

High-fidelity reconstruction of large-area damaged turbulent fields with a physically constrained generative adversarial network

Qinmin Zheng,^{1,2} Tianyi Li³, Benteng Ma¹, Lin Fu^{2,4,5,*} and Xiaomeng Li^{1,†}

¹*Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong*

²*Department of Mechanical and Aerospace Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong*

³*Department of Physics and INFN, University of Rome "Tor Vergata",
Via della Ricerca Scientifica 1, Rome 00133, Italy*

⁴*Department of Mathematics and Center for Ocean Research in Hong Kong and Macau (CORE), The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong*

⁵*HKUST Shenzhen-Hong Kong Collaborative Innovation Research Institute, Futian, Shenzhen, People's Republic of China*



(Received 9 October 2023; accepted 17 January 2024; published 29 February 2024)

The reconstruction of incomplete information in flow fields is a pervasive challenge in numerous turbulence-related applications. This paper proposes a framework for the high-fidelity reconstruction of large-area damaged turbulent fields with high resolution based on a physically constrained generative adversarial network. The network incorporates some special designs, such as leveraging complete or sparse fields of velocity components as physical constraints, employing a ResNet-like network with a fast Fourier convolution module, and adopting a PatchGAN discriminator network. To train the network, we employ a combination of loss functions, which comprise mean absolute error, network-level loss, and feature matching loss. We validate our method on a dataset of compressible isotropic turbulent flow and investigate three distinct damaged distributions with gap ratios of 39.06% and 56.25%. The proposed reconstruction framework has been shown to achieve excellent reconstruction performance. The reconstructed flow fields are consistent with the raw flow fields in terms of magnitude, power spectrum, and two-point correlation function. We also provide extensive ablation studies to validate our approach. The results indicate that the utilization of physical constraints significantly improves the reconstruction performance in damaged regions, particularly in cases that involve a large damaged area. Furthermore, in the proposed reconstruction framework, the size of patch in the PatchGAN discriminator network can be flexibly adjusted according to the scale of turbulence structures, and the perceptual loss function plays a crucial role in evaluating differences between feature maps of flow fields at network levels based on a pretrained network. These special designs ensure consistency across diverse scales of turbulence structures and improve the accuracy and efficiency of network training.

DOI: [10.1103/PhysRevFluids.9.024608](https://doi.org/10.1103/PhysRevFluids.9.024608)

*linfu@ust.hk

†eexmli@ust.hk

I. INTRODUCTION

The challenge of reconstructing incomplete information of flow fields, caused by the inadequate spatial resolution in numerical simulations and/or the measurement limitations in experiments, is a pervasive issue in nearly all significant applications related to turbulence [1–4]. The particle image velocimetry (PIV) experiment is one of the notable examples in which incomplete data can arise from various sources, such as out-of-plane particles, object shadows, or light reflections [5,6]. Additionally, oceanographic and meteorological observations also suffer from similar issues, where dead pixels and thick cloud cover may lead to missing information [7–9]. Various approaches can be followed to reconstruct a full-field description from incomplete measurements, depending on the degree of incompleteness and the number of available samples. If there are a sufficiently large number of samples and the degree of incompleteness is limited, reconstruction techniques that rely on extracting flow features, such as proper orthogonal decomposition (POD) [10,11], dynamic mode decomposition [12,13], and their variants, can be utilized. For instance, the Gappy-POD [14–17] reconstruction procedures estimate the nongappy POD modes and time coefficients by the least-squares regression on the available incomplete fields; the reconstructed full fields are later calculated based on the estimated POD modes. However, these model reduction approaches essentially rely on linear interpolation, which poses a great challenge when dealing with complex multiscale flows (such as fully developed turbulence) and scenarios in which extreme events occurred in damaged regions [18,19].

In recent years, convolutional neural networks have achieved great success in computer vision tasks and have gradually been applied to the fields of fluid dynamics [20–22], geophysics [23,24], and astrophysics [25,26], due to their ability to handle nonlinearities [27,28]. A typical case is the successful application of generative adversarial network (GAN)-based super-resolution frameworks to turbulent flows and climate data, to improve spatial resolution and/or fill gaps in incomplete measurements [3,23,29–33]. Recently, GAN-based deep learning approaches have also been used to reconstruct three-dimensional turbulent flows from two-dimensional observations [34]. A generative adversarial network consists of two convolutional neural networks, i.e., a generator network and a discriminator network, which operate in an adversarial manner to achieve the task of generative modeling. Numerous evidence suggests that the introduction of a discriminator network significantly improves the accuracy of reconstructed fields [35]. For instance, Kim *et al.* [36] developed an unsupervised learning model (CycleGAN) for the super-resolution reconstruction with unpaired turbulence data. The model was validated using the filtered data from the direct numerical simulation (DNS) of homogeneous isotropic turbulence and the partially measured data from DNS of turbulent channel flows. The results suggested that this unsupervised model demonstrated qualitatively and quantitatively similar performances as that of the best supervised learning model. In addition, Güemes *et al.* [1] introduced a super-resolution GAN framework, named RaSeedGAN, which can effectively estimate field quantities from random sparse sensors without relying on full-field high-resolution representations during training. The proposed approach was evaluated on synthetic databases of flow simulations, ocean surface temperature measurements, and PIV data of turbulent boundary layer flows. The results demonstrated the excellent performance of RaSeedGAN, even in scenarios with high levels of sparsity or noise. Unlike the reconstruction from sparse measurements, as an extreme example, Li *et al.* [18] and Buzzicotti *et al.* [19] addressed the problem of reconstructing a series of two-dimensional large-area damaged snapshots of three-dimensional rotating turbulence with the deep-GAN approach. They further compared their results with those obtained from linear data-driven techniques, namely Gappy-POD and Extend-POD. The results revealed that the deep-GAN approach exhibited significant advantages over linear data-driven methods in the reconstruction of turbulence, particularly in scenarios where extreme events occurred in damaged regions.

In prior research, most reconstructed flows have exhibited pronounced coherent structures, such as the Kármán vortex streets in turbulent or laminar wakes of bluff bodies or the hairpin vortices and low-speed streaks in wall-bounded turbulent flows. Coherent structures have statistically clear

motion patterns that can be identified and analyzed. For instance, the normalized shedding frequency of the Kármán vortex is relatively fixed depending on the freestream Reynolds number [37–39], while the streamwise inclination angles of wall-attached eddies in turbulent channel flows tend to saturate at 45° as the Reynolds number increases [40,41]. For such flows, current deep learning methods are typically capable of effectively reconstructing the dominant large-scale coherent structures in the flow fields while tending to overlook the small-scale turbulence structures. This is manifested as a loss of small-scale structures in the reconstructed flow fields and discrepancies in the high-wave-number regions of the power spectrum [34,36,42]. Small-scale structures in turbulent flows are often obscured and attenuated by deep learning-based reconstruction frameworks due to the lack of explicit motion rules, which are required to minimize the loss function. However, in the context of turbulence research, small-scale structures are of particular significance, especially for multiscale energy transport and high-order statistics of turbulence [43,44]. In contrast to previous studies, the main contribution of the present investigation is the development of a more refined reconstruction framework for the high-fidelity reconstruction of large-area damaged turbulent fields without obvious coherent structures. For complex systems where the governing equations are unclear or where the governing equations are difficult to implant into deep learning networks, it is one of the effective approaches to add physical constraints to deep learning networks by integrating physical information into the loss function, such as the flow gradient, Reynolds stress tensor, and spectral content [45–47]. Different from existing literature, in the current framework, we propose a flexible approach to integrate additional physical constraints into the reconstruction process. Reference fields are incorporated into the network as physical constraints to improve the accuracy of reconstructed flow fields. Our approach stems from application scenarios where measurement information is missing in certain regions due to factors such as shadows or obstructions. To address this issue, advanced deep learning techniques can be leveraged to infer missing information based on available measurements within the same region. A concrete example is satellite-based measurement of sea level temperature, where cloud obstruction often leads to data gaps in certain regions. In such cases, we can utilize the wind speed and pressure data obtained from the measurement buoys on the sea surface located within the same region to reconstruct the temperature information for the affected region. The proposed framework achieves excellent performance in reconstructing small-scale structures in turbulent flows, accurately reproducing the original flow fields at high wave numbers in the power spectrum. To the best of our knowledge, the reconstruction performance of the proposed framework surpasses existing methods in the literature.

The paper is organized as follows. The dataset and the reconstruction problem are described in Sec. II A, while the reconstruction framework, the combination of loss functions and the implementation details are introduced in Secs. II B, II C, and II D, respectively. In Sec. III, we present a detailed analysis of the performance of the proposed framework. This analysis includes qualitative comparisons of instantaneous contours (Sec. III A), as well as quantitative evaluations of single-point statistics (Sec. III B), power spectrum and two-point correlation functions (Sec. III C). To establish a benchmark for evaluation, in Sec. IV A, we replicate the deep convolutional generative adversarial network (DCGAN) that has been widely used in the prior literature. We then proceed to assess the quantitative and qualitative performance of DCGAN in the current reconstruction problem. Additionally, we investigated the impacts of loss functions and physical constraints within the proposed reconstruction framework on reconstruction performance in Sec. IV B. We highlight the importance of the PatchGAN discriminator and the perceptual loss function to achieve accurate reconstruction, and discuss the effects of training data amount and the network’s robustness against noisy input in Sec. V. Finally, a comprehensive conclusion about our findings and contributions is provided in Sec. VI.

II. METHODOLOGY

A. Dataset and reconstruction problem set-up

In this work, we employed two-dimensional slices of an isotropic turbulent flow as our dataset to validate the performance of the proposed framework in reconstructing multiscale structures in turbulent flows. Isotropic turbulence, first proposed by G. I. Taylor in 1935, is a well-established and idealized turbulence system that has been widely investigated in the fluid community. Here, we use the data of Case II₅ in Ref. [48] as a dataset, which is a compressible isotropic turbulent flow with vibrational nonequilibrium and has been analyzed in our previous publication [48–50]. The following is a brief introduction to the dataset, and the readers are referred to Ref. [48] for details of computational methods and statistical information. The governing equations of compressible turbulence are numerically solved in a cubic domain with a side length of 2π and a grid resolution of 512^3 . To maintain the total kinetic energy in the first two shells at the prescribed levels, a large-scale momentum forcing is applied to the solenoidal velocity component $\hat{\mathbf{u}}^S(\mathbf{k}, t)$, while the dilatational velocity component $\hat{\mathbf{u}}^D(\mathbf{k}, t)$ is left untouched. That is, the forced velocity field $\hat{\mathbf{u}}^f(\mathbf{k}, t) = \xi \hat{\mathbf{u}}^S(\mathbf{k}, t) + \hat{\mathbf{u}}^D(\mathbf{k}, t)$, where ξ is the scaling factor. Furthermore, to enhance the flow compressibility and maintain the internal energy of the system in a statistically steady state, we employ a large-scale thermal forcing that acts on the translational-rotational energy, as well as a spatially uniform thermal cooling function. After the system reaches the statistically stationary state, 121 flow fields, uniformly spanning the time period of $9.01 \leq t/T_e \leq 28.82$, are sampled every 1000 numerical time steps. Here, T_e is the large eddy turnover time. For each three-dimensional flow field sampled, 128 two-dimensional slices are uniformly taken along a single direction. This results in a total of 15,488 slices, of which 12,480 slices are used as the training dataset, while 1,500 and 1,500 slices are used as validation and testing datasets, respectively. Note that, the training, validation, and testing datasets are extracted in chronological order without overlapping with each other. Furthermore, during network training, the slices are randomly sampled from the training dataset to minimize their temporal correlation. Figure 1(a) depicts the instantaneous isosurfaces of velocity divergence at a given time. Regions with positive and negative velocity divergence represent expansion and compression regions, respectively. The compression regions exhibit sheetlike structures, while the expansion regions have blocklike structures. The sheetlike compression regions are commonly referred to as shocklet structures, which induce discontinuities in physical quantities, such as density and temperature, within the flow field. As illustrated in Fig. 1(a), the lack of discernible coherent structures suggests a high degree of complexity and irregularity in the compressible isotropic turbulent flow. Figure 1(b) displays the velocity power spectrum, which exhibits a distinct inertial subrange, indicating the cascade process in turbulent kinetic energy transport. The shaded region in Fig. 1(b) represents the scale range of energy injection. From a statistical perspective, within the inertial subrange, the injected energy undergoes a gradual transfer from large-scale to small-scale eddies through the cascade process [44]. This is the charm of isotropic turbulence—it contains rich statistical properties of turbulence within the highly irregular and multiscale flow motion. As illustrated in Fig. 1(c), three different types of masks are utilized in our experiments, including a single-damaged mask, a multiple-damaged mask, and a random-damaged mask. For each type of mask, we used two different gap ratios, namely 39.06% and 56.25%. As an example, for the single-damaged mask with a gap ratio of 56.25% [Fig. 1(c₁)], the damaged region (384×384 elements) is set to 0.0, while the rest of the mask is given to be 1.0. In this work, the shapes of damaged regions are identical for the multiple-damaged mask [Fig. 1(c₂)] and the random-damaged mask [Fig. 1(c₃)], being 192×192 elements and 16×16 elements, respectively. The first two types of masks [Figs. 1(c₁) and 1(c₂)] are used to assess the performance of the proposed framework in reconstructing the flow field when the area of damaged regions is extensive and their locations remain fixed. The third type of mask [Fig. 1(c₃)] is employed to evaluate the network’s reconstruction performance when the damaged regions are small in size but their positions change randomly in different slices. Consequently, during network training, the location of the damaged regions is randomly altered for this third case. However, for the validation




FIG. 1. (a) The instantaneous isosurface of velocity divergence, and (b) the velocity power spectrum of compressible isotropic turbulent flow; (c) three different masks with a gap ratio of 56.25%: a single-damaged mask, a multiple-damaged mask, and a random-damaged mask.

and testing datasets, we deliberately keep the position of damaged regions fixed. This enables us to compare the network's reconstruction performance under different network configurations, such as using various loss functions. It is worth highlighting that, for the third type of mask, once the network training is completed, the damaged flow field can be effectively reconstructed even the position of damaged regions changes randomly in different slices.

B. Physically constrained generative adversarial network (PCGAN)

Figure 2 illustrates the proposed reconstruction framework, namely PCGAN, which consists of the generator and the discriminator networks. This framework aims to reconstruct a turbulent field x masked by a binary mask m , with the assistance of two complete or sparse reference fields R_{ref} as physical constraints. Specifically, the objective is to reconstruct the damaged density field with the two complete or sparse reference fields of velocity components (u and v). It is worth highlighting that, for homogeneous turbulence, although the different velocity components exhibit a statistical consistency, the difference between instantaneous fields of different velocity components at the same moment can be substantial. In addition, because turbulence is governed by a set of highly coupled physical quantities, the damaged and reference fields may refer to other quantities in practical applications, such as temperature or pressure. The improvement of reconstruction performance based on the reference field largely depends on the correlation between the physical quantities of the damaged and reference fields, and this correlation is closely related to flow systems. Furthermore, these




FIG. 2. (a) Framework of generator network for flow field reconstruction. The framework is based on ResNet-like feedforward networks with a fast Fourier convolution module; (b) specific implementation of the fast Fourier convolution; (c) specific implementation of the PatchGAN discriminator network.

reference fields can be complete flow fields or sparse measurement data, whose completeness will be determined based on practical considerations, such as the resolution of the measurement equipment. This framework provides a flexible approach to integrating additional physical constraints into the reconstruction process. The information provided by the additional physical constraints has a significant impact on the reconstruction performance, which will be discussed in Sec. 4.2.2. The masked field is denoted as $x' = x \odot m$. Taking x' and R_{ref} , the feedforward reconstructing network $\mathcal{F}_\theta(\cdot)$ processes the input in a fully convolutional manner, and produces a reconstructed flow field $\hat{x} = \mathcal{F}_\theta(x', R_{ref})$ [Fig. 2(a)]. The discriminator network takes the real data (ground truths) and the fake data (reconstructed fields) as inputs and determines their authenticity. The generator network attempts to deceive the discriminator network, while the discriminator network strives to

make accurate judgments. Consequently, both networks learn to produce “realistic” information spontaneously.

As shown in Fig. 2(a), the generator network $\mathcal{F}_\theta(\cdot)$ includes three downsampling blocks, nine residual blocks, and three upsampling blocks. By employing a consistent convolutional kernel size throughout the generator network, downsampling operations based on convolution and max pooling facilitate the extraction of global characteristics and high-level abstractions from flow fields by reducing its spatial resolution, while upsampling operations with deconvolution are instrumental in restoring fine-grained details and spatial resolution to the flow field, while also inferring flow field information within damaged regions. The residual block is based on ResNet-like feedforward networks, in which the standard convolution module is replaced by the fast Fourier convolution (FFC) module [51,52]. The FFC module is based on a channel-wise fast Fourier transform (FFT) and has a receptive field that covers the entire field [53,54]. As depicted in Fig. 2(b), the FFC module splits channels into two parallel branches, namely local and global branches. The local branch employs conventional convolution, while the global branch utilizes the real FFT to account for the global context. Subsequently, outputs of the local and global branches are fused together, resulting in a more informed and accurate representation of the input data.

The discriminator network is fed with the real and the reconstructed flow fields. As depicted in Fig. 2(c), the network initialization is accomplished by performing a convolution operation using a filter size of 4×4 and 64 feature maps, followed by a LeakyReLU activation function. Subsequently, four discriminator blocks are sequentially applied, each consisting of a convolution operation, a BatchNorm operation, and a LeakyReLU activation function. Last, the network concludes with a convolution operation that uses a filter size of 4×4 and a single feature map. In the discriminator network, we utilize the “PatchGAN” discriminator, which differs from the global GAN discriminator that evaluates the entire field as a whole. The PatchGAN discriminator imposes a penalty solely on the structure at the patch level [55,56]. The primary objective of the discriminator is to classify each input patch as real or fake. Subsequently, the discriminator is convolved across the entire field, and the responses are averaged to produce the final result. It is worth noting that the utilization of the PatchGAN discriminator is of utmost significance in the accurate reconstruction of flow field details. A comprehensive explanation will be given in Sec. V.

C. Combination of loss functions

1. Perceptual loss function

Reconstructing the flow field in a damaged region is inherently ill-posed. This means that there could be many plausible solutions for the same damaged region, especially when the damaged region becomes larger in size. Conventional element-wise losses, such as mean absolute error (MAE) or mean-squared error (MSE), demand the generator network to precisely reconstruct the ground truth, but the unmasked parts of flow fields often do not contain enough information to exactly reconstruct the flow field within the damaged regions. Consequently, using conventional element-wise losses leads to blurry results due to the averaging of multiple plausible solutions of the reconstructed content [56]. Here, we employ the perceptual loss function, which evaluates the difference between features extracted from the reconstructed and target flow fields by a pretrained network $\phi(\cdot)$ [57]. It does not require an exact reconstruction, which allows variations in damaged regions. In the proposed framework, the pretrained network uses the ResNet50dilated module as encoder, and is pretrained on the ADE20K dataset [58,59]. It is worth mentioning that the pretrained network is solely employed to calculate the perceptual loss function with its weights held fixed. The pretrained network computes the feature maps of the input fields from the training dataset and their corresponding output across its diverse modules, while the perceptual loss is utilized to assess the differences within these feature maps. Consequently, the specific training dataset employed for the pretrained network is not a critical factor. The perceptual loss associated with the j th layer of

the pretrained network is given as

$$\ell_{\text{feat}}^{\phi,j}(x, \hat{x}) = \frac{1}{C_j H_j W_j} \|\phi_j(x) - \phi_j(\hat{x})\|_2^2, \quad (1)$$

where x and $\hat{x}$ are the original and the reconstructed flow fields, respectively. ϕ_j represents the feature map of the j th layer in the pretrained network. C_j , H_j and W_j are the dimensions of the feature map. The overall perceptual loss $\mathcal{L}_{\text{PL}}$ corresponds to the mean of the individual perceptual loss across various layers of the pretrained network. For the specific calculation details of the perceptual loss, the readers are referred to Sec. VB.

2. Adversarial loss function

The adversarial loss is included to enhance the fidelity of the flow field details generated by the reconstruction network $\mathcal{F}_\theta(\cdot)$. As mentioned above, the PatchGAN discriminator network $D_\xi(\cdot)$ is defined to work at a local patch level, discriminating between “real” and “fake” patches. Due to the introduction of perceptual loss, the generator network is expected to quickly copy the unmasked parts of the input field. Therefore, the unmasked parts of the reconstructed flow field are labeled as “real,” and only patches that interact with damaged regions get the “fake” label. Finally, the nonsaturating adversarial loss $\mathcal{L}_{\text{Adv}}$ is defined as

$$\begin{aligned} \mathcal{L}_{\text{Dis}} = & -\mathbb{E}_x[\log D_\xi(x)] - \mathbb{E}_{x,m}[\log D_\xi(\hat{x}) \odot m] \\ & - \mathbb{E}_{x,m}[\log(1 - D_\xi(\hat{x})) \odot (1 - m)], \end{aligned} \quad (2)$$

$$\mathcal{L}_{\text{Gen}} = -\mathbb{E}_{x,m}[\log D_\xi(\hat{x})], \quad (3)$$

$$\mathcal{L}_{\text{Adv}} = \text{sg}_\theta(\mathcal{L}_{\text{Dis}}) + \text{sg}_\xi(\mathcal{L}_{\text{Gen}}) \rightarrow \min_{\theta, \xi}, \quad (4)$$

where $\mathbb{E}[\cdot]$ represents the operation of taking the average in mini-batches, x is a sample from the dataset, m is the mask, $\hat{x} = \mathcal{F}_\theta(x', R_{\text{ref}})$ is the reconstructed field, sg_{var} represents stopping gradients with respect to var, and $\mathcal{L}_{\text{Adv}}$ is the joint adversarial loss to be optimized.

3. Combination of loss functions

In the final loss function, the MAE loss $\mathcal{L}_1 = \mathbb{E}_{x,m}[\|x - \hat{x}\|_1]$, the R_1 gradient penalty ($R_1 = \mathbb{E}_x[\|\nabla D_\xi(x)\|_2]$) [60], and the feature matching loss $\mathcal{L}_{\text{FML}}$ (a discriminator-based perceptual loss, [61]) are also included. Here, $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the L_1 and L_2 norms, respectively. The final loss function for the reconstruction system, i.e.,

$$\mathcal{L}_{\text{final}} = \alpha \mathcal{L}_1 + \beta \mathcal{L}_{\text{PL}} + \gamma \mathcal{L}_{\text{Adv}} + \kappa \mathcal{L}_{\text{FML}} + \lambda R_1, \quad (5)$$

is the weighted sum of the discussed losses, where $\mathcal{L}_1$ is employed to guide the generator network to converge its reconstructed field to the target field, $\mathcal{L}_{\text{PL}}$ is responsible for the supervised signal and consistency of the global structure, while $\mathcal{L}_{\text{Adv}}$ and $\mathcal{L}_{\text{FML}}$ are responsible for generating the natural-looking local details, and R_1 is used to stabilize the training process and prevent overfitting to a certain degree. Weights α , β , γ , κ , and λ control the importance of each loss function.

D. Implementation details

The proposed reconstruction framework is implemented in the open-source PyTorch framework [62]. All modules are initialized from scratch using random weights. The weights of the generator and discriminator networks have been updated using Adam optimizer [63] with fixed learning rates of 0.001 and 0.0001, respectively. All models are trained for 120 epochs with a batch size of 8, and the generator and discriminator states have been stored every two epochs to find the optimum point. The training is performed on an NVIDIA V100 GPU and takes approximately 72 h for each model. The optimized weights of loss functions in the reconstruction framework are listed as




FIG. 3. (a) Three different masks with a gap ratio of 39.06%; (b) reconstructed density fields corresponding to the three different masks; (c) original density field; (d) gradient fields corresponding to the reconstructed density fields; (e) original density gradient field.

follows: $\alpha = 5.0$, $\beta = 5.0$, $\gamma = 30.0$, $\kappa = 100.0$ and $\lambda = 0.001$. These hyperparameters are used for training all models, except those described in the loss ablation study. For example, to assess the impact of perceptual loss and feature matching loss on reconstruction performance, their weights are set to zero (i.e., $\gamma = 0.0$ or $\kappa = 0.0$) for some specific cases, while maintaining the weights of other loss functions.

III. PERFORMANCE OF RECONSTRUCTION FRAMEWORK

In this section, we adopt the complete velocity fields as additional physical constraints, and demonstrate the reconstruction performance of the proposed framework based on qualitative visualization, single-point statistics, power spectrum and two-point correlation functions. Furthermore, SDF, MDF, and RDF are used to represent the single-damaged form, the multiple-damaged form and the random-damaged form, respectively.

A. Qualitative visualization

Reconstructed density fields with different damaged forms are presented in Figs. 3(b) and 4(b), respectively, with gap ratios of 39.06% and 56.25%. The corresponding masks and ground truth are also provided in Figs. 3(a), 3(c) and 4(a), 4(c). In this paper, $\text{PCGAN}_Y^{X\%}$ represents the reconstruction framework PCGAN that contains $X\%$ information of reference fields as additional physical constraints; subscript Y represents the damaged form, where S, M, and R represent




FIG. 4. (a) Three different masks with a gap ratio of 56.25%; (b) reconstructed density fields corresponding to the three different masks; (c) original density field; (d) gradient fields corresponding to the reconstructed density fields; (e) original density gradient field.

the single-damaged, multiple-damaged and random-damaged forms, respectively. Obviously, the proposed framework exhibits excellent performance in reconstructing the flow information of damaged regions, particularly in terms of local details, even when applied to the single-damaged form with a gap ratio of 56.25% [Fig. 4(b₁)]. To the best of our knowledge, the proposed framework outperforms existing approaches in the literature on turbulent field reconstruction (especially in reconstructing the local details of flow structures), including traditional linear data-driven methods and some recently proposed deep learning-based frameworks [14, 18, 19]. Comparing the reconstruction performance of the three damaged forms, the single-damaged form [Figs. 3(b₁) and 4(b₁)] exhibits the worst performance in reconstructing the density field, while the random-damaged form yields the best results [Figs. 3(b₃) and 4(b₃)]. This observation should be attributed to the fact that the reconstruction of the flow field within damaged regions is highly dependent on the physical constraints obtained from the boundaries of damaged regions, and the physical constraints become weaker as the distance from the boundaries increases. As mentioned above, the reconstruction network tends to average a variety of multiple plausible solutions in damaged regions to minimize the loss function. In the case of the single-damaged form, the physical constraints offered by the boundaries of damaged regions are evidently weaker than those of the other two damaged forms. Therefore, the reconstructed field of the single-damaged form becomes increasingly blurred as the distance from the boundary increases. On the contrary, the difference between the reconstructed field and the ground truth is barely discernible for the random-damaged form [Figs. 3(b₃) and 4(b₃)].

TABLE I. Normalized MAE values of density and its gradient in damaged regions. $()_1$ and $()_2$ correspond to two different situations with gap ratios of 39.06% and 56.25%, respectively.

Case	(SDF) ₁	(MDF) ₁	(RDF) ₁	(SDF) ₂	(MDF) ₂	(RDF) ₂
MAE_d^*	0.351	0.279	0.024	0.363	0.294	0.028
MAE_g^*	0.434	0.432	0.107	0.463	0.448	0.118

Figures 3(d) and 4(d) depict the gradient ($\nabla \rho = \partial \rho / \partial x + \partial \rho / \partial y$) fields corresponding to the reconstructed and the original density fields with gap ratios of 39.06% and 56.25%, respectively. The compressible flow field considered is characterized by pervasive shocklet structures [Fig. 1(a)], which causes obvious discontinuities in density. As such, pronounced discontinuous structures are discernible from the density gradient contours [Figs. 3(d) and 4(d)]. The blurring of the reconstructed density field would lead to a weakening of discontinuous structures in the gradient field. Here should be highlighted that checking the gradient of reconstructed fields is a very strict assessment due to the error accumulation from the reconstructed variables [46]. Obviously, compared to the other two damaged forms, the discontinuous structures in the damaged region corresponding to the single-damaged form are slightly weakened [Figs. 3(d₁) and 4(d₁)], especially in the center area of the damaged region. In addition, the density gradient fields corresponding to the random-damaged form [Figs. 3(d₃) and 4(d₃)] are almost the same as the ground truth [Figs. 3(e) and 4(e)].

B. Single-point statistics

Table I presents the normalized MAE values of the reconstructed fields with different damaged forms, which serves as an indicator of reconstruction performance. Notably, the MAE values of density and its gradient are computed solely in damaged regions, and their definitions are

$$\text{MAE}_d = \frac{1}{N_{\text{sam}} N_{\text{gap}}} \left[\sum_{k=1}^{N_{\text{sam}}} \sum_{l=1}^{N_{\text{gap}}} |\rho_R^k(l) - \rho_{GT}^k(l)| \right], \quad (6)$$

$$\text{MAE}_g = \frac{1}{N_{\text{sam}} N_{\text{gap}}} \left[\sum_{k=1}^{N_{\text{sam}}} \sum_{l=1}^{N_{\text{gap}}} |\nabla \rho_R^k(l) - \nabla \rho_{GT}^k(l)| \right], \quad (7)$$

where $\rho^k(l)$ and $\nabla \rho^k(l)$ represent values of the density and its gradient on the snapshot k and the element l , respectively; while the subscripts R and GT represent the reconstructed and the original fields, respectively. N_{gap} is the number of elements in damaged regions, while N_{sam} is the number of samples. We normalize the MAE values using the root mean-square values of the density and its gradient, i.e., $\text{MAE}_d^* = \text{MAE}_d / \rho_{\text{rms},GT}$ and $\text{MAE}_g^* = \text{MAE}_g / \nabla \rho_{\text{rms},GT}$. Consistent with the instantaneous contours depicted in Figs. 3(b) and 4(b), the MAE_d^* values for the single-damaged form (approximately 0.351 and 0.363 with gap ratios of 39.06% and 56.25%, respectively) are observed to be larger than those for the other two damaged forms. For example, the MAE_d^* values for the random-damaged form (approximately 0.024 and 0.028, respectively) are one order of magnitude smaller than those for the single-damaged form. Similarly to the performance of MAE_d^* , the values of MAE_g^* corresponding to the random-damaged form (approximately 0.107 and 0.118 with gap ratios of 39.06% and 56.25%, respectively) are much smaller than those for the other two damaged forms (Table I).




FIG. 5. $\text{MAE}_{\text{elem},d}^*$ contour for the single-damaged form with gap ratios of (a) 39.06% and (b) 56.25%.

Similarly, we can also define the mean absolute error across elements (MAE_{elem}) to examine the distribution of discrepancies between the reconstructed and the original fields. The definitions of MAE_{elem} for the density and its gradient are

$$\text{MAE}_{\text{elem},d} = \frac{1}{N_{\text{sam}}} \left[\sum_{k=1}^{N_{\text{sam}}} |\rho_R^k(l) - \rho_{GT}^k(l)| \right], \quad (8)$$

$$\text{MAE}_{\text{elem},g} = \frac{1}{N_{\text{sam}}} \left[\sum_{k=1}^{N_{\text{sam}}} |\nabla \rho_R^k(l) - \nabla \rho_{GT}^k(l)| \right]. \quad (9)$$

We normalize the $\text{MAE}_{\text{elem},d}$ and $\text{MAE}_{\text{elem},g}$ values using the root mean-square values of density and its gradient, i.e., $\text{MAE}_{\text{elem},d}^* = \text{MAE}_{\text{elem},d} / \rho_{\text{rms},GT}$ and $\text{MAE}_{\text{elem},g}^* = \text{MAE}_{\text{elem},g} / \nabla \rho_{\text{rms},GT}$. As an example, Fig. 5 illustrates the $\text{MAE}_{\text{elem},d}^*$ contour for the single-damaged form with gap ratios of 39.06% and 56.25%, respectively. The $\text{MAE}_{\text{elem},d}^*$ values are close to zero in the unmasked region, and the $\text{MAE}_{\text{elem},d}^*$ values in the damaged region are closely related to the distance from the boundary. The closer to the center of the damaged region, the larger the $\text{MAE}_{\text{elem},d}^*$ values. However, Fig. 6 shows the $\text{MAE}_{\text{elem},g}^*$ contour for the single-damaged form with gap ratios of 39.06% and 56.25%, respectively. Unlike $\text{MAE}_{\text{elem},d}^*$, the values of $\text{MAE}_{\text{elem},g}^*$ in the damaged region do not exhibit an evident spatial dependence on the distance from the boundary. Specifically, the values of $\text{MAE}_{\text{elem},g}^*$ remain relatively constant within the core of the damaged region. This observation also




FIG. 6. $\text{MAE}_{\text{elem},g}^*$ contour for the single-damaged form with gap ratios of (a) 39.06% and (b) 56.25%.




FIG. 7. (a), (b) Power spectra of the reconstructed fields with various damaged forms, as well as that of the original and the random-damaged fields; (c), (d) power spectrum ratio of the reconstructed to the original fields [$E_R(k)/E_{GT}(k)$].

suggests that the $MAE_{\text{elem},d}^*$ values follow an approximately linear trend with respect to the distance from the boundary.

C. Power spectrum and spatial correlation

Figures 7(a) and 7(b) depict the power spectra of the reconstructed fields for various damaged forms under two distinct gap ratios, as well as the power spectra of the original and the random-damaged fields. Differently from single-point statistics, the power spectrum represents the information of the whole field, especially the fluctuation information at different wave numbers. It is observed from Figs. 7(a) and 7(b) that the power spectrum of the random-damaged fields is far from that of the original fields. Severe oscillations can be observed in the high-wave-number region. However, the results demonstrate that the power spectra of the reconstructed fields with different damaged forms closely resemble that of the original fields, with only minor deviations observed in the high-wave-number region. In particular, Figs. 7(c) and 7(d) present the power spectrum ratio of the reconstructed to the original fields [$E_R(k)/E_{GT}(k)$] at different wave numbers. Within the wave-number range $k \leq 170$, the power spectrum ratio is approximately equal to 1.0. These findings indicate that, in a statistical sense, the present reconstruction framework can effectively reproduce the flow structure at different scales of the original density fields. In a convolutional network, the convolution operation can extract scale information from flow fields, and it is closely related to the convolutional kernel size. The convolutional kernel condenses the information from the overlapping regions of the flow field into a single value. Simultaneously, the resolution of the feature map undergoes successive reductions through the application of max-pooling. As a result, various network modules encompass different scale information of the flow field, with subsequent modules incorporating larger-scale information. In the proposed framework, the generator network utilizes a convolutional kernel of size 3×3 . Consequently, the convolutional kernel has the capability to




FIG. 8. Two-point correlation functions of density in the reconstructed and the original fields with different damaged forms. (a) Gap Ratio = 39.06%; (b) Gap Ratio = 56.25%.

resolve turbulence structures with a maximum wave number of $\kappa = 512/3 \approx 170.67$, corresponding to the minimum length scale $l_{min} = 2\pi/\kappa \approx 0.0117\pi$. Turbulence structures with length scales smaller than the minimum length scale l_{min} are eliminated by the convolution operation. We also compare the two-point correlation function of density in the reconstructed and the original fields, which is defined as

$$R(\delta x) = \mathbb{E}_f\{[\rho(\mathbf{x}) - \langle \rho \rangle][\rho(\mathbf{x} + \delta x \mathbf{e}_1) - \langle \rho \rangle]\}, \quad (10)$$

where $\mathbb{E}_f[\cdot]$ represents the operation of taking average in the testing datasets, δx represents the separation along the x direction, $\mathbf{e}_1$ represents the unit vector and $\langle \cdot \rangle$ represents the average operation. Since the two-point correlation at zero separation ($\delta x = 0$) is the variance of density, the variance of density is thus utilized for normalization. Figure 8 depicts the normalized two-point correlation function $[R(\delta x)/\rho_{rms}^2]$ of density in the original and the reconstructed fields with different damaged forms. $R(\delta x)/\rho_{rms}^2$ in the original field gradually decreases as δx increases from 0.0 to π , and attains a minimum value at $\delta x = \pi$ due to the application of periodic boundary conditions in all three directions. As shown in Fig. 8, although the two-point correlation functions of the reconstructed and the original fields do not perfectly overlap, their deviation can only be observed when δx is large enough, especially for cases with a gap ratio of 39.26% [Fig. 8(a)]. Furthermore, consistent with previous results, $R(\delta x)/\rho_{rms}^2$ corresponding to the reconstructed field with the single-damaged form has the largest deviation from that of the original field, while the deviation of $R(\delta x)/\rho_{rms}^2$ between the reconstructed and the original fields with the random-damaged form is almost negligible, even for the gap ratio of 56.25% [Fig. 8(b)].

IV. PERFORMANCE OF DCGAN AND ABLATION STUDY

The advent of generative adversarial networks in 2014 marked a significant milestone in the field of computer vision. However, the deep convolutional generative adversarial network (DCGAN) was not applied to the reconstruction of turbulent flow fields until year 2019 due to its inherent shortcomings, such as mode collapse, difficulty in training, and lacking details in the generated field [56,64]. Since then, continuous efforts have been made to enhance DCGAN performance by incorporating advanced modules and optimizing loss functions. As illustrated in Sec. III, we proposed a state-of-the-art reconstruction framework (PCGAN), which achieves excellent performance in reconstructing multiscale turbulent fields. In this section, we will validate the efficacy of DCGAN in the current reconstruction problem and evaluate the role of different network modules and loss functions in our proposed framework.




FIG. 9. (a) Framework of generator network for flow reconstruction, which is based on encoder and decoder networks; (b), (c) specific implementation of the encoder and decoder networks; (d) specific implementation of the global GAN discriminator network.

A. Performance of DCGAN

In Fig. 9, we have replicated the DCGAN framework proposed by [19], with certain modifications to accommodate the disparate resolution of flow fields and the inclusion of two complete fields of velocity components as additional physical constraints. As shown in Fig. 9, DCGAN comprises a generator network and a discriminator network. The reconstructing network $\mathcal{F}_\theta(\cdot)$ processes the masked field x' and the reference fields R_{ref} as input and produces a reconstructed field $\hat{x} = \mathcal{F}_\theta(x', R_{\text{ref}})$ [Fig. 9(a)]. The DCGAN generator network comprises an encoder network and a decoder network. Specific implementations of the encoder and decoder networks are shown in Figs. 9(b) and 9(c), respectively. The encoder and decoder networks each comprise seven convolutional modules, consisting of convolutional or deconvolutional operations, a BatchNorm operation, and a LeakyReLU activation function. Finally, the generator network $\mathcal{F}_\theta(\cdot)$ ends with




FIG. 10. (a) Random-damaged mask with a gap ratio of 56.25%; (b) reconstructed density field with the random-damaged form by DCGAN; (c) original density field.

a convolutional layer with a single feature map and a Sigmoid activation function. However, the discriminator network [Fig. 9(d)] is the same as that utilized in the present study, except for the last two modules [Fig. 2(c)]. Unlike the PatchGAN discriminator, it evaluates the genuineness of the entire field rather than discriminating between local patches. To this end, flatten and fully connected modules have been sequentially integrated into the end of the discriminator network.

In the DCGAN framework, the loss function for the generator network $\mathcal{L}_{\text{Gen}}$ is defined as

$$\mathcal{L}_{\text{Gen}} = (1 - \lambda_{\text{adv}})\mathcal{L}_2 + \lambda_{\text{adv}}\mathcal{L}_{\text{adv}}, \quad (11)$$

where $\mathcal{L}_2 = \mathbb{E}[\|x - \hat{x}\|_2]$ is the MSE between the reconstructed and the original flow fields, λ_{adv} is called the adversarial ratio, and $\mathcal{L}_{\text{adv}} (= \mathbb{E}[\log D_{\xi}(\hat{x})])$ is the adversarial loss. Meanwhile, the loss function for the discriminator network $\mathcal{L}_{\text{Dis}}$ is defined as

$$\mathcal{L}_{\text{Dis}} = -\mathbb{E}[\log D_{\xi}(x)] - \mathbb{E}[\log(1 - D_{\xi}(\hat{x}))]. \quad (12)$$

The generator and discriminator networks are trained together with the Adam optimizer, where the learning rate of the generator network is twice that of the discriminator network. To improve the stability of training, a staircase-decay schedule is adopted for the learning rate. It decays at a rate of 0.5 every 50 epochs. A batch size of 128 and an initial learning rate of 0.001 are utilized for the generator network, while the adversarial ratio is set to 0.001.

Figure 10 presents the instantaneous contour of the DCGAN-reconstructed density field, along with the corresponding mask and the original field for comparison. To limit the length of the article, here we just show results corresponding to the random-damaged form with a gap ratio of 56.25%, because this type of damaged form is more common than the other two forms in practical applications. The results indicate that the reconstruction performance of DCGAN is weaker than that of PCGAN [Fig. 4(b₃)]. The reconstructed field by DCGAN only reproduces large-scale structures in the turbulent field, even with a random-damaged form. Obviously, DCGAN fails to reconstruct small-scale structures and flow field details. Figure 11 presents the gradient field based on the DCGAN-reconstructed density field. Similarly, the apparent discontinuities in the original gradient field [Fig. 11(c)] are difficult to observe from the reconstructed gradient field [Fig. 11(b)]. This limitation of DCGAN can be attributed to the utilization of the global GAN discriminator and the conventional element-wise loss function (MSE). Table II compares the normalized MAE values of the density and its gradient corresponding to the two different frameworks. Consistent with results of the instantaneous contours, MAE_d^* and MAE_g^* corresponding to DCGAN ($\text{MAE}_d^* = 0.246$, $\text{MAE}_g^* = 0.810$) are almost one order of magnitude larger than those corresponding to PCGAN ($\text{MAE}_d^* = 0.028$, $\text{MAE}_g^* = 0.118$).

Figure 12 compares the reconstruction performance of two frameworks, DCGAN and PCGAN, based on the power spectrum function and the two-point correlation function. As shown




FIG. 11. (a) Random-damaged mask with a gap ratio of 56.25%; (b) gradient field based on the reconstructed density field by DCGAN with the random-damaged form; (c) original density gradient field.

in Fig. 12(a), DCGAN is only capable of reproducing the flow field at low wave numbers, with the power spectrum of the reconstructed field in this range almost overlapping that of the original field. However, as the wave number increases and exceeds 10, the power spectra of the reconstructed and the original fields gradually deviate from each other. However, the power spectrum of the reconstructed field based on PCGAN consistently overlaps that of the original field within the range of wave number $k \leq 170$. In addition, as shown in Fig. 12(b), although the reconstructed field of DCGAN only reproduces the large-scale structure of the original field, the two-point correlation function of the reconstructed field is in good agreement with that of the original field. This also suggests that the two-point correlation function exhibits a larger sensitivity to the large-scale structure in the flow field, while being less suitable for assessing the reconstruction of flow details, compared with the power spectrum function.

The underperformance of the DCGAN model utilized in Refs. [18,19] in the current problem primarily stems from its utilization of only MSE and adversarial loss in the loss function, without considering additional physical constraints. The existing literature has demonstrated that incorporating appropriate physical constraints into the loss function can significantly enhance the network's reconstruction performance [46,47,65]. Furthermore, employing a perceptual loss function in networks ensures accurate reconstruction of the flow field in unmasked regions, while the adoption of a PatchGAN discriminator network facilitates the reconstruction of small-scale structures within turbulence fields. We will delve into these topics extensively in subsequent sections.

B. Ablation study

In this section, we will mainly discuss the influence of perceptual loss, feature marching loss, and physical constraints from reference fields on the reconstruction performance in the proposed reconstruction framework. The effect is evaluated by comparing the reconstruction performance after weakening the corresponding loss functions or physical constraints. Table III shows the abbreviations of the reduced frameworks (PCGAN₁–PCGAN₃). Among them, PCGAN₁ represents the reconstruction framework that removes both the feature matching loss and the perceptual loss;

TABLE II. Normalized MAE values of density and its gradient in damaged regions. $()_{DC}$ and $()_{PC}$ correspond to results of different frameworks, DCGAN and PCGAN, respectively.

$(MAE_d^*)_{DC}$	$(MAE_g^*)_{DC}$	$(MAE_d^*)_{PC}$	$(MAE_g^*)_{PC}$
0.246	0.810	0.028	0.118




FIG. 12. Comparison of reconstruction performance of DCGAN and PCGAN for the random-damaged form with a gap ratio of 56.25%, based on (a) the power spectrum function and (b) the two-point correlation function.

PCGAN₂ and PCGAN₃ are the reconstruction frameworks that remove the perceptual loss and the feature matching loss, respectively.

1. Impact of perceptual and feature-matching loss functions

As mentioned in Sec. II C, the proposed framework employs the combination of feature matching and perceptual loss functions. The perceptual loss function evaluates the distance between feature maps of the reconstructed and the original flow fields through a pretrained network, while the feature matching loss is the perceptual loss of the discriminator network. In this section, we discuss the impact of perceptual and feature matching loss functions on the reconstruction performance.

Figure 13(a) depicts the power spectrum and the two-point correlation functions of the reconstructed flow fields with the random-damaged form and a gap ratio of 56.26%. The power spectra of the reconstructed flow fields by the proposed framework and its reduced variants exhibit a high degree of similarity to that of the original flow field in the wave-number range $k \leq 100$. However, discernible differences between the power spectra of the reconstructed and the original fields emerge at the wave number of $k \approx 100, 120$, and 170 for PCGAN₁, PCGAN₂, and PCGAN₃, respectively. Furthermore, as illustrated in Fig. 13(b), while the two-point correlation function of the reconstructed flow fields obtained using PCGAN₁ exhibits a slight deviation from that of the original flow fields, the two-point correlation functions of the flow fields reconstructed by other frameworks show excellent agreement with that of the original flow fields.

Figure 14 displays the instantaneous contours of flow fields reconstructed by PCGAN and its reduced versions (PCGAN₁–PCGAN₃), along with the corresponding masked and original flow fields. It is observed that the instantaneous contours of the reconstructed fields generated by PCGAN, PCGAN₂, and PCGAN₃ exhibit negligible differences from the original field. In addition, the performance of PCGAN₁ is comparatively worse than that of other reduced frameworks

TABLE III. Abbreviations of the reduced frameworks after removing part of loss functions.

Abbreviation	Meaning
PCGAN ₁	Proposed framework without feature matching and perceptual loss functions
PCGAN ₂	Proposed framework without perceptual loss function
PCGAN ₃	Proposed framework without feature matching loss function




FIG. 13. Comparison of reconstruction performance of PCGAN and its reduced frameworks ($PCGAN_1$ – $PCGAN_3$) with the random-damaged form and a gap ratio of 56.25%, based on (a) the power spectrum function and (b) the two-point correlation function.

[Figs. 14(c₁)–14(c₃)] in the reconstruction of flow field details, but still better than that of DCGAN [Fig. 10(b)]. Table IV presents the normalized MAE of density and density gradient of the damaged regions in the reconstructed flow fields, and quantitatively assesses the impacts of feature matching and perceptual loss functions on the reconstruction performance. The results demonstrate that the removal of both the feature matching loss and the perceptual loss leads to an increase in the values of MAE_d^* and MAE_g^* from 0.028 and 0.119 to 0.080 and 0.361, respectively. When only the perceptual loss is removed, the values of MAE_d^* and MAE_g^* increase from 0.028 and 0.118 to 0.046 and 0.186, respectively. On the contrary, the removal of feature matching loss has little effect on the values of MAE_d^* and MAE_g^* . Evidently, the variations in MAE_d^* and MAE_g^* are consistent with the changes




FIG. 14. (a) Random-damaged mask with a gap ratio of 56.25%; (b) original density field (ground truth); (c) reconstructed density field based on the proposed reconstruction framework (PCGAN) and its reduced versions ($PCGAN_1$ – $PCGAN_3$) with the random-damaged form.

TABLE IV. Normalized MAE values of density and its gradient in damaged regions. PCGAN₁–PCGAN₃ are the reduced frameworks of PCGAN after removing part of the loss functions.

	PCGAN _{1R} ^{100%}	PCGAN _{2R} ^{100%}	PCGAN _{3R} ^{100%}	PCGAN _R ^{100%}
MAE _d *	0.079	0.046	0.031	0.028
MAE _g *	0.358	0.186	0.116	0.118

observed in the power spectra as depicted in Fig. 13(a). From the above analysis, it is clear that the perceptual loss function has a significant impact on the reconstruction performance, particularly in accurately reconstructing small-scale structures in turbulent fields.

2. Impact of physical constraints of reference fields

The essence of flow field reconstruction via deep learning approaches is leveraging the powerful ability of networks to handle nonlinearities to construct a surrogate model of physical laws governing flows based on incomplete flow fields. Missing information within damaged regions is expected to be inferred by utilizing the surrogate model and the physical information at the boundaries of damaged regions. However, turbulent flows are governed by a set of highly coupled physical quantities, and information obtained solely from the boundaries is inadequate to precisely reconstruct the flow fields in damaged regions. As a result, the deep learning network typically produces blurred reconstruction results by averaging multiple plausible solutions. Furthermore, it is challenging to incorporate the physical laws into the reconstruction framework for two-dimensional turbulent fields, owing to the absence of temporal information and the limited information concerning the third spatial dimension. To address this issue, in the proposed framework, we utilize the two complete or sparse fields of velocity components as additional physical constraints for the damaged density field. This approach offers a flexible approach to incorporate additional physical constraints into the reconstruction process. The completeness of the additional physical constraint will be determined based on practical considerations, such as the resolution of the measurement equipment. In this section, we examine the importance of physical constraints in reconstruction performance. We consider three scenarios, i.e., using complete velocity fields or 50% random sampling of the velocity fields as additional physical constraints, or without additional physical constraints. In Fig. 15, the power spectrum and two-point correlation functions of the reconstructed flow fields are presented with a gap ratio of 56.26%, and compared with those of the original and damaged fields. The results demonstrate that the power spectra of the reconstructed fields with different damaged forms closely resemble that of the original field, although the power spectra of the damaged fields are far from that of the original field. For the single-damaged and multiple-damaged forms [Figs. 15(a) and 15(b)], noticeable differences in magnitude are observed between the power spectra of the reconstructed and original fields at various wave numbers when no additional physical constraints are considered. The difference between the power spectrum is significantly reduced when 50% or complete velocity field information is employed as physical constraints. For the random-damaged form [Fig. 15(c)], the power spectra of the reconstructed and original flow fields are nearly indistinguishable in the wave-number range $k \leq 170$, regardless of whether additional physical constraints are employed. Figures 15(b) and 15(d) depict the two-point correlation functions of flow fields. For the single-damaged form [Fig. 15(b)], a notable discrepancy is observed between the two-point correlation functions of the reconstructed and original flow fields when the separation δx is large enough. However, utilizing the velocity fields as physical constraints notably mitigates the deviation between the two fields. The deviation of the two-point correlation functions between the reconstructed and original fields is significantly reduced for the multiple-damaged form [Fig. 15(d)], and is nearly negligible for the random-damaged form [Fig. 15(f)] throughout all separations δx . Figure 16 presents the instantaneous contours of the reconstructed flow fields, as well as the corresponding mask and original flow field, for cases where velocity fields are utilized as physical




FIG. 15. Comparison of reconstruction performance of PCGAN for different damaged forms with a gap ratio of 56.25% and additional physical constraints that contain flow information extracted randomly from the complete reference field at different proportions. (a), (c), (e) Power spectrum function; (b), (d), (f) two-point correlation function.

constraints and where they are not. A comparison of Figs. 16(d₁) and 16(a) reveals that, in the cases of single-damaged form, the absence of physical constraints results in a reconstructed flow field that significantly deviates from the original within the damaged region. However, this deviation can be effectively mitigated by incorporating the velocity fields as physical constraints. The reconstruction performance is closely related to the proportion of flow information contained in the velocity field.




FIG. 16. (a) Original density field (ground truth); (b) additional physical constraints that contain flow information extracted randomly from the complete reference field at different proportions; (c) single-damaged, multiple-damaged and random-damaged masks with a gap ratio of 56.25%; (d)–(f) reconstructed density fields with corresponding masks and additional physical constraints.

The more flow information included, the better the reconstruction performance. Similar observations can be obtained for multiple-damaged and random-damaged forms [Figs. 16(e) and 16(f)]. For the random-damaged form [Figs. 16(e₁)–16(e₃)], the difference between the reconstructed and original flow fields is minimal, making it challenging to distinguish the effect of utilizing additional physical constraints. Clearly, these observations are consistent with the results of the power spectrum and the two-point correlation function in Fig. 15. Table V presents values of MAE_d^* and MAE_g^* , providing a quantitative comparison of the influence of physical constraints on reconstruction performance. It is clear that the reconstruction performance gradually improves as more flow information is included in the additional physical constraints. For the cases of single-damaged form, incorporation of physical constraints leads to a significant reduction in the MAE_d^* value, from 0.606 to 0.363, representing an improvement of 40.1%. Furthermore, the value of MAE_g^* decreases from 0.727 to 0.463 (improvement of 36.3%). This trend is similarly observed in the cases of multiple-damaged and random-damaged forms, where both MAE_d^* and MAE_g^* are accordingly reduced. In addition, it can be found from the table that the reconstruction performance of the random-damaged form

TABLE V. Normalized MAE values of density and its gradient in damaged regions.

	PCGAN _S ^{0%}	PCGAN _S ^{50%}	PCGAN _S ^{100%}
MAE _d [*]	0.606	0.481	0.363
MAE _g [*]	0.727	0.546	0.463
	PCGAN _M ^{0%}	PCGAN _M ^{50%}	PCGAN _M ^{100%}
MAE _d [*]	0.480	0.349	0.294
MAE _g [*]	0.676	0.535	0.448
	PCGAN _R ^{0%}	PCGAN _R ^{50%}	PCGAN _R ^{100%}
MAE _d [*]	0.035	0.033	0.028
MAE _g [*]	0.170	0.151	0.118

is significantly better than that of the single-damaged and multiple-damaged forms, regardless of whether additional physical constraints are implanted. Based on the aforementioned analysis, it can be concluded that additional physical constraints exert a substantial influence on reconstruction performance, particularly in cases where the damaged regions are large in size.

V. DISCUSSION

Since 2019, a significant body of literature has focused on the application of deep learning approaches for the super-resolution reconstruction of turbulent fields, i.e., generating high-resolution flow fields from low-resolution counterparts [1,66,67]. However, unlike previous research, the present work aims to reconstruct damaged turbulent fields with high resolution. This objective entails reproducing small-scale turbulence structures within damaged regions while ensuring the overall consistency of the reconstructed fields with the original fields in terms of magnitudes and statistical laws. Specifically, we hope that the power spectrum function and the two-point correlation function of the reconstructed fields align with those of the original fields, as these two functions are commonly employed to characterize the statistical properties of turbulence. To achieve our stated objectives, we have employed several specialized designs for networks and loss functions. In this section, we primarily explore how the PatchGAN discriminator network and the perceptual loss function improve the reconstruction performance.

A. PatchGAN discriminator network

As mentioned in Sec. II B, the PatchGAN discriminator network is adopted in the proposed framework. The schematic diagrams of the global GAN and PatchGAN discriminator networks are illustrated in Fig. 17. At the end of the global GAN network [Fig. 17(a)], a fully connected module and the Sigmoid activation function are used to assess the authenticity of the input field as a whole. On the contrary, at the end of the PatchGAN network [Fig. 17(b)], the feature maps are transformed into a two-dimensional array by convolution operations, and the discriminator evaluates the authenticity of the input field within each patch region. The global GAN discriminator is widely used in generative adversarial networks for super-resolution reconstruction, especially when the reconstructed field is small in size and contains fewer flow details [1]. However, when dealing with high-resolution turbulent fields containing diverse vortex structures across different scales and lacking prominent coherent structures [68], the PatchGAN discriminator proves to be more advantageous in enhancing the detailed reconstruction of turbulent fields. As depicted in Fig. 17(b), patches in the input field are uniform in size and cover the entire field without overlapping each other. More importantly, the PatchGAN discriminator network can flexibly adjust the size of patches by adjusting the network structure, and judge the authenticity of flow fields in an appropriate local




FIG. 17. Schematic diagrams of (a) the global GAN discriminator network and (b) the PatchGAN discriminator network.

range. Because the spatial scales of flow structures are closely related to the Reynolds number, this property is particularly important for the turbulent field reconstruction at different Reynolds numbers. Furthermore, the utilization of the PatchGAN discriminator has been found to alleviate the difficulty of network training for high-resolution turbulent field reconstruction. It is worth noting that in practical applications, the selection of patch size needs to take into account factors such as the complexity of flow field in various flow systems, the desired reconstruction performance, and the training difficulty of the network.

B. Perceptual loss function

As mentioned in Sec. IV A, the conventional element-wise loss functions often lead to blurry outcomes for turbulent field reconstruction. Figure 10 suggests that the MSE loss function and the global GAN discriminator only ensure the reconstruction of large-scale turbulence structures. Specifically, the power spectra of the reconstructed and the original fields demonstrate good agreement in the low-wave-number range and significant difference in the high-wave-number range [Fig. 12(a)]. Consequently, we have adopted a hybrid loss function consisting of the MAE and perceptual loss functions in the proposed framework. As mentioned above, the perceptual loss evaluates the distance between feature maps of the reconstructed and the original fields at different stages of the pretrained network, thereby ensuring consistency across different scales of the flow fields. Figure 18(a) shows the schematic diagram of the pretrained network and the perceptual loss calculation method. The pretrained network adopts the ResNet50dilated module as encoder. The ResNet50dilated module includes five stages (stages 0 ~ 4), and the sizes of the feature map produced by the last four stages are (128, 128, 256), (64, 64, 512), (64, 64, 1024), and (64, 64, 2048), respectively [52]. The perceptual loss calculates the average of multiple MSEs ($MSE_1 \sim MSE_4$) between the feature maps generated by the last four stages. That is, unlike the element-wise loss functions, the perceptual loss leverages the pretrained network to map the flow field to different levels and compares the reconstructed and the original fields on the feature maps. Such a training approach can significantly enhance the network's training accuracy and efficiency. Additionally, perceptual loss can ensure accurate replication of the flow field in unmasked regions, providing correct physical constraints at boundaries for the reconstruction of flow fields in damaged regions.

To illustrate how the perceptual loss works on the flow field reconstruction at the network level, we present an input field along with its representative feature maps generated by the pretrained network (stage 1) in Figs. 18(b) and 18(c). Evidently, the feature maps depicted in Figs. 18(c₁)




FIG. 18. (a) Schematic diagram of the pretrained network and the calculation of perceptual loss; (b) instantaneous contour of the input field and (c) its representative feature maps generated in stage 1 of the pretrained network.

and $18(c_2)$ delineate the profile of the input field. The pretrained network employs specific rules to separate the input field into two distinct components based on the magnitude of physical quantity, in which Figs. $18(c_1)$ and $18(c_2)$ represent regions with lower and higher magnitudes, respectively. Furthermore, as depicted in Figs. $18(c_2)$ – $18(c_5)$, a portion of the feature maps generated in stage 1 comprises mappings of the input field at different spatial scales. Here, Figs. $18(c_2)$ – $18(c_5)$ exhibit the




FIG. 19. Effects of the amount of training data on the reconstruction performance with the random-damaged form and a gap ratio of 56.25%; (a) the power spectrum function and (b) the two-point correlation function.

mappings with their corresponding spatial scales successively refined. That is, the perceptual loss evaluates the difference of feature maps of the reconstructed and the original fields at the network level, including the flow field profile and the mappings at different spatial scales. Consequently, compared with conventional element-wise loss functions, the perceptual loss imposes more stringent and refined constraints, thereby facilitating accurate flow field reconstruction over a broader range of wave numbers.

C. Effects of training data amount

The success of flow field reconstruction based on deep learning methods relies heavily on the training data. It is evident that for different damaged forms, the amount of training data required to achieve a similar reconstruction performance is different; it is intricately linked to the extent of damage in the flow field and the correlation between the training data and the prediction data. Therefore, in practical applications, it is difficult to determine the amount of training data required to obtain a satisfactory reconstruction performance without trial and error. To demonstrate the influence of the amount of training data on reconstruction performance for a specific damaged form, we consider the random-damaged form with a gap ratio of 56.25%. We employed three distinct amounts of training data, namely 3120 slices, 6240 slices, and 12480 slices. Once the model training converged, we utilized the identical testing dataset to assess its reconstruction performance. Figure 19 illustrates the power spectrum and the two-point correlation function of the reconstructed and original flow fields. As shown in Fig. 19, the difference in reconstruction performance among models trained with different amounts of data is practically imperceptible. The power spectrum curves corresponding to different amounts of training data almost overlap within the range of $\kappa \leq 200$ [Fig. 19(a)]. Similarly, the curves of the two-point correlation function overlap each other [Fig. 19(b)]. Table VI presents a comparison of the normalized MAE values of the density and its

TABLE VI. Normalized MAE values of density and its gradient in damaged regions. PCGAN_R^{100%}(X) denotes the PCGAN_R^{100%} model trained using X slices from the training data.

	PCGAN _R ^{100%} (3120)	PCGAN _R ^{100%} (6240)	PCGAN _R ^{100%} (12480)
MAE _d [*]	0.039	0.030	0.028
MAE _g [*]	0.163	0.137	0.118




FIG. 20. (a) Original density field (ground truth); (b) original density field combined with Gaussian noise of different variances; (c) random-damaged masks with a gap ratio of 56.25%; (d), (e) reconstructed density fields with corresponding masks and additional physical constraints.

gradient. It is evident that there is a decrease in the normalized MAE values of the density and its gradient as the amount of training data increases. Taking MAE_d^* as an example, a reduction in the amount of training data by half and three-quarters results in an increase in MAE_d^* by 7.14% and 38.97%, respectively. Consequently, for the random-damaged form with a gap ratio of 56.25%, a satisfactory reconstruction performance can be achieved using only half the amount of training data (6240 slices).

D. Robustness against noisy input

In this section, we will explore the reconstruction performance of a trained network for noisy input. As an illustration, we consider the random-damaged form with a gap ratio of 56.25%, and introduce Gaussian noise to the testing dataset. Specifically, the Gaussian noise introduced has variances of 0.00001, 0.0001, and 0.001, respectively, while the average value is set to zero. Figure 20 presents the reconstructed fields of the trained network for input containing Gaussian noise. To investigate the influence of the reference field in this scenario, we employed two distinct frameworks, namely $\text{PCGAN}_R^{100\%}$ and $\text{PCGAN}_R^{0\%}$. As shown in Figs. 20(d) and 20(e), it is evident that noise has a significant impact on the reconstruction performance of the trained network. On the one hand, it can be observed that for inputs containing Gaussian noise, although the trained network reconstructs large-scale turbulence structures, it also produces many distorted small-scale structures in response to the Gaussian noise. On the other hand, for inputs containing Gaussian noise with the same variance, the reconstruction performance of $\text{PCGAN}_R^{100\%}$ is significantly better than that




FIG. 21. Power spectrum of the reconstructed fields corresponding to input containing Gaussian noise, and that of the original field and the noisy input; (a) $\text{PCGAN}_R^{100\%}$ results and (b) $\text{PCGAN}_R^{0\%}$ results.

of $\text{PCGAN}_R^{0\%}$. In other words, the existence of the reference field inhibits the generation of such distorted small-scale structures to a certain extent.

Figure 21 illustrates the power spectrum of the reconstructed field corresponding to the input containing Gaussian noise. Additionally, it also displays the power spectra of the original field and the noisy input. As shown in Fig. 21, as the variance of Gaussian noise increases, although the power spectrum of the noisy input aligns with that of the original field in the low-wave-number region, it progressively diverges in the high-wave-number region. Similarly, the difference between power spectra of the reconstructed and original fields intensifies in the high-wave-number region as the variance of noise increases. Furthermore, compared with the power spectrum of the reconstructed field corresponding to $\text{PCGAN}_R^{0\%}$, the power spectrum of the reconstructed field corresponding to $\text{PCGAN}_R^{100\%}$ is closer to that of the original field for noisy inputs of the same variance. In other words, a network with reference fields as physical constraints has a stronger inhibitory effect on the impact of noise. This observation is consistent with the results shown in Fig. 20.

VI. CONCLUSIONS

The reconstruction of incomplete information in flow fields is a pervasive challenge in numerous applications concerning turbulence, and such incompleteness may arise from measurement constraints in experiments or inadequate spatial resolution in simulations. The present study represents a departure from previous research on super-resolution reconstruction in two fundamental aspects. First, we aim to reconstruct the large-area damaged turbulent fields with high resolution. Second, the study employs a dataset of compressible isotropic turbulent flow, which does not have any clearly identifiable coherent structures. The present objective is to reproduce small-scale turbulence structures within damaged regions while ensuring the overall consistency of the reconstructed fields with the original fields in terms of magnitudes and statistical laws.

In this work, we propose a framework for the high-fidelity reconstruction of large-area damaged turbulent fields with high resolution based on a physically constrained generative adversarial network. To improve the reconstruction performance, special designs have been incorporated into the proposed framework. Specific designs involve leveraging complete or sparse reference fields as additional physical constraints, employing a ResNet-like network with an FFC module, and adopting a PatchGAN discriminator network. To train the network, we employ a combination of loss functions, comprising MAE, perceptual loss, and feature matching loss. In terms of the dataset, we have considered three different damaged forms, namely the single-damaged form, the

multiple-damaged form, and the random-damaged form, and evaluated two gap ratios of 39.06% and 56.25%.

The proposed framework demonstrates excellent reconstruction performance, particularly in scenarios involving the random-damaged form. The reconstructed fields are consistent with the original fields in terms of magnitudes, power spectrum and two-point correlation functions. To provide a comparative analysis, we have reproduced the DCGAN framework, which is widely used in super-resolution reconstruction tasks, and have assessed its performance in the current reconstruction problem. Our analysis indicates that DCGAN can only effectively reconstruct large-scale structures and tends to overlook small-scale structures due to utilizing the global GAN discriminator network and the conventional element-wise loss function.

Extensive ablation studies are conducted to examine impacts of the loss functions, the utilization of reference fields as physical constraints, and the PatchGAN discriminator network on reconstruction performance. Specific analyses and results are listed below:

(1) The perceptual loss evaluates the differences between feature maps of the reconstructed and the original fields at network level based on a pretrained network, thus ensuring consistency across diverse flow scales. This training method also improves the accuracy and efficiency of network training. Furthermore, perceptual loss ensures accurate replication of the flow field in unmasked regions, which in turn provides correct boundary physical constraints for the flow field reconstruction within damaged regions.

(2) For reconstruction tasks with two-dimensional turbulent fields, the absence of temporal information and the limited information concerning the third spatial dimension make it challenging to integrate physical laws into the reconstruction framework. The proposed framework offers a flexible approach to incorporate additional physical constraints into the reconstruction process, and the additional physical constraints have a great impact on improving reconstruction performance, particularly for scenarios involving a large damaged area. For example, for the single-damaged form with a gap ratio of 56.25%, the MAE statistics in the damaged region improve by 40.0%.

(3) The PatchGAN discriminator network is utilized in the proposed framework. When reconstructing high-resolution turbulent fields, particularly in scenarios involving multiscale structures and the absence of distinct coherent structures, the PatchGAN discriminator demonstrates distinct advantages in accurately reconstructing the flow field details. Furthermore, the patch size of the discriminator network can be flexibly adjusted according to the spatial scale of turbulence structures, which significantly improves the accuracy of reconstruction and the efficiency of network training.

(4) Noise has a significant impact on the reconstruction performance of the trained network. Taking the random-damaged form as an example, for inputs containing Gaussian noise, although the trained network accurately reconstructs large-scale turbulence structures, it also produces many distorted small-scale structures in response to the Gaussian noise. The present study reveals that a network with reference fields as physical constraints has a stronger inhibitory effect on the impact of noise. This also highlights the important role of physical constraints in flow field reconstruction.

Overall, we firmly believe that the reconstruction framework proposed in this work, and the deeper insights into the contribution of the perceptual loss function and the PatchGAN discriminator, will provide a new pathway toward the practical use of neural networks for high-fidelity reconstruction of multiscale turbulent fields. Furthermore, the proposed framework can be easily extended to other applications, such as oceanographic and meteorological observations.

The code used in this study is available from the corresponding authors upon reasonable request.

ACKNOWLEDGMENTS

L.F. acknowledges the fund from the Research Grants Council (RGC) of the Government of Hong Kong Special Administrative Region (HKSAR) with RGC/ECS Project (Grant No. 26200222), RGC/GRF Project (Grant No. 16201023), and RGC/STG Project (Grant No. STG2/E-605/23-N), the fund from Guangdong Basic and Applied Basic Research Foundation, the fund

from Guangdong Province Science and Technology Plan Project (Grant No. 2023A0505030005), the fund from Center for Ocean Research in Hong Kong and Macau, a joint research center between Laoshan Laboratory and HKUST, and the fund from the Project of Hetao Shenzhen–Hong Kong Science and Technology Innovation Cooperation Zone (Grant No. HZQB-KCZYB-2020083). X.L. acknowledges the fund from the National Natural Science Foundation of China (Grant No. 62306254). Q.Z. acknowledges the fund from the National Natural Science Foundation of China (Grant No. 11902139).

The authors report no conflict of interest.

APPENDIX: RECONSTRUCTION PERFORMANCE WITH MASKED REFERENCE FIELDS

In the present study, we propose to utilize the information of other quantities as physical constraints to reconstruct the turbulence field within damaged regions. In Sec. IV B, we discuss the reconstruction performance of the proposed framework under physical constraints, which contains




FIG. 22. (a) Original density field (ground truth); (b) additional physical constraints from reference fields, where figure b_2 as an example denotes that reference fields have the same damaged form as the density input; (c) single-damaged, multiple-damaged and random-damaged masks with a gap ratio of 56.25%; (d)–(f) reconstructed density fields with corresponding masks and additional physical constraints.




FIG. 23. (a) Power spectra of the reconstructed fields with various damaged forms, as well as that of the original and the random-damaged fields; (b) two-point correlation functions of density in the reconstructed and the original fields.

flow information randomly extracted from the complete reference field at different proportions. In this appendix, we will further discuss the reconstruction performance of the proposed framework when the reference fields have the same damaged regions with the density input. Here, we consider different damaged forms with a gap ratio of 56.25%.

Figure 22 presents the instantaneous contours of the reconstructed fields, as well as the corresponding masks and the original field. To facilitate comparison, Fig. 22 also includes the reconstructed fields for cases where complete velocity fields are utilized as physical constraints and where they are not. Here, $(PCGAN_X^{100\%})_{MR}$ denotes the $PCGAN_X^{100\%}$ framework which employs masked reference fields as physical constraints. As shown in Fig. 22, the proposed framework effectively reconstructs the unmarked region, regardless of whether the reference fields are employed, owing to the utilization of the perceptual loss function. Furthermore, due to the absence of velocity fields as physical constraints in damaged regions, the reconstruction performance of $PCGAN_X^{0\%}$ and $(PCGAN_X^{100\%})_{MR}$ is notably inferior to that of $PCGAN_X^{100\%}$.

Figure 23 depicts the power spectrum and the two-point correlation functions of the reconstructed fields, and compared with that of the original field. As shown in Fig. 23(a), although the power spectra of the reconstructed and original fields exhibit a consistent trend within the range of $\kappa \leq 170$, for the single-damaged and multiple-damaged forms, a noticeable difference in magnitude is observed between the power spectra at various wave numbers. For the random-damaged form, the power spectrum curves corresponding to the reconstructed and original fields overlap each other within the range of $\kappa \leq 170$. Similarly, for the single-damaged and multiple-damaged forms, obvious deviations can be observed in the two-point correlation function curves between the reconstructed and original fields when the separation δx is large enough [Fig. 23(b)]. However, the curves overlap each other for the random-damaged form.

TABLE VII. Normalized MAE values of density and its gradient in damaged regions. $(PCGAN_X^{100\%})_{MR}$ denotes that the $PCGAN_X^{100\%}$ framework employs masked reference fields as physical constraints.

	$(PCGAN_S^{100\%})_{MR}$	$(PCGAN_M^{100\%})_{MR}$	$(PCGAN_R^{100\%})_{MR}$
MAE_d^*	0.615	0.477	0.034
MAE_g^*	0.723	0.689	0.164

Table VII presents the values of MAE_d^* and MAE_g^* to quantitatively evaluate the network's reconstruction performance with the marked reference fields. Compared with Table V, when the reference fields have the same damaged regions with the density input, the values of MAE_d^* and MAE_g^* are close to those corresponding to the framework without considering physical constraints. It suggests that its reconstruction performance is equivalent to the case where no reference fields are employed.

-
- [1] A. Güemes, C. Sanmiguel Vila, and S. Discetti, Super-resolution generative adversarial networks of randomly-seeded fields, *Nat. Mach. Intell.* **4**, 1165 (2022).
 - [2] L. Fu, S. Bose, and P. Moin, Prediction of aerothermal characteristics of a generic hypersonic inlet flow, *Theor. Comput. Fluid Dyn.* **36**, 345 (2022).
 - [3] K. Fukami, R. Maulik, N. Ramachandra, K. Fukagata, and K. Taira, Global field reconstruction from sparse sensors with Voronoi tessellation-assisted deep learning, *Nat. Mach. Intell.* **3**, 945 (2021).
 - [4] S. L. Brunton, B. R. Noack, and P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid Mech.* **52**, 477 (2020).
 - [5] X. Wen, Z. Li, D. Peng, W. Zhou, and Y. Liu, Missing data recovery using data fusion of incomplete complementary data sets: A particle image velocimetry application, *Phys. Fluids* **31**, 025105 (2019).
 - [6] C. Wang, Q. Gao, H. Wang, R. Wei, T. Li, and J. Wang, Divergence-free smoothing for volumetric PIV data, *Exp. Fluids* **57**, 1 (2016).
 - [7] C. Irrgang, N. Boers, M. Sonnewald, E. A. Barnes, C. Kadow, J. Staneva, and J. Saynisch-Wagner, Towards neural earth system modelling by integrating artificial intelligence in earth system science, *Nat. Mach. Intell.* **3**, 667 (2021).
 - [8] S. Yu and J. Ma, Deep learning for geophysics: Current and future trends, *Rev. Geophys.* **59**, e2021RG000742 (2021).
 - [9] C. Kadow, D. M. Hall, and U. Ulbrich, Artificial intelligence reconstructs missing climate information, *Nature Geosci.* **13**, 408 (2020).
 - [10] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge Monographs on Mechanics (Cambridge University Press, Cambridge, UK, 2012).
 - [11] G. Berkooz, P. Holmes, and J. L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.* **25**, 539 (1993).
 - [12] P. J. Schmid, Dynamic mode decomposition and its variants, *Annu. Rev. Fluid Mech.* **54**, 225 (2022).
 - [13] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
 - [14] S. G. Raben, J. J. Charonko, and P. P. Vlachos, Adaptive gappy proper orthogonal decomposition for particle image velocimetry data reconstruction, *Meas. Sci. Technol.* **23**, 025303 (2012).
 - [15] N. Murray and J. Seiner, The effects of gappy POD on higher-order turbulence quantities, in *Proceedings of the 46th AIAA Aerospace Sciences Meeting and Exhibit* (AIAA, Reston, VA, 2008), p. 241.
 - [16] H. Gunes, S. Sirisup, and G. E. Karniadakis, Gappy data: To Krig or not to Krig? *J. Comput. Phys.* **212**, 358 (2006).
 - [17] D. Venturi and G. E. Karniadakis, Gappy data and reconstruction procedures for flow past a cylinder, *J. Fluid Mech.* **519**, 315 (1999).
 - [18] T. Li, M. Buzzicotti, L. Biferale, F. Bonaccorso, S. Chen, and M. Wan, Multi-scale reconstruction of turbulent rotating flows with proper orthogonal decomposition and generative adversarial networks, *J. Fluid Mech.* **971**, A3 (2023).
 - [19] M. Buzzicotti, F. Bonaccorso, P. C. Di Leoni, and L. Biferale, Reconstruction of turbulent data with deep generative models for semantic inpainting from TURB-Rot database, *Phys. Rev. Fluids* **6**, 050503 (2021).
 - [20] E. Jagodinski, X. Zhu, and S. Verma, Inverse identification of dynamically important regions in turbulent flows using three-dimensional convolutional neural networks, *Phys. Rev. Fluids* **8**, 094605 (2023).

- [21] J. Kim and C. Lee, Prediction of turbulent heat transfer using convolutional neural networks, *J. Fluid Mech.* **882**, A18 (2020).
- [22] S. Bhatnagar, Y. Afshar, S. Pan, K. Duraisamy, and S. Kaushik, Prediction of aerodynamic flow fields using convolutional neural networks, *Comput. Mech.* **64**, 525 (2019).
- [23] K. Stengel, A. Glaws, D. Hettinger, and R. N. King, Adversarial super-resolution of climatological wind and solar data, *Proc. Natl. Acad. Sci. USA* **117**, 16805 (2020).
- [24] T. Bolton and L. Zanna, Applications of deep learning to ocean data inference and subgrid parameterization, *J. Adv. Model. Earth Syst.* **11**, 376 (2019).
- [25] A. Asensio Ramos and C. J. Díaz Baso, Stokes inversion based on convolutional neural networks, *Astron. Astrophys.* **626**, A102 (2019).
- [26] C. Schaefer, M. Geiger, T. Kuntzer, and J.-P. Kneib, Deep convolutional neural networks as strong gravitational lens detectors, *Astron. Astrophys.* **611**, A2 (2018).
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Commun. ACM* **60**, 84 (2017).
- [28] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [29] C. Drygala, B. Winhart, F. di Mare, and H. Gottschalk, Generative modeling of turbulence, *Phys. Fluids* **34**, 035114 (2022).
- [30] M. Z. Yousif, L. Yu, and H.-C. Lim, High-fidelity reconstruction of turbulent flow from spatially limited data using enhanced super-resolution generative adversarial network, *Phys. Fluids* **33**, 125119 (2021).
- [31] H. Gao, L. Sun, and J.-X. Wang, Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels, *Phys. Fluids* **33**, 073603 (2021).
- [32] L. Sun and J.-X. Wang, Physics-constrained Bayesian neural network for fluid flow reconstruction with sparse and noisy data, *Theor. Appl. Mech. Lett.* **10**, 161 (2020).
- [33] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial networks, *Commun. ACM* **63**, 139 (2020).
- [34] M. Z. Yousif, L. Yu, S. Hoyas, R. Vinuesa, and H. Lim, A deep-learning approach for reconstructing 3D turbulent flows from 2D observation data, *Sci. Rep.* **13**, 2529 (2023).
- [35] A. Subramaniam, M. L. Wong, R. D. Borker, S. Nimmagadda, and S. K. Lele, Turbulence enrichment using physics-informed generative adversarial networks, *arXiv:2003.01907*.
- [36] H. Kim, J. Kim, S. Won, and C. Lee, Unsupervised deep learning for super-resolution reconstruction of turbulence, *J. Fluid Mech.* **910**, A29 (2021).
- [37] Q. Zheng and M. M. Alam, Evolution of the wake of three inline square prisms, *Phys. Rev. Fluids* **4**, 104701 (2019).
- [38] H. Choi, W.-P. Jeon, and J. Kim, Control of flow over a bluff body, *Annu. Rev. Fluid Mech.* **40**, 113 (2008).
- [39] C. H. Williamson, Vortex dynamics in the cylinder wake, *Annu. Rev. Fluid Mech.* **28**, 477 (1996).
- [40] C. Cheng, W. Shyy, and L. Fu, Streamwise inclination angle of wall-attached eddies in turbulent channel flows, *J. Fluid Mech.* **946**, A49 (2022).
- [41] C. Cheng and L. Fu, Consistency between the attached-eddy model and the inner–outer interaction model: A study of streamwise wall-shear stress fluctuations in a turbulent channel flow, *J. Fluid Mech.* **942**, R9 (2022).
- [42] K. Fukami, K. Fukagata, and K. Taira, Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows, *J. Fluid Mech.* **909**, A9 (2021).
- [43] D. Buaria, M. P. Clay, K. R. Sreenivasan, and P. K. Yeung, Small-scale isotropy and ramp-cliff structures in scalar turbulence, *Phys. Rev. Lett.* **126**, 034504 (2021).
- [44] S. B. Pope, *Turbulent Flows* (Cambridge University Press, Cambridge, UK, 2000).
- [45] M. Z. G. Yousif, L. Yu, and H. Lim, Physics-guided deep learning for generating turbulent inflow conditions, *J. Fluid Mech.* **936**, A21 (2022).
- [46] 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).
- [47] P. Sharma, W. T. Chung, B. Akoush, and M. Ihme, A review of physics-informed machine learning in fluid mechanics, *Energies* **16**, 2343 (2023).

- [48] Q. Zheng, Y. Yang, J. Wang, and S. Chen, Enstrophy production and flow topology in compressible isotropic turbulence with vibrational nonequilibrium, *J. Fluid Mech.* **950**, A21 (2022).
- [49] Q. Zheng, J. Wang, M. M. Alam, B. R. Noack, H. Li, and S. Chen, Transfer of internal energy fluctuation in compressible isotropic turbulence with vibrational nonequilibrium, *J. Fluid Mech.* **919**, A26 (2021).
- [50] Q. Zheng, J. Wang, B. R. Noack, H. Li, M. Wan, and S. Chen, Vibrational relaxation in compressible isotropic turbulence with thermal nonequilibrium, *Phys. Rev. Fluids* **5**, 044602 (2020).
- [51] L. Chi, B. Jiang, and Y. Mu, Fast fourier convolution, *Adv. Neural Inf. Process. Syst.* **33**, 4479 (2020).
- [52] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2016), pp. 770–778.
- [53] D. Zhang, F. Huang, S. Liu, X. Wang, and Z. Jin, Swinfr: Revisiting the swinir with fast fourier convolution and improved training for image super-resolution, [arXiv:2208.11247](https://arxiv.org/abs/2208.11247).
- [54] R. Suvorov, E. Logacheva, A. Mashikhin, A. Remizova, A. Ashukha, A. Silvestrov, N. Kong, H. Goka, K. Park, and V. Lempitsky, Resolution-robust large mask inpainting with fourier convolutions, in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (IEEE, Piscataway, NJ, 2022), pp. 2149–2159.
- [55] U. Demir and G. Unal, Patch-based image inpainting with generative adversarial networks, [arXiv:1803.07422](https://arxiv.org/abs/1803.07422).
- [56] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, Image-to-image translation with conditional adversarial networks, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2017), pp. 1125–1134.
- [57] J. Johnson, A. Alahi, and L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution, in *Proceedings of the 14th European Conference on Computer Vision (ECCV)*, Part II 14 (Springer, Berlin, 2016), pp. 694–711.
- [58] B. Zhou, H. Zhao, X. Puig, T. Xiao, S. Fidler, A. Barriuso, and A. Torralba, Semantic understanding of scenes through the ade20k dataset, *Int. J. Comput. Vision* **127**, 302 (2019).
- [59] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, Scene parsing through ade20k dataset, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2017), pp. 633–641.
- [60] A. Ross and F. Doshi-Velez, Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients, in *Proceedings of the AAAI Conference on Artificial Intelligence* (AIAA, Reston, VA, 2018), Vol. 32.
- [61] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, High-resolution image synthesis and semantic manipulation with conditional gans, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2018), pp. 8798–8807.
- [62] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, Pytorch: An imperative style, high-performance deep learning library, *Adv. Neural Inf. Process. Syst.* **32** (2019).
- [63] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [64] V. Kushwaha, G. Nandi *et al.*, Study of prevention of mode collapse in generative adversarial network (GAN), in *Proceedings of the IEEE 4th Conference on Information and Communication Technology (CICT)* (IEEE, Piscataway, NJ, 2020), pp. 1–6.
- [65] M. Bode, M. Gauding, Z. Lian, D. Denker, M. Davidovic, K. Kleinheinz, J. Jitsev, and H. Pitsch, Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows, *Proc. Combust. Inst.* **38**, 2617 (2021).
- [66] K. Fukami, K. Fukagata, and K. Taira, Super-resolution reconstruction of turbulent flows with machine learning, *J. Fluid Mech.* **870**, 106 (2019).
- [67] Z. Deng, C. He, Y. Liu, and K. C. Kim, Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework, *Phys. Fluids* **31**, 125111 (2019).
- [68] L. Fu, M. Karp, S. T. Bose, P. Moin, and J. Urzay, Shock-induced heating and transition to turbulence in a hypersonic boundary layer, *J. Fluid Mech.* **909**, A8 (2021).