time, whereas along a trajectory moving with speed  $U_c$  coherence persists. This path corresponds to the ridge of maximum streamwise space-time correlation as discussed by Na and Moin [58]. To assess whether the reconstructions preserve this organization, we compute the streamwise space-time correlation coefficient  $R_{pp}(\Delta x_1, \Delta t)$  at several stations, defined as

$$R_{pp}(\Delta x_1, \Delta t; x) = \frac{\langle p_w(x_1, x_3, t) p_w(x_1 + \Delta x_1, x_3, t + \Delta t) \rangle}{\sqrt{\langle p_w^2(x_1, x_3, t) \rangle} \sqrt{\langle p_w^2(x_1 + \Delta x_1, x_3, t) \rangle}}, \quad (14)$$

where  $\langle \cdot \rangle$  denotes time and spanwise averaging. The convection speed is estimated from the ridge slope,  $U_c \simeq \Delta x_1 / \Delta t$ , along the locus of maximal correlation. Figure 9 (station  $x_1/\theta_{in} = 100$ ) shows that the reconstructions reproduce the DNS ridge angle, anisotropy, and decorrelation rates in both




FIG. 9. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 100$  for training ZPG and APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .

$\Delta x_1$  and  $\Delta t$ . As pressure-gradient strength increases ( $V0 \rightarrow V40 \rightarrow V80$ ), the ridge flattens, indicating a decrease in  $U_c$ ; the reconstructed fields match this trend quantitatively.

At the bubble center  $x_1/\theta_{in} = 200$ , the DNS exhibits a markedly reduced  $U_c$  and longer temporal coherence for V80 due to the recirculation region. Our model recovers both the reduced ridge slope and the elongated time correlation, as shown in Fig. 33. Downstream at  $x_1/\theta_{in} = 300$  (Fig. 34), where the flow reattaches to the wall, the convection velocities for all cases become nearly identical, indicating a recovery to a more uniform flow regime, again captured by the reconstruction. These observations confirm that the proposed model not only reconstructs instantaneous flow fields but also faithfully preserves the temporal/spatial coherence properties inherent to convection-dominated wall-pressure dynamics.

### C. Generative wall-pressure reconstruction under different pressure gradients (cross-regime testing)

We now test generalization to unseen pressure-gradient conditions with a single unified model. Unlike Sec. III B, where a separate model was trained per regime, here we jointly train on V0, V40, and V80 while conditioning on a continuous flow descriptor (Condition II): the mean wall-pressure profile  $C_p$  embedded via CFG. During inference, no retraining is performed: we supply the appropriate descriptor  $c$  and assimilate sparse pressure sensor data (Condition I) via DPS, enabling zero-shot reconstruction to new APG/FPG strengths. We evaluate on two interpolated cases (V20, V60) and two extrapolated cases (V90, V100), none used in training.

For the interpolated cases V20 and V60, the unified model reconstructs the instantaneous wall-pressure fields with high fidelity. As shown in Figs. 10 and 11, the instantaneous wall pressure fields are accurately reconstructed, with small-scale coherent structures effectively captured. In Case V60, where the flow remains attached to the wall, larger structures that appear downstream are also well captured. Notably, for the extrapolated strong APG/FPG cases V90 and V100, the model successfully captures multiscale organization of separation and reattachment footprints, as




FIG. 10. Reconstructed wall-pressure fluctuations for interpolated Case V20 (APG/FPG,  $\beta = 2.14$ ), conditioned on mean profiles and sensor arrays.

illustrated in Figs. 12 and 39; mild blurring appears near reattachment in V100, but the spatial placement and temporal evolution of large-scale structures remain accurate.

The generative capability of the model is further supported by the wall-pressure time series sampled at unconditioned spatial locations. As shown in Figs. 13, 35, 38, and 41, the variations across different realizations reflect the model's ability to generate diverse outputs that remain consistent with the physical characteristics of the flow.

The reconstructed fields reproduce wall-pressure statistics with high fidelity. For V60, the PDFs of wall-pressure fluctuations at three stations (Fig. 14) closely match DNS in mean, variance, and tails; similar agreement is observed for V20, V90, and V100 (Figs. 36, 37, and 40). Although the mean profile  $C_p$  is provided at inference, it is incorporated only via unconditional sampling with probability  $p_{\text{uncond}} = 0.2$  rather than enforced as a hard constraint. Accordingly, we recompute  $C_p$  from the generated fields [Fig. 15(a)]: predictions are accurate for the interpolated cases (V20,




FIG. 11. Reconstructed wall-pressure fluctuations for interpolated Case V60 (APG/FPG,  $\beta = 22.31$ ), conditioned on mean profiles and sensor arrays.




FIG. 12. Reconstructed wall-pressure fluctuations for extrapolated Case V90 (APG/FPG,  $\beta = 90.19$ ), conditioned on mean profiles and sensor arrays.

V60) and remain so for the slightly extrapolated V90; a mild overprediction appears near the center of the recirculation bubble for the strongly extrapolated V100. Second-order statistics are also well captured: the peaks of root-mean-square wall-pressure fluctuations  $C_{p,\text{rms}}$  near separation and reattachment are generally reproduced across cases [Fig. 15(b)]. For V100, magnitudes are slightly underestimated, which is consistent with its being far outside the training regime. In




FIG. 13. Time series of wall-pressure fluctuations at three unconditioned sensor locations for interpolated Case V60 (APG/FPG,  $\beta = 22.31$ ), including three different sample realizations.






FIG. 14. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for interpolated Case V60 (APG/FPG,  $\beta = 22.31$ ): (a) station locations; (b) PDF at  $x_1/\theta_{\text{in}} = 100$ ; (c) PDF at  $x_1/\theta_{\text{in}} = 200$ ; (d) PDF at  $x_1/\theta_{\text{in}} = 300$ .

training, the strongest APG/FPG (V80) exhibits a smaller second peak near reattachment than the first near separation [Fig. 8(b)]; V90 behaves similarly, whereas V100 shows the opposite trend (second peak exceeding the first), indicating a regime change and a modest reduction in accuracy that remains acceptable. The zigzag pattern of  $C_{p,\text{rms}}$  in V20 mirrors the training case V40, reflecting increased unsteadiness near separation; this local complexity occasionally introduces minor discontinuities at patch boundaries in the generated fields. The frequency spectra are further analyzed to assess the model's ability to reconstruct coherent structures across a broad range of spatial and temporal scales. As shown in Fig. 15(c), the low-frequency components, associated with large-scale convective coherence, are accurately recovered. This behavior will be examined in greater detail through the subsequent analysis of space-time correlations. At higher frequencies, the reconstructed wall-pressure fields capture the majority of the small-scale fluctuations, although a slight overestimation is observed in regions corresponding to extremely low energy content. Should enhanced accuracy at small scales be required, the spectrally decomposed framework previously proposed by Fan *et al.* [43] may provide an effective strategy for improving model fidelity at finer scales.

Following the analysis of the frequency spectra, the space-time correlation coefficients for all testing cases are presented in Fig. 16. The convection velocities  $U_c$  inferred from the reconstructed wall-pressure fields show strong agreement with those obtained from DNS. Near the separation point  $x_1/\theta_{\text{in}} = 100$ , the model accurately captures the decreasing trend of  $U_c$  with increasing APG/FPG, reflecting the expected deceleration of coherent structures in the presence of stronger pressure gradients. Further validation is provided by the results near the bubble center  $x_1/\theta_{\text{in}} = 200$  (Fig. 42) and the reattachment point  $x_1/\theta_{\text{in}} = 300$  (Fig. 43), where the predicted convection velocities remain in close agreement with DNS. These results confirm that the model not only reconstructs




FIG. 15. Turbulence statistics for all testing APG/FPG regimes: (a) mean wall-pressure coefficient; (b) root-mean-square of wall-pressure fluctuations; (c) wall-pressure frequency spectra at three representative streamwise locations.

the instantaneous flow structures but also faithfully preserves the underlying convective transport dynamics, which are essential for accurately representing convection-dominated wall-pressure fields in turbulent flows. The quantitative and qualitative evaluations presented above collectively demonstrate the generalization capability of the proposed model in reconstructing wall-pressure fields across a range of APG/FPG conditions.

#### D. Comparison with baselines

The proposed framework is compared with both classical spatial interpolation and Gaussian process regression (GPR) based reconstructions to demonstrate the advantages. For a deterministic baseline, we compare with bicubic interpolation, which is a straightforward method for spatial reconstruction from distributed sparse sensors. As shown in Figs. 17 and 18, bicubic interpolation is able to approximately recover the mean profiles [panel (a)], but its second-order statistics are significantly underestimated [panel (b)]. In particular, the reconstructed  $C_{p,rms}$  exhibits clear oscillatory behavior, indicating that bicubic interpolation cannot adequately recover the spatial correlation structure of the wall pressure field from sparse measurements. More importantly, it fails to recover the high wave number content [panel (c)], whereas the proposed model reconstructs the small-scale structures much more accurately.

For a probabilistic baseline, we compare with GPR [59] using Case V80 as a representative example. We employ a standard exact Gaussian process with a Gaussian likelihood, a constant mean function, and a squared-exponential (RBF) kernel with separate length scales in the two spatial directions. The hyperparameters are learned by maximizing the exact marginal log-likelihood using Adam, with at most 2500 randomly selected training points to control the computational cost. In the present analysis, only the posterior mean field is used for comparison, since full-field




FIG. 16. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 100$  for unseen APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .




FIG. 17. Comparison of the proposed method and bicubic interpolation for three training cases: (a) mean wall-pressure coefficient; (b) root-mean-square of wall-pressure fluctuations; (c) wall-pressure wave number spectra at three representative streamwise locations.




FIG. 18. Comparison of the proposed method and bicubic interpolation for four testing cases: (a) mean wall-pressure coefficient; (b) root-mean-square of wall-pressure fluctuations; (c) wall-pressure wave number spectra at three representative streamwise locations.

posterior sampling is computationally expensive for this high-dimensional problem. As shown in Fig. 19, although GPR also provides a probabilistic reconstruction from sparse sensors, the reconstructed field quality remains similar to that of bicubic interpolation and is clearly inferior to the proposed method. This observation is further supported by the statistical comparisons in Fig. 20.




FIG. 19. Comparison of reconstructed wall pressure fluctuation contours obtained by the proposed method, bicubic interpolation, and GPR for Case V80.




FIG. 20. Comparison of turbulence statistics obtained by the proposed method, bicubic interpolation, and GPR for Case V80: (a) mean wall pressure coefficient; (b) root-mean-square of wall-pressure fluctuations; (c) wall-pressure wave number spectra at three representative streamwise locations.

The conclusions are largely consistent with those for bicubic interpolation, in that the high-wave number features are not recovered. The only partial improvement is that GPR provides a somewhat better prediction of  $C_{p,rms}$ , as shown in Fig. 20(b), compared with bicubic interpolation, although at a substantially higher computational cost.

Overall, these additional baseline comparisons show that both conventional deterministic interpolation and probabilistic GPR reconstruction are limited in recovering the complex multiscale and convection-dominated wall pressure field in turbulent boundary layers. In contrast, the proposed generative framework provides substantially improved reconstruction quality, especially in capturing higher-order statistics and small-scale flow structures.

### E. Computational cost (forward vs. inverse)

Figure 21 reports wall-clock time to produce one flow-through of wall-pressure fluctuations. DNS timings are **forward** runs on 64 and 3000 CPUs. The D-CNF model is trained on four NVIDIA RTX A100 GPUs, the latent diffusion model is trained on one NVIDIA RTX A100 GPU, and all inference is performed on one NVIDIA RTX A100 GPU. Counting training plus one inference [“Ours (Tot)”], our approach is about  $3\times$  faster than DNS on 64 CPUs; inference alone [“Ours (I)”] is about  $36\times$  faster than the 64-CPU DNS and still roughly  $2\times$  faster than DNS on 3000 CPUs. Because training is amortized across many reconstructions (different initial conditions, APGs, or long realizations), the effective speedup increases with reuse.

Crucially, DNS-based pressure-field reconstruction is an **inverse** problem. With variational adjoint [60,61] or Kalman-type inversion [62,63], the cost scales with the number of forward/adjoint






FIG. 21. Wall-clock time to produce one flow-through of wall-pressure fluctuations. “Ours (Tot)” includes training plus one inference; “Ours (I)” is inference only. DNS bars reflect forward runs on 64 and 3000 CPUs. For DNS-based inverse reconstruction (e.g., adjoint or Kalman inversion), the total cost scales as multiples of the forward cost (typically  $\mathcal{O}(\text{tens}) T_{\text{fwd}}$ ) and is not plotted for clarity.

evaluations,

$$T_{\text{inv}}^{\text{DNS}} \approx \begin{cases} (N_{\text{iter}}+1) T_{\text{fwd}} & (\text{adjoint/4D-Var}) \\ N_{\text{ens}} T_{\text{fwd}} & (\text{EnKF/EnKS}) \end{cases}, \quad (15)$$

which is typically on the order of tens of forward solves for practical windows. Thus a DNS-based inverse reconstruction would incur an order-of-magnitude (or more) higher cost than the forward bars shown, whereas our method’s reported cost already corresponds to solving the inverse task directly (conditioning on  $C_p$  and generating  $p_w(x_1, x_3, t)$ ). This distinction makes the practical computational advantage of the proposed approach substantially larger than suggested by forward-only comparisons.

## IV. DISCUSSION

### A. Generalizability of the domain-decomposed CNF

Our framework relies on a domain-decomposed conditional neural field (D-CNF) to map each instantaneous wall-pressure field to a compact latent and decode latents back to the physical domain. Because the diffusion model operates entirely in latent space, the D-CNF with a shared SIREN backbone and fully projected hypernetwork must generalize across flow regimes and times. We therefore evaluate the D-CNF as a standalone representation model.

As described in Sec. III A, the training dataset covers a time window from zero to four flow-through times, corresponding to 1000 snapshots. All cases share the same container parameters ( $\zeta^*$ ,  $\gamma^*$ ), while per-tile, per-snapshot latents are optimized during training (no amortized encoder). Reconstruction accuracy is measured by the relative  $L_2$  error for case  $k$ ,

$$r^k = \frac{\|\hat{p}_w^k(\mathbf{s}^*, \zeta^*, \gamma^*) - p_w^k\|_2}{\|p_w^k\|_2}. \quad (16)$$

Table II shows errors of 1%–4% for all APG/FPG cases both inside (zero to four  $T_{\text{flow}}$ ) and outside (eight to 12  $T_{\text{flow}}$ ) the training window, indicating good temporal generalization once the container is fixed. The larger error for V0 ( $\approx 16\%$ ) is consistent with its nearly homogeneous, small-scale-dominated fluctuations, which offer fewer coherent patterns for the latent model to exploit. This interpretation is supported by the t-distributed stochastic neighbor embedding (t-SNE) visualizations in Fig. 22: V0 yields a single diffuse cluster, whereas increasing APG/FPG produces well-separated clusters (two at V20 and multiple beyond V40), mirroring the emergence of larger coherent wall-pressure structures in the physical domain. To assess generalization capability, we freeze the base

TABLE II. Relative  $L_2$  error of the D-CNF (%).

| Case                                | V0    | V20  | V40  | V60  | V80  | V90  | V100 |
|-------------------------------------|-------|------|------|------|------|------|------|
| Training (0 - 4 $T_{\text{flow}}$ ) | 15.97 | 3.43 | 1.78 | 1.24 | 1.45 | 1.11 | 1.19 |
| Testing (8 - 12 $T_{\text{flow}}$ ) | 15.70 | 3.45 | 2.24 | 1.24 | 1.15 | 1.07 | 1.15 |

SIREN network parameters ( $\zeta^*$ ,  $\gamma^*$ ) and optimize the latents (i.e., autodecoding) for completely unseen snapshots from eight to 12  $T_{\text{flow}}$ , which are fully decorrelated from the training data. The latent codes  $s_{i,j}$  are initialized as zero vectors and then optimized using the Adam optimizer with a step-decay learning-rate scheduler. The optimization is terminated after a predefined number of iterations, consistent with the training setting, i.e., 20 000 iterations. The resulting errors remain comparable to those in the training window (Table II), and the reconstructed fields preserve global patterns and statistics (Fig. 44). These results indicate that the D-CNF with nonoverlapping tiling provides a robust, reusable container for convection-dominated wall-pressure fields. Empirically, a monolithic CNF without domain decomposition does not attain similar generalization in these cases, underscoring the importance of the patching strategy [46].

### B. Analysis of diffusion-generated latents

We examine the quality of the diffusion prior by comparing its generated latents (conditioned by sensors and/or the mean-profile descriptor) with the reference latents from the DNS snapshots (eight to 12 flow-through times) projected by the D-CNF. To visualize the data structure, we apply t-SNE to the combined sets for three representative regimes: the zero-pressure-gradient case V0, an interpolated case V60, and an extrapolated case V90 (Fig. 23; reference D-CNF latents are outlined).

Across all three regimes, the generated latents lie largely within the support of the reference distribution, while remaining distinct from the exact D-CNF latents. For V0, the diffusion samples reproduce the unimodal, diffuse cloud without mode collapse. For V60, the multimodal cluster pattern is recovered and populated with comparable density, with only small cluster-wise offsets, consistent with posterior sampling under DPS and sensor noise. For V90 (extrapolated), the cluster layout is preserved with a mild covariate shift, yet samples remain inside the envelope of reference latents. Taken together, these plots indicate that the diffusion model samples new realizations on the learned D-CNF manifold (no memorization) while staying close enough to enable faithful decoding




FIG. 22. T-SNE analysis of the compressed latent representations from the D-CNF for training snapshots (zero to four flow-through times): (a) zero-pressure-gradient Case V0; (b) moderate pressure-gradient Case V20, V40, and V60; and (c) strong pressure-gradient Case V80, V90, and V100.




FIG. 23. Comparison of the generated latent vectors and D-CNF latent representations in the T-SNE analysis for testing snapshots (eight to 12 flow-through times): (a) Case V0; (b) Case V60; and (c) Case V90.

back to the physical space. We emphasize that t-SNE is used for qualitative assessment of support and mode coverage; distances are not interpreted quantitatively.

### C. Novelty, limitation, and future prospects

Reconstructing flow fields from sparse sensors or low-resolution data has been studied for decades, with most methods being deterministic and focused on velocity. In contrast, our framework is probabilistic: a latent diffusion model, coupled with a D-CNF, reconstructs the full spatiotemporal wall-pressure field with calibrated uncertainty. We demonstrate two primary capabilities: (i) generating new flow realizations under fixed parameters, and (ii) zero-/few-shot generalization across APG/FPG regimes via classifier-free guidance and DPS assimilation. As a result, the method serves as an efficient surrogate able to produce many physically plausible wall-pressure fields conditioned on available sensors.

The main limitation stems from the local factorization used for efficiency. To handle high-dimensional, convection-dominated structure, the D-CNF compresses the domain into non-overlapping patches and the diffusion prior operates on per-patch latents. While effective, this design weakens cross-patch coupling: without sufficient measurements, reconstructions can be realistic within each patch yet exhibit misalignment across patches. To quantify this issue, we introduce a boundary error metric that characterizes the cross-patch misalignment at patch interfaces. Since the main artifact is the artificial discontinuity introduced between adjacent patches, we measure it through the jump of the reconstructed wall pressure across each interface, relative to the corresponding jump in the ground truth. Specifically, the boundary error is defined as

$$\varepsilon_b = \frac{1}{N_p N_b} \sum_{j=1}^{N_p} \sum_{i=1}^{N_b} \frac{|(\hat{p}_{w,i,j}^+ - \hat{p}_{w,i,j}^-) - (p_{w,i,j}^+ - p_{w,i,j}^-)|}{p_{w,rms}}, \quad (17)$$

where  $N_p$  is the number of patch interfaces,  $N_b$  is the number of boundary points on each interface,  $\hat{p}_{w,i,j}^+$  and  $\hat{p}_{w,i,j}^-$  denote the reconstructed wall pressure values on the two sides of the  $j$ -th interface at the  $i$ -th boundary point,  $p_{w,i,j}^+$  and  $p_{w,i,j}^-$  are the corresponding ground truth values, and  $p_{w,rms}$  is the root-mean-square wall pressure used for normalization. A larger value of  $\varepsilon_b$  indicates stronger cross-patch misalignment.

The boundary errors for all cases considered in this work are listed in Table III. The results show that cross-patch misalignment becomes more pronounced as the pressure gradient decreases and the flow becomes increasingly dominated by small-scale structures. This trend is also consistent with the overall error behavior reported for D-CNF in Table II, indicating that the discontinuity is closely associated with the increased complexity of the flow field and the greater difficulty of patch-wise

TABLE III. Relative boundary error ( $\varepsilon_b$ ) of the reconstructed wall-pressure field (%).

| Case            | V0    | V20   | V40  | V60  | V80  | V90  | V100 |
|-----------------|-------|-------|------|------|------|------|------|
| $\varepsilon_b$ | 33.54 | 10.29 | 4.53 | 3.45 | 3.53 | 3.37 | 4.31 |

generation under sparse sensing. In addition, we compare the boundary errors under different sensor layouts for Case V80. As shown in Fig. 45, the error increases significantly when the number of sensors is reduced, such that continuity across patch boundaries is no longer sufficiently constrained. In the current implementation, we therefore use a relatively dense  $16 \times 16$  sensor array to anchor each patch. Future work will focus on enforcing global spatial coherence with fewer sensors, for example, overlapping tiles with interface penalties on  $p_w$  and  $\nabla p_w$ , spatial autoregression or attention that exchanges boundary context between neighboring patches, and hierarchical latents that combine a coarse global code with fine local codes.

A second limitation is computational. Compared with our earlier global reconstruction variants [41,42], patch-based diffusion has higher inference cost, dominated by denoising diffusion implicit model (DDIM) sampling (typically  $\geq 300$  steps for accurate latents). Promising accelerations include single/few-step diffusion [64] or consistency models [65], progressive/distillation-style samplers, and improved latent compression (e.g., coarser global latents with patch refinements). These directions should reduce runtime while preserving fidelity and coherence at the domain scale.

## V. CONCLUSIONS

Reconstructing fully resolved, spatiotemporal wall-pressure fields in turbulent boundary layers is a challenging, ill-posed inverse problem. We presented a probabilistic generative framework that infers complete wall-pressure fields from sparse sensors, rather than predicting spectra alone. The method couples a patchwise, domain-decomposed conditional neural field (D-CNF) for compact, high-resolution representations with a latent diffusion prior that samples physically plausible realizations. Two complementary conditioning mechanisms enable generalization and training-free data assimilation. (i) Diffusion posterior sampling (DPS) incorporates arbitrary sensor layouts only at inference, decoupling training from any specific instrumentation. (ii) Classifier-free regime guidance integrates a low-cost flow descriptor during training and inference, promoting zero-/few-shot transfer across different APG/FPG conditions.

Across TBL cases spanning weak to strong pressure gradients, the framework reconstructs instantaneous fields with high fidelity and recovers key statistics, including PDFs,  $C_p$ ,  $C_{p,rms}$ , frequency spectra, space-time correlation, and convection velocity, while providing calibrated uncertainty via ensembles. These results indicate that the approach functions as an efficient, generalizable surrogate for wall-pressure dynamics under sparse sensing.

Limitations point to clear next steps. The patchwise factorization can introduce cross-patch seams under very sparse sensing; enforcing continuity via overlapping tiles with interface penalties, boundary-aware attention/autoregression, or hierarchical (global-local) latents is a priority. Inference cost is dominated by diffusion sampling; adopting distilled or consistency-based few-/single-step solvers, together with improved latent compression, should further reduce runtime. Extensions to more complex geometries, higher Reynolds numbers, and multiphysics couplings are natural directions for future work.

## ACKNOWLEDGMENTS

The authors would like to acknowledge the funds from the Office of Naval Research under Award No. N00014-23-1-2071 and the National Science Foundation under Award No. OAC-2047127.

The authors declare that they have no conflict of interest.




FIG. 24. Architecture of the domain-decomposed CNF. The top section represents the conditional components, while the bottom section depicts the container.

### DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

### APPENDIX A: ARCHITECTURE AND PARAMETERS OF DOMAIN-DECOMPOSED CNF

The overall architecture of the D-CNF used in this study is illustrated in Fig. 24. The latent dimension is set to 64. Both the container and the conditional components consist of six layers, each with 64 neurons.

The top component is the modulation network, which is a trainable neural network that adjusts the weights and biases of the container. Meanwhile, the latent variable  $z$  is learned through autodecoding, enabling the overall model to achieve strong expressivity in compressing the complex wall-pressure field. Owing to the high flexibility introduced by the decoding process, the latent space generally does not admit a direct physical interpretation. The container in D-CNF effectively learns an average wall-pressure representation from the training dataset across all patches. This can be directly observed by manually setting the latent variable to  $z = 0$  and removing the modulation, as shown in Fig. 25, where the resulting base mode is identical across all patches. This behavior is also consistent with our recent findings in [46]. The modulation network and the autodecoded latent variable are therefore primarily responsible for representing the fluctuating components around this average wall-pressure representation.




FIG. 25. Average wall pressure representation learned by the D-CNF when the modulation and latent variable are manually set to zero.



## APPENDIX B: ARCHITECTURE AND PARAMETERS OF LATENT DIFFUSION MODEL

In the latent diffusion model, a 2D U-Net with attention modules is employed as the backbone architecture for denoising, as illustrated in Fig. 26.




FIG. 26. U-Net architecture employed as the denoising component within the latent diffusion model.

## APPENDIX C: ADDITIONAL RESULTS

The instantaneous contours of the streamwise velocity component  $u$  in the  $x_1$ - $x_2$  plane are shown in Fig. 27. It is clear that, as the pressure gradient increases, the flow gradually separates from the wall and forms a recirculation bubble in the central region, exhibiting increasingly pronounced nonequilibrium flow features.




FIG. 27. Instantaneous streamwise velocity contours for varying pressure gradients.




FIG. 28. Time series of wall-pressure fluctuations at three unconditioned sensor locations for Case V0 (ZPG,  $\beta = 0$ ), including three different sample realizations.




FIG. 29. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for Case V0 (ZPG,  $\beta = 0$ ): (a) diagram showing the three locations; (b) PDF at  $x_1/\theta_{in} = 100$ ; (c) PDF at  $x_1/\theta_{in} = 200$ ; (d) PDF at  $x_1/\theta_{in} = 300$ .




FIG. 30. Time series of wall-pressure fluctuations at three unconditioned sensor locations for Case V40 (APG/FPG,  $\beta = 6.92$ ), including three different sample realizations.




FIG. 31. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for Case V40 (APG/FPG,  $\beta = 6.92$ ): (a) diagram showing the three locations; (b) PDF at  $x_1/\theta_{in} = 100$ ; (c) PDF at  $x_1/\theta_{in} = 200$ ; (d) PDF at  $x_1/\theta_{in} = 300$ .




FIG. 32. Reconstructed wall pressure fields for different sensor arrangements. The top panel shows comparisons between the reconstructed fields and the ground truth, while the bottom panel presents the corresponding turbulence statistics.




FIG. 33. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 200$  for training ZPG and APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .




FIG. 34. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 300$  for training ZPG and APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .




FIG. 35. Time series of wall-pressure fluctuations at three unconditioned sensor locations for interpolated Case V20 (APG/FPG,  $\beta = 2.14$ ), including three different sample realizations.






FIG. 36. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for interpolated Case V20 (APG/FPG,  $\beta = 2.14$ ): (a) diagram showing the three locations; (b) PDF at  $x_1/\theta_{\text{in}} = 100$ ; (c) PDF at  $x_1/\theta_{\text{in}} = 200$ ; (d) PDF at  $x_1/\theta_{\text{in}} = 300$ .




FIG. 37. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for extrapolated Case V90 (APG/FPG,  $\beta = 90.19$ ): (a) diagram showing the three locations; (b) PDF near the separation point  $x_1/\theta_{\text{in}} = 100$ ; (c) PDF near the bubble center  $x_1/\theta_{\text{in}} = 200$ ; (d) PDF near the reattachment point  $x_1/\theta_{\text{in}} = 300$ .




FIG. 38. Time series of wall-pressure fluctuations at three unconditioned sensor locations for extrapolated Case V90 (APG/FPG,  $\beta = 90.19$ ), including three different sample realizations.




FIG. 39. Reconstructed wall-pressure fluctuations for extrapolated Case V100 (APG/FPG,  $\beta = 302.55$ ), conditioned on mean profiles and sensor arrays.




FIG. 40. Probability density functions (PDFs) of wall-pressure fluctuations at three streamwise locations for extrapolated Case V100 (APG/FPG,  $\beta = 302.55$ ): (a) diagram showing the three locations; (b) PDF near the separation point  $x_1/\theta_{\text{in}} = 100$ ; (c) PDF near the bubble center  $x_1/\theta_{\text{in}} = 200$ ; (d) PDF near the reattachment point  $x_1/\theta_{\text{in}} = 300$ .




FIG. 41. Time series of wall-pressure fluctuations at three unconditioned sensor locations for extrapolated Case V100 (APG/FPG,  $\beta = 302.55$ ), including three different sample realizations.




FIG. 42. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 200$  for unseen APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .

From the streamwise space-time correlation coefficients shown in Fig. 33, it is evident that the convection velocity at the mid-streamwise location decreases significantly as the pressure gradient increases. It approaches zero when the flow becomes fully separated from the wall and a recirculation bubble forms at this location. In contrast, near the reattachment point, as shown in Fig. 34, the convection velocity remains nearly unchanged, suggesting that the flow has recovered and reattached to the wall.

The convection velocities for all the testing cases are evaluated from the streamwise space-time correlation. As shown in Fig. 42, the convection velocity decreases as the pressure gradient increases, and becomes close to zero for Case V100, where the flow is fully separated from the wall and forms a nearly steady recirculation bubble, as also indicated by the velocity contours in Fig. 27. Once the flow reattaches to the wall at  $x_1/\theta_{in} = 300$ , as shown in Fig. 43, the convection-dominated feature is recovered, and all cases exhibit a similar convection speed.




FIG. 43. Contours of the streamwise space-time correlation coefficient at the streamwise location  $x_1/\theta_{in} = 300$  for unseen APG/FPG cases: (a) DNS results; (b) reconstructed field. Slopes of the dashed lines indicate the convection velocities  $U_c$ .




FIG. 44. Results from the D-CNF applied to testing snapshots corresponding to eight to 12 flow-through times: (a) mean pressure coefficient  $C_p$ ; (b) root-mean-square of pressure fluctuations  $C_{p,rms}$ ; and (c) selected instantaneous wall-pressure contours from DNS and decoded by the D-CNF. Solid lines represent DNS results, while dashed lines denote the results obtained from the D-CNF.

The accuracy of D-CNF is thoroughly evaluated in Fig. 44. As shown in panel (a), the mean  $C_p$  profiles are well reconstructed, indicating that D-CNF captures the mean features accurately. In panel (b), the root-mean-square wall pressure fluctuations are generally consistent with the DNS results, except near the separation region ( $x_1/\theta_{in} = 100 - 150$ ), where the flow is highly unsteady, and the D-CNF prediction exhibits some oscillatory behavior. The overall field patterns are nevertheless well captured for all cases, as illustrated by Case V20 and Case V100 in panels (c) and (d).




FIG. 45. Effect of sensor layouts on the relative boundary error of the reconstructed wall-pressure field for Case V80.

- [1] Y. Na and P. Moin, The structure of wall-pressure fluctuations in turbulent boundary layers with adverse pressure gradient and separation, *J. Fluid Mech.* **377**, 347 (1998).
- [2] W. Devenport, N. Alexander, S. Glegg, and M. Wang, The sound of flow over rigid walls, *Annu. Rev. Fluid Mech.* **50**, 435 (2018).
- [3] S. Lee, L. Ayton, F. Bertagnolio, S. Moreau, T. P. Chong, and P. Joseph, Turbulent boundary layer trailing-edge noise: Theory, computation, experiment, and application, *Prog. Aerosp. Sci.* **126**, 100737 (2021).
- [4] M. Wang, J. B. Freund, and S. K. Lele, Computational prediction of flow-generated sound, *Annu. Rev. Fluid Mech.* **38**, 483 (2006).
- [5] B. Yang and Z. Yang, On the wavenumber–frequency spectrum of the wall pressure fluctuations in turbulent channel flow, *J. Fluid Mech.* **937**, A39 (2022).
- [6] G. Grasso, P. Jaiswal, H. Wu, S. Moreau, and M. Roger, Analytical models of the wall-pressure spectrum under a turbulent boundary layer with adverse pressure gradient, *J. Fluid Mech.* **877**, 1007 (2019).
- [7] S. L. Prigent, É. Salze, and C. Bailly, Deconvolution of wave-number-frequency spectra of wall pressure fluctuations, *AIAA J.* **58**, 164 (2020).
- [8] H. Choi and P. Moin, Grid-point requirements for large eddy simulation: Chapman’s estimates revisited, *Phys. Fluids* **24**, 011702 (2012).
- [9] M. Slama, C. Leblond, and P. Sagaut, A Kriging-based elliptic extended anisotropic model for the turbulent boundary layer wall pressure spectrum, *J. Fluid Mech.* **840**, 25 (2018).
- [10] L. T. L. Pereira, F. Avallone, D. Ragni, and F. Scarano, A physics-based description and modelling of the wall-pressure fluctuations on a serrated trailing edge, *J. Fluid Mech.* **938**, A28 (2022).
- [11] R. L. Panton and J. H. Linebarger, Wall pressure spectra calculations for equilibrium boundary layers, *J. Fluid Mech.* **65**, 261 (1974).
- [12] N. Thomson and J. Rocha, Comparison of semi-empirical single point wall pressure spectrum models with experimental data, *Fluids* **6**, 270 (2021).
- [13] D. J. Fritsch, V. Vishwanathan, C. J. Roy, K. Todd Lowe, W. J. Devenport, P. Croaker, O. Tkachenko, D. Pook, G. Lane, S. Shubham, *et al.*, Modeling the surface pressure spectrum beneath turbulent boundary layers in pressure gradients, *AIAA J.* **61**, 2002 (2023).
- [14] J. Dominique, J. Christophe, C. Schram, and R. D. Sandberg, Inferring empirical wall pressure spectral models with gene expression programming, *J. Sound Vib.* **506**, 116162 (2021).
- [15] S. Shubham, S. Pargal, S. Moreau, R. D. Sandberg, J. Yuan, A. Kushari, and M. Sanjose, Data-driven empirical wall pressure spectrum models for fan noise prediction, in: *AIAA Aviation 2023 Forum* (American Institute of Aeronautics and Astronautics, San Diego, CA, 2023), p. 3508.



- [16] S. Shubham, R. D. Sandberg, and A. Kushari, More general wall pressure spectra models: Combining feature engineering with evolutionary algorithms, *AIAA J.* **62**, 2622 (2024).
- [17] Y. Jin, Y. Li, F. Liao, and J. Cai, Assessment of recent empirical wall-pressure auto-spectrum models under various pressure gradient conditions, *Appl. Math. Modell.* **146**, 116154 (2025).
- [18] L. Botero-BolÁvar, D. Huerger, F. L. dos Santos, C. H. Venner, L. D. de Santana, and E. Ferrer, An empirical wall-pressure spectrum model for aeroacoustic predictions based on symbolic regression, *Appl. Acoust.* **240**, 110876 (2025).
- [19] N. N. Kurhade, N. R. Vadlamani, and A. Haridas, Artificial neural networks and guided gene expression programming to predict wall pressure spectra beneath turbulent boundary layers, *Phys. Fluids* **35**, 085123 (2023).
- [20] J. Dominique, J. Van den Berghe, C. Schram, and M. Mendez, Artificial neural networks modeling of wall pressure spectra beneath turbulent boundary layers, *Phys. Fluids* **34**, 035119 (2022).
- [21] S. Zeng, K. Liu, W. Yang, M. Fu, G. Wu, and J. Zheng, Rapid prediction of rocket wall pressure fluctuation spectra using machine learning, *Phys. Fluids* **37**, 035191 (2025).
- [22] K. Poulinakis, D. Drikakis, I. W. Kokkinakis, and S. M. Spottswood, Deep learning reconstruction of pressure fluctuations in supersonic shock–boundary layer interaction, *Phys. Fluids* **35**, 076117 (2023).
- [23] K. Fukami, K. Fukagata, and K. Taira, Super-resolution reconstruction of turbulent flows with machine learning, *J. Fluid Mech.* **870**, 106 (2019).
- [24] K. Fukami, K. Fukagata, and K. Taira, Super-resolution analysis via machine learning: A survey for fluid flows, *Theor. Comput. Fluid Dyn.* **37**, 421 (2023).
- [25] K. Fukami, B. An, M. Nohmi, M. Obuchi, and K. Taira, Machine-learning-based reconstruction of turbulent vortices from sparse pressure sensors in a pump sump, *J. Fluids Eng.* **144**, 121501 (2022).
- [26] B. A. Danciu, V. A. Pagone, B. Böhm, M. Schmidt, and C. E. Frouzakis, Flow reconstruction in time-varying geometries using graph neural networks, *arXiv:2411.08764*.
- [27] F. Sofos, D. Drikakis, and I. W. Kokkinakis, Comparison of super-resolution deep learning models for flow imaging, *Comput. Fluids* **283**, 106396 (2024).
- [28] R. Cheng, A. Shamooni, T. Zirwes, and A. Kronenburg, Improved super-resolution reconstruction of turbulent flows with spectral loss function, *Phys. Fluids* **37**, 035208 (2025).
- [29] K. Zeng, Y. Zhang, H. Xu, and X. Feng, Super-resolution reconstruction of turbulent flows with a hybrid framework of attention, *Phys. Fluids* **36**, 065107 (2024).
- [30] M. Diop, P. Dubois, H. Toubin, L. Planckaert, J.-F. Le Roy, and E. Garnier, Reconstruction of flow around a high-rise building from wake measurements using machine learning techniques, *J. Wind Eng. Ind. Aerodyn.* **230**, 105149 (2022).
- [31] J.-W. Hu and W.-W. Zhang, Mesh-Conv: Convolution operator with mesh resolution independence for flow field modeling, *J. Comput. Phys.* **452**, 110896 (2022).
- [32] V. Yadav, M. Casel, and A. Ghani, RF-PINNs: Reactive flow physics-informed neural networks for field reconstruction of laminar and turbulent flames using sparse data, *J. Comput. Phys.* **524**, 113698 (2025).
- [33] L. Zhu, X. Jiang, A. Lefauve, R. R. Kerswell, and P. Linden, New insights into experimental stratified flows obtained through physics-informed neural networks, *J. Fluid Mech.* **981**, R1 (2024).
- [34] C. Lee, Y. Ozawa, T. Nagata, T. Colonius, and T. Nonomura, Superresolution and analysis of three-dimensional velocity fields of underexpanded jets in different screech modes, *Phys. Rev. Fluids* **9**, 104604 (2024).
- [35] S. Rühling Cachay, B. Zhao, H. Joren, and R. Yu, Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting, *Adv. Neural Inf. Process. Syst.* **36**, 45259 (2023).
- [36] R. Molinaro, S. Lanthaler, B. Raonić, T. Rohner, V. Armegioiu, S. Simonis, D. Grund, Y. Ramic, Z. Y. Wan, F. Sha, *et al.*, Generative AI for fast and accurate statistical computation of fluids, *arXiv:2409.18359*.
- [37] D. Shu, Z. Li, and A. B. Farimani, A physics-informed diffusion model for high-fidelity flow field reconstruction, *J. Comput. Phys.* **478**, 111972 (2023).
- [38] G. Kohl, L.-W. Chen, and N. Thuerey, Benchmarking autoregressive conditional diffusion models for turbulent flow simulation, in ICML 2024 AI for Science Workshop (OpenReview.net, Vienna, Austria, 2024).

- [39] H. Gao, X. Han, X. Fan, L. Sun, L.-P. Liu, L. Duan, and J.-X. Wang, Bayesian conditional diffusion models for versatile spatiotemporal turbulence generation, [Comput. Methods Appl. Mech. Eng. \*\*427\*\*, 117023 \(2024\)](#).
- [40] Y. Shehata, B. Holzschuh, and N. Thuerey, Improved sampling of diffusion models in fluid dynamics with Tweedie’s formula, in: *The Thirteenth International Conference on Learning Representations* (OpenReview.net, Singapore, 2025).
- [41] P. Du, M. H. Parikh, X. Fan, X.-Y. Liu, and J.-X. Wang, Conditional neural field latent diffusion model for generating spatiotemporal turbulence, [Nat. Commun. \*\*15\*\*, 10416 \(2024\)](#).
- [42] X.-Y. Liu, M. H. Parikh, X. Fan, P. Du, Q. Wang, Y.-F. Chen, and J.-X. Wang, CoNFILD-inlet: Synthetic turbulence inflow using generative latent diffusion models with neural fields, [Phys. Rev. Fluids \*\*10\*\*, 054901 \(2025\)](#).
- [43] X. Fan, D. Akhare, and J.-X. Wang, Neural differentiable modeling with diffusion-based super-resolution for two-dimensional spatiotemporal turbulence, [Comput. Methods Appl. Mech. Eng. \*\*433\*\*, 117478 \(2025\)](#).
- [44] M. Hemant Parikh, X. Fan, and J.-X. Wang, Conditional flow matching for generative modeling of near-wall turbulence with quantified uncertainty, [J. Fluid Mech. \*\*1029\*\*, A32 \(2026\)](#).
- [45] S. Meloni, F. Centracchio, E. De Paola, R. Camussi, and U. Iemma, Experimental characterisation and data-driven modelling of unsteady wall pressure fields induced by a supersonic jet over a tangential flat plate, [J. Fluid Mech. \*\*958\*\*, A27 \(2023\)](#).
- [46] J. Guo, P. Du, X. Fan, Y. Li, and J.-X. Wang, Conditional neural field for spatial dimension reduction of turbulence data: A comparison study, [Phys. Fluids \*\*38\*\*, 025153 \(2026\)](#).
- [47] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, Implicit neural representations with periodic activation functions, [Adv. Neural Inf. Process. Syst. \*\*33\*\*, 7462 \(2020\)](#).
- [48] J. Ho and T. Salimans, Classifier-free diffusion guidance, [arXiv:2207.12598](#).
- [49] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, Diffusion posterior sampling for general noisy inverse problems, in *The Eleventh International Conference on Learning Representations* (OpenReview.net, Kigali, Rwanda, 2023).
- [50] B. Efron, Tweedie’s formula and selection bias, [J. Am. Stat. Assoc. \*\*106\*\*, 1602 \(2011\)](#).
- [51] K. Kim and J. C. Ye, Noise2Score: Tweedie’s approach to self-supervised image denoising without clean images, [Adv. Neural Inf. Process. Syst. \*\*34\*\*, 864 \(2021\)](#).
- [52] D. You, F. Ham, and P. Moin, Discrete conservation principles in large-eddy simulation with application to separation control over an airfoil, [Phys. Fluids \*\*20\*\*, 101515 \(2008\)](#).
- [53] H. Abe, Reynolds-number dependence of wall-pressure fluctuations in a pressure-induced turbulent separation bubble, [J. Fluid Mech. \*\*833\*\*, 563 \(2017\)](#).
- [54] H. Abe, Direct numerical simulation of a turbulent boundary layer with separation and reattachment over a range of Reynolds numbers, [Fluid Dyn. Res. \*\*51\*\*, 011409 \(2019\)](#).
- [55] T. S. Lund, X. Wu, and K. D. Squires, Generation of turbulent inflow data for spatially-developing boundary layer simulations, [J. Comput. Phys. \*\*140\*\*, 233 \(1998\)](#).
- [56] V. Kitsios, A. Sekimoto, C. Atkinson, J. A. Sillero, G. Borrell, A. G. Gungor, J. Jiménez, and J. Soria, Direct numerical simulation of a self-similar adverse pressure gradient turbulent boundary layer at the verge of separation, [J. Fluid Mech. \*\*829\*\*, 392 \(2017\)](#).
- [57] P. R. Spalart, Direct simulation of a turbulent boundary layer up to  $R_\theta = 1410$ , [J. Fluid Mech. \*\*187\*\*, 61 \(1988\)](#).
- [58] Y. Na and P. Moin, Direct numerical simulation of a separated turbulent boundary layer, [J. Fluid Mech. \*\*374\*\*, 379 \(1998\)](#).
- [59] J. Wang, An intuitive tutorial to Gaussian process regression, [Comput. Sci. Eng. \*\*25\*\*, 4 \(2023\)](#).
- [60] O. Ashtari and T. M. Schneider, Identifying invariant solutions of wall-bounded three-dimensional shear flows using robust adjoint-based variational techniques, [J. Fluid Mech. \*\*977\*\*, A7 \(2023\)](#).
- [61] M. Wang and T. A. Zaki, Variational data assimilation in wall turbulence: From outer observations to wall stress and pressure, [J. Fluid Mech. \*\*1008\*\*, A26 \(2025\)](#).
- [62] J.-X. Wang and H. Xiao, Data-driven CFD modeling of turbulent flows through complex structures, [Int. J. Heat Fluid Flow \*\*62\*\*, 138 \(2016\)](#).

- [63] J.-X. Wang, R. Sun, and H. Xiao, Quantification of uncertainties in turbulence modeling: A comparison of physics-based and random matrix theoretic approaches, [Int. J. Heat Fluid Flow](#) **62**, 577 (2016).
- [64] T. Yin, M. Gharbi, R. Zhang, E. Shechtman, F. Durand, W. T. Freeman, and T. Park, One-step diffusion with distribution matching distillation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, Seattle, WA, 2024), pp. 6613–6623.
- [65] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever, Consistency models, [arXiv:2303.01469](#).