

Ensemble flow reconstruction in the atmospheric boundary layer from spatially limited measurements through latent diffusion models

Cite as: Phys. Fluids **35**, 126604 (2023); doi: 10.1063/5.0172559

Submitted: 16 August 2023 · Accepted: 16 November 2023 ·

Published Online: 11 December 2023



View Online



Export Citation



CrossMark

Alex Rybchuk,^{1,a)} Malik Hassanaly,¹ Nicholas Hamilton,¹ Paula Doubrawa,¹ Mitchell J. Fulton,² and Luis A. Martínez-Tossas¹

AFFILIATIONS

¹National Wind Technology Center, National Renewable Energy Laboratory, Golden, Colorado 80401, USA

²Department of Mechanical Engineering, University of Colorado Boulder, Boulder, Colorado 80309-0552, USA

^{a)}Author to whom correspondence should be addressed: alex.rybchuk@nrel.gov

ABSTRACT

Due to costs and practical constraints, field campaigns in the atmospheric boundary layer typically only measure a fraction of the atmospheric volume of interest. Machine learning techniques have previously successfully reconstructed unobserved regions of flow in canonical fluid mechanics problems and two-dimensional geophysical flows, but these techniques have not yet been demonstrated in the three-dimensional atmospheric boundary layer. Here, we conduct a numerical analogue of a field campaign with spatially limited measurements using large-eddy simulation. We pose flow reconstruction as an inpainting problem, and reconstruct realistic samples of turbulent, three-dimensional flow with the use of a latent diffusion model. The diffusion model generates physically plausible turbulent structures on larger spatial scales, even when input observations cover less than 1% of the volume. Through a combination of qualitative visualization and quantitative assessment, we demonstrate that the diffusion model generates meaningfully diverse samples when conditioned on just one observation. These samples successfully serve as initial conditions for a large-eddy simulation code. We find that diffusion models show promise and potential for other applications for other turbulent flow reconstruction problems.

© 2023 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0172559>

I. INTRODUCTION

Atmospheric field campaigns aim to characterize the complicated state of the atmosphere by deploying several measurement systems, often supplementing observations with modeling. Field campaigns in the atmospheric boundary layer (ABL, approximately the lowest 1 km) are critical for different research areas, such as wind energy,^{1,2} air quality,^{3,4} and wildland fires.⁵ While field campaigns in the ABL typically strive to measure as much of the lower atmosphere as possible, they measure only a small fraction of the atmosphere due to practical constraints and the costs associated with observation systems. Models can be used to fill the unmeasured areas between the spatially limited measurements so that atmospheric dynamics can be characterized across a wide range of scales. This combination of measurements and models enables what is referred to as flow reconstruction: turning spatially and temporally sparse flow information into a highly resolved field.

Machine learning has emerged as a powerful technique for turbulent flow reconstruction,^{6,7} joining the ranks of more traditional

approaches such as data assimilation^{8–12} and nudging.^{13,14} The flow reconstruction problem can be formulated in a variety of approaches, depending on the available data and choice of machine learning architecture. Perhaps the most common framing is that of a “super-resolution” problem (see Fukami *et al.*¹⁵ for a recent review). In super-resolution, low-resolution data are available, and data are upsampled to higher resolution through the use of an algorithm. Turbulence super-resolution has been carried out through the use of convolutional neural network architectures^{16–18} as well as generative adversarial network (GAN) architectures.^{19–24} In an alternative to super-resolution, turbulent flow reconstruction has been posed as an “inpainting” problem.²⁵ In this scenario, part of an image is masked, and an algorithm plausibly reconstructs the missing data. Others have specifically examined the problem of reconstruction given sparse measurements, often through the use of convolutional neural networks.^{26,27} Finally, instead of using machine learning architectures rooted in the field of computer vision, others have drawn inspiration from the intersection of partial

differential equations and machine learning. Flow reconstruction approaches in these categories employ architectures such as “physics-informed neural networks” (PINNs)^{28–30} and Deep Operator Networks.³¹ All in all, the aforementioned flow reconstruction techniques have been applied to a variety of turbulent flow environments, ranging from two-dimensional²⁰ and three-dimensional³² canonical fluid mechanics problems (e.g., flow behind a cylinder) to two-dimensional geophysical problems²⁶ (e.g., the state of the sea surface).

While flow reconstruction based on machine learning has been shown to be powerful, these techniques have not yet been applied to real-world, three-dimensional geophysical flows, such as an ABL. Real-world, ABL flow reconstruction comes with important challenges that must be accounted for. For example, many observation systems (e.g., Doppler lidar) measure line-of-sight velocity, which often serves as an accurate proxy for just one velocity component (e.g., streamwise velocity u) instead of measuring all three velocity components (u , v , and w). So far, only PINNs have demonstrated the ability to reconstruct all velocity components from measurements of just a scalar field.²⁸ While PINNs show great potential, they often struggle for multi-scale problems³³ such as high Re three-dimensional turbulence, and their success has not yet been demonstrated in an ABL. Another issue in reconstructing real-world flows is the sparseness of ABL measurements, as mentioned earlier. The atmosphere is highly chaotic and flow reconstruction from sparse measurements is typically an ill-posed problem, as many non-unique states could correspond to the same observation. As such, geoscientists often characterize its state probabilistically (e.g., in ensemble-based weather forecasting³⁴). However, most flow reconstruction studies have been practically deterministic (e.g., using GANs, which suffer from the “mode collapse” problem³⁵). Only two studies thus far, Gunderson *et al.*²⁷ using variational autoencoders (AEs) and Hassanaly *et al.*²³ using modified GANs, have reconstructed turbulent flow in a probabilistic manner, though it is not clear if their two-dimensional architectures can computationally scale to a three-dimensional, highly turbulent ABL.

Recently, a new neural architecture known as a diffusion model (DM)^{36,37} has achieved state-of-the-art status for generating high-resolution imagery,³⁸ showing promise for turbulent flow reconstruction in the ABL. A specific category of DM, known as a latent diffusion model (LDM),³⁹ is computationally efficient, and LDMs have been used to generate high-resolution two-dimensional imagery³⁹ and three-dimensional medical imagery.⁴⁰ DMs are inherently stochastic, and they have been shown to generate diverse, photorealistic imagery when a random seed is changed. Finally, while DMs have shown strong performance for the super-resolution problem,³⁹ they have also excelled at inpainting.^{41,42} We posit that the inpainting perspective is a natural approach to pose the problem of flow reconstruction from spatially limited measurements.

In this paper, we ask: given measurements in one small region of the atmosphere, can we estimate the instantaneous, unmeasured state of the ABL nearby through inpainting with an LDM? Here, we study this problem in the context of a synthetic field campaign that is conducted through large-eddy simulation (LES). In this work, our flow reconstruction strategy is to recreate instantaneous, volumetric flow fields. This strategy shares similarities with data assimilation approaches that are commonly applied for mesoscale and synoptic scale atmospheric reconstruction, namely, 4D variational data assimilation (4DVAR)⁴³ and the ensemble Kalman filter,⁴⁴ which produce plausible initial conditions for a dynamical solver given a time-history

of observations. Our investigation here addresses three key ABL reconstruction challenges in a synthetic environment: (1) probabilistic reconstruction, (2) dealing with spatially limited measurements, and (3) translating from scalar measurements to all three velocity components. For the time being, we omit two additional challenges of ABL reconstruction, namely, dealing with (1) noisy measurements and (2) a time-history of measurements. An optimal ABL reconstruction strategy would likely be able to account for all these challenges.

We find that LDMs can successfully reconstruct ABL flow given limited measurements, and we believe that LDMs can be applied for other turbulent flow reconstruction problems as well. In this paper, we demonstrate the following contributions:

- Through the use of LDMs, we can generate diverse three-dimensional turbulent flow fields. We characterize the quality of LDM samples through the use of qualitative visualization and quantitative assessment.
- We show that LDMs can reconstruct all three velocity components (u , v , w), even when given observations of almost exclusively u .
- While LDM studies in the field of computer vision have masked upwards of 75% of an image, we show that LDM reconstructions can be conditioned, even when minimal observations (<1% of the volume) are provided.
- We demonstrate that a LDM sample can successfully be used as an initial condition for an LES. To our knowledge, this is the first time that a machine learning sample has demonstrated this kind of compatibility with an LES code.

The rest of this manuscript is structured as follows. In Sec. II, we describe our synthetic field campaign, the configuration of the flow reconstruction problem, and the LES dataset that is used to train and test the LDM. In Sec. III, we provide details of the LDM architecture. In Sec. IV, we study the performance of the LDMs. Finally, in Sec. V, we conclude the paper and discuss potential future lines of inquiry.

II. DATA

In this work, we explore flow reconstruction in the context of the real-world Rotor Aerodynamics Aeroelastics and Wakes (RAAW) field campaign,² which seeks to thoroughly characterize the behavior of a single, utility-scale wind turbine with respect to the incoming atmospheric inflow, particularly on shorter timescales (1 s to 10 min). Here, we conduct a synthetic version of the field campaign, which is sometimes referred to as an “observing system simulation experiment.”

We simulate a thermally neutral atmospheric boundary layer using the LES code AMR-Wind.⁴⁵ AMR-Wind solves the incompressible Navier–Stokes equations using a finite-volume approach with second-order accuracy in space and time. Details of the spatial and temporal discretization can be found in Almgren *et al.*⁴⁶ The simulation is run with a Smagorinsky subgrid-scale model.⁴⁷ The domain is forced by $U = 10 \text{ m s}^{-1}$ geostrophic winds in the x -direction, as well as Coriolis forcing at a latitude of 90° . The LES domain is sized $(x, y, z) = (1920, 1920, 960) \text{ m}$ with $(nx, ny, nz) = (128, 128, 64)$ grid points, leading to 15 m resolution in all three directions. The simulation uses a fixed 0.5 s time step. The LES is initially run for 6 h to allow turbulence to spin up and become stationary. Afterward, we generate data for the training set by running the LES for 3.5 days of simulated time and saving the (u, v, w) fields everywhere in the domain every minute,

totaling 5040 samples. After the training set is generated, we generate a test dataset by simulating another 3.5 days after the end of the training dataset period. By using test data after the training period, we ensure there is no cross-contamination between datasets while also achieving good statistical agreement between the two datasets in our heights of interest, rotor span heights of 56.5–183.5 m (Sec. IV).

Synthetic measurements are generated by masking regions of the three-dimensional LES output, and in the context of this paper, we refer to unmasked regions as “observations.” We test the LDM on two sets of masks (Fig. 1). The “box mask” is commonly used in two-dimensional image inpainting problems.^{41,42} Here, we mask a cube at the center of the domain spanning (64, 64, 32) grid points, in order to qualitatively demonstrate flow reconstruction in a scenario with many observations. For the box mask, we use observations of all three components velocity. The “field campaign (FC) mask” hides data at locations specific to the RAAW field campaign instrumentation layout, observing data at 1525 of 1 048 576 LES grid cells. This mask could be modified for any other layout of field measurements. All instruments are aligned to point into the incoming wind, as they would be in the field. The turbine can be thought of as sitting at $(x, y) = (1800, 960)$ m. The horizontal scanning lidar reaches 1000 m upwind of the turbine, covering an azimuthal range of 18° at a height of 120 m. The vertical spinner lidar scan covers the rotor area and sits 120 m upwind of the turbine. The meteorological mast sits 360 m upwind of the turbine and measures in a vertical column between 0 and 180 m.

In this study, we take steps toward developing an algorithm that could be applied to real-world measurements, but we still make a number of simplifying assumptions that will be relaxed in future work.

- We assume that measurements are noise-free.
- We assume that both lidars directly measure the u -component of velocity and only this component of velocity. In practice, lidars measure line-of-sight velocity, which is then projected to other directions to obtain specific components of the velocity vector. As our lidars measure directly upstream into the oncoming flow, we believe this assumption is warranted. The meteorological mast measures all three components of velocity.
- We assume that lidars instantaneously scan their 2D x - y plane of interest. In practice, both lidars scan across their 2D planes in approximately 2–5 s.

III. METHODS

A. Background

In order to generate synthetic atmospheric states, we employ an LDM³⁹ (Figs. 2 and 3, Table I), an architecture based on DMs. In recent years, DMs have emerged as a new category of deep generative model. They were originally developed through the perspective of non-equilibrium statistical physics,³⁶ but they can be derived through the lens of score-based modeling³⁷ or Markovian hierarchical variational autoencoders.⁴⁸ From one perspective,⁴⁹ DMs are trained with the assistance of a prescribed degradation process, in which Gaussian noise is repeatedly added to a sample \mathbf{x} over an interval $t \in [0, T]$ until the final sample is indistinguishable from pure Gaussian noise. This process can be thought of as a forward stochastic differential equation,

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w}, \tag{1}$$

where \mathbf{f} is the drift coefficient, g is the diffusion coefficient, and \mathbf{w} is the standard Wiener process. This process can be undone such that the final Gaussian noise state can be reverted to the initial sample using the reverse stochastic differential equation,

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t)dt - g(t)^2 \nabla_x \log p_t(\mathbf{x})]dt + g(t)d\bar{\mathbf{w}}, \tag{2}$$

where $\bar{\mathbf{w}}$ is the standard Wiener process for the reverse equation and $\nabla \log p_t(\mathbf{x})$ is the score function, which is challenging to estimate. DMs learn the score function, after which they can be used to draw samples from the probability distribution function (pdf) $p_0(\mathbf{x})$ (which we will simply refer to as $p(\mathbf{x})$ from here on) that characterizes the training dataset by initially starting with Gaussian noise. Given supplemental information, such as observations or class labels, DMs can draw samples from a conditional probability density function $p(\mathbf{x}|\mathbf{y})$ through a number of approaches—for example, training on paired data to directly learn to sample from $p(\mathbf{x}|\mathbf{y})$ as is done by Song and Ermon (2020),³⁷ or by first learning to sample from the unconditional $p(\mathbf{x})$ and then learning an additional conditioning mechanism.⁵⁰ We employ the first approach in this study.

While traditional DMs have been shown to generate high-quality images, they can be prohibitively expensive to train for large samples, and three-dimensional samples are often large. LDMs address this issue by compressing raw, pixel-space samples into a latent space with the use of an autoencoder (AE) [Fig. 2(a)]. The AE is traditionally trained

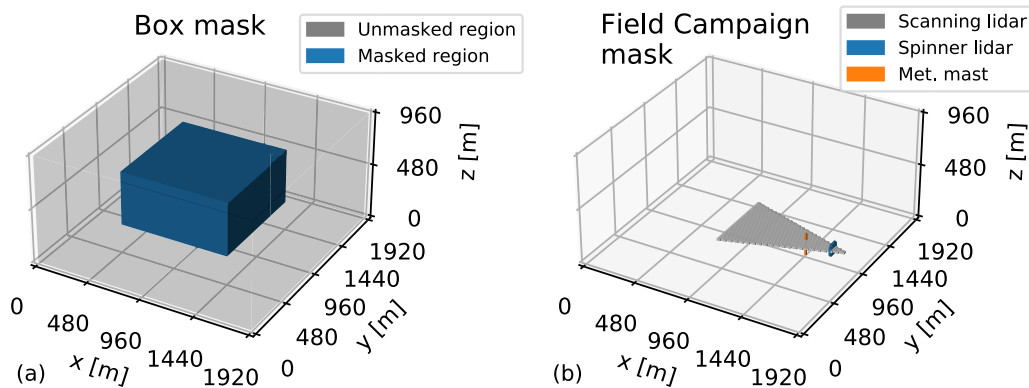


FIG. 1. (a) The masked region and observed region associated with the box mask. (b) The observations associated with the field campaign (FC) mask, with each instrument highlighted. The regions without observations are masked.

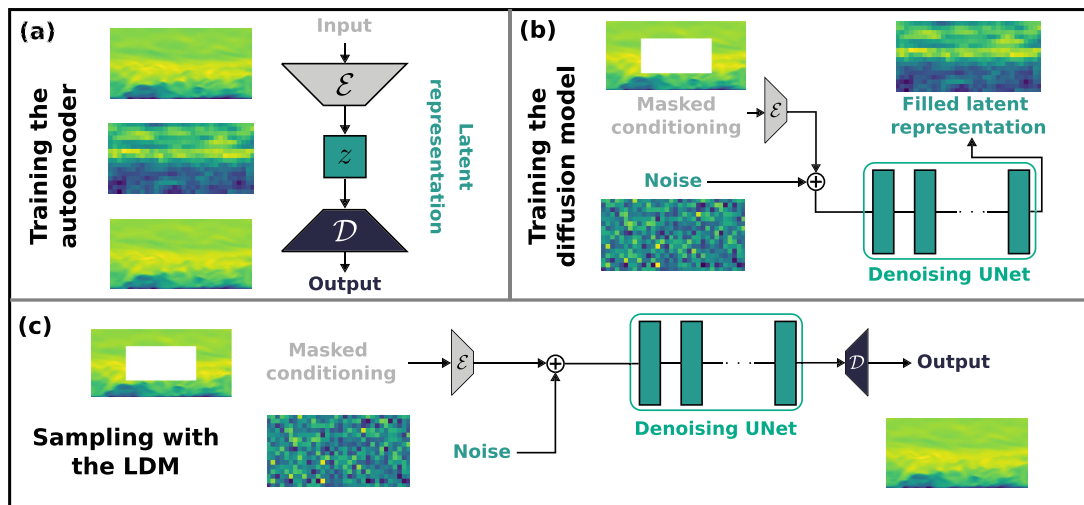


FIG. 2. A schematic depicting the major components of an LDM and their function. (a) First, the encoder and decoder of an autoencoder are trained. (b) Next, a diffusion model is trained with the help of the encoder from the autoencoder. The conditioning is optional and is not included for unconditional networks. (c) The trained encoder, diffusion model, and decoder are combined into an LDM and used to generate samples, possibly given optional conditioning. Additional details can be found in Fig. 3 and Table I.

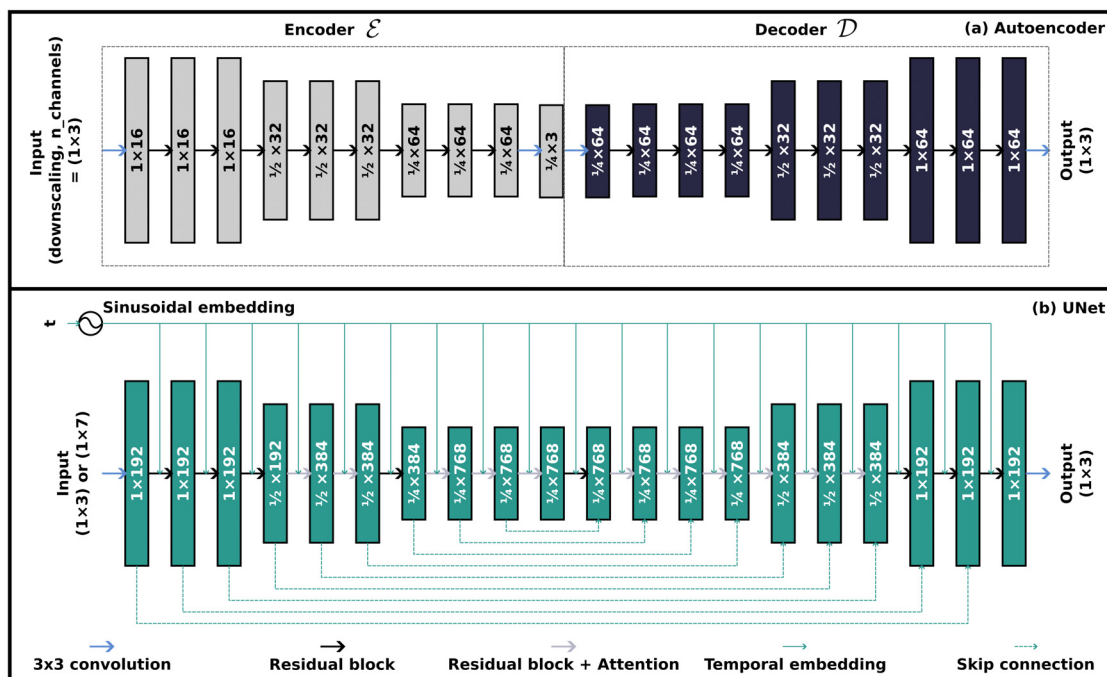


FIG. 3. A schematic showing the internal components of the LDM in greater detail, including (a) in the autoencoder and (b) in the UNet of the diffusion model.

with a mixture of L_1 or L_2 losses, Kullback–Leibler divergence losses as in a variational AE,⁵¹ perceptual losses based on the VGG network,⁵² and patch-based adversarial losses.⁵³ After the AE is trained, a diffusion model is trained to generate samples in latent space [Fig. 2(b)]. After the components of the autoencoder and diffusion model are trained, they can be combined to generate samples in pixel-space [Fig. 2(c)].

B. Modifications to the original LDM

We modify the original LDM architecture so that it can be applied to generate synthetic LES data. The exact implementation of the architecture can be found at <https://zenodo.org/doi/10.5281/zenodo.10206880>, and we provide a summary of the modifications to the LDM code here.

TABLE I. Additional parameters used during training of the autoencoder and the diffusion model.

Parameter	Value
Batch size, AE	2
Batch size, DM	2
Base learning rate, AE	4.5×10^{-6}
Base learning rate, DM	5×10^{-5}
Noise schedule	Linear
N noise steps	1000
Start variance	0.0015
End variance	0.0155
N attention heads	8

- The original LDM was developed for 2D images, so we modify the architecture to work on 3D data through the use of operations like 3D convolutions and 3D normalization.
- While the original LDM uses several attention blocks⁵⁴ throughout the AE, we omit their use there due to the computational demands of a 3D attention block.
- We replace group normalization⁵⁵ with instance normalization⁵⁶ everywhere except within attention blocks, as we found the latter performed better.
- In the process of selecting the finalized architecture for this paper, we experimented with different weights for each of the loss components in the AE. We also added the option for a physics-based loss that assesses mass conservation by calculating the divergence of the velocity field. In the end, we found the best performance by using an L_1 term, Kullback–Leibler divergence term, and whole-sample adversarial term in the loss function. We omit the VGG-based adversarial loss as well as the mass conservation loss.
- In the original LDM architecture,³⁹ image inpainting is accomplished by first drawing a sample from $p(x|y)$ using unmasked regions as conditioning information. In this scenario, there is no guarantee that the generated sample exactly agrees with the conditioning information y , so inpainting is achieved by overlaying y onto the sample in a postprocessing step. In effect, this treats conditioning information as a hard constraint, meaning the observation is exactly matched in the reconstruction, which makes sense when inpainting images. However, we found this hard constraint would lead to artifacts at the boundaries of FC mask observations, likely because these observations are only one pixel wide in certain dimensions. As such, we do not postprocess LDM output to exactly match the observations, thereby treating observations as soft constraints, meaning observations are not exactly matched in reconstructions. In other common atmospheric reconstruction techniques, namely, data assimilation,⁵⁷ observations are also treated as soft constraints. We note that in data assimilation the relative influence of observations on reconstructions can be controlled by prescribing a certain measurement noise magnitude, whereas in our LDM technique here, we cannot modulate the influence of observations.

C. Network configurations

The LDM requires an AE network for compression to the latent space and decompression [Fig. 3(a)]. Here, the AE has three input,

latent space, and output channels that correspond to the velocity variables (u , v , w). Future work could expand this to learn other quantities like temperature, though we prioritize the three velocity components in this study as these have the greatest relevance to wind turbine dynamics. The encoder and decoder of the AE have three internal levels, each with three residual blocks,⁵⁸ compressing data with spatial dimensions of (128, 128, 64) to (32, 32, 16). The first internal level uses 16 channels, followed by 32 channels, and then 64 channels. We use a batch size of 2 due to the high memory use of 3D convolutions. This network is much smaller than the original LDM used in Rombach *et al.*,³⁹ <https://github.com/CompVis/latent-diffusion>, which, for example, had six internal layers with 128 base channels. Our network is smaller due to computational constraints (we train on two 16 GB V100s) as well as our large three-dimensional samples. Following common practice, we rescale the input data so that the values in each channel fall in the range of -1 to 1 .

While our primary goal is to generate candidate velocity fields given FC-style observations, we train both a conditional DM and an unconditional DM to provide more context on network behavior. Both DMs use the same AE for conversion between pixel space and latent space. The DMs use a UNet architecture⁵⁹ with three internal layers, respectively, using (192, 384, 768) channels, and two residual blocks per layer. Attention blocks are used in the 384 channel layers, 768 channel layers, and the center of the UNet. The DM noise schedule uses 1000 diffusion steps and is linear. The only difference between the unconditional and conditional DM architecture is the input layer—the unconditional network uses three channels, and the conditional network has an additional four channels, three of which correspond to the compressed observation in latent space and the last of which corresponds to a compressed binary mask corresponding to observed pixels. The conditional DM is exposed to several different masks during training, including box-style and the FC-style observations, in order to increase robustness and potentially accuracy.⁴² The same conditional DM is then used to produce both box-style and FC-style samples.

IV. RESULTS

Below, we assess the ability of the LDM to generate plausible atmospheric states. We begin by visualizing unconditional and conditional LDM samples and qualitatively assessing them, as is common in both computer vision manuscripts,^{39,49} as well as turbulent flow manuscripts.^{16,18,60} Next, we quantitatively assess the LDM samples by calculating mean profiles, statistics, spectra, and divergence. Finally, we quantitatively assess the ability of the LDM to generate diverse samples from a single observation.

A. Qualitative assessment of LDM samples

First, in order to broadly demonstrate that LDMs can generate samples that look like an ABL, we compare cross sections of u from the unconditional LDM to those from the ground-truth ABL training data (Fig. 4). Additional qualitative visualizations of v and w as well as three-dimensional isocontours are available in Appendix A. Similar to the training data, the LDM samples have winds that are generally slow near the ground and stronger aloft until they hit the capping inversion near 480 m, above which the winds are less turbulent. The LDM turbulent structures look like turbulent structures in the ground truth data. The horizontal boundaries of the samples are periodic, just as they are in the testing data. Crucially, however, closer inspection reveals that the LDM

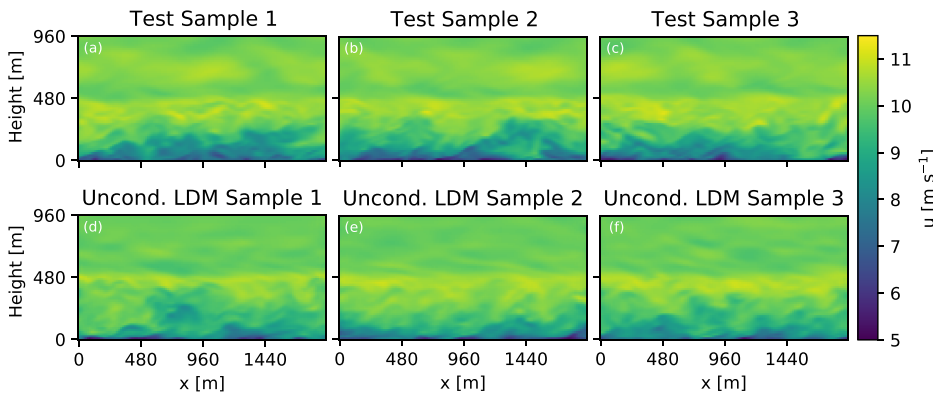


FIG. 4. (a)–(c) Streamwise cross sections of u at $y = 960$ m from ground truth, test data. (d)–(f) Same as above, except for samples from an unconditional LDM.

samples appear slightly smoother than the ground truth data. This discrepancy on smaller scales is further characterized in Sec. IV B.

We next show that conditional LDMs can inpaint well when provided abundant data, as is the scenario with the box max (Fig. 5). We visualize a u cross section of a ground truth sample, the associated box mask observation, one conditionally generated sample, and the mean and standard deviation from $n = 100$ conditionally generated samples. The conditional samples look similar to the unconditional samples, except the perimeter of the conditional sample and the averaged field is relatively unchanged. The transition between the unmasked and the masked region is smooth and does not display any artifacts. This smooth transition is in part achieved by treating observations as a soft constraint instead of a traditional inpainting hard constraint, in which the exact values of the observations would be superimposed over the unmasked region. The statistical behavior of the conditional predictions qualitatively matches expected behavior. The mean prediction clearly displays the capping inversion, above which the mean flow structures inside the masked region agree well with the structures in the unmasked region. Beneath the capping inversion, the mean prediction is smoothed out when compared to individual turbulent realizations. Inside the unmasked region, the standard deviation of LDM predictions is largest at the bottom of the mask and smaller aloft. This aligns with physical behavior and is expected, as mean shear gradients at the surface encourage mechanical generation of turbulent kinetic energy and therefore wind speed variance. The LDM also displays small, nonzero standard deviation in the unmasked regions, which would not be the case if we treated observations as a hard constraint.

Finally, we find that conditional LDMs can generate samples that look realistic even when provided minimal observations, as is the case for the FC mask (Fig. 6). We include additional visualizations, including visualizations of v and w which are only conditioned on meteorological mast measurements, in Appendix B. For the particular observation shown in Fig. 6, there is a region of slow u -wind in the vicinity of the meteorological mast. Correspondingly, each of the samples also has a patch of slow wind in the same area. This slow flow structure is especially clear in the mean prediction [Figs. 6(d) and 6(j)], and qualitatively, the mean prediction aligns well with the ground truth in the vicinity of the observations. This behavior illustrates that the LDM includes the conditional information of the measurements when generating samples. Away from the observations, the mean prediction becomes much smoother, showing that the conditioning only has local effects, as would be expected in a turbulent flow. The standard deviation of the predictions for the FC mask [Figs. 6(e) and 6(k)] also shows a smaller standard deviation in the vicinity of the observations, especially near the horizontal lidar measurement. Thus, the conditioning is reducing the sample-to-sample spread near the measurements, as desired. The FC samples exhibit much higher variance than the box mask samples [Fig. 5(e)], as would be expected because the FC samples are conditioned on a smaller region. As we discuss in Sec. IV C in greater detail, the sample standard deviation qualitatively matches up with a different estimate of the standard deviation that comes from another tool [Fig. 6(f)].

After assessing the visual quality of samples, we further demonstrate the quality of LDM samples by showing that they can successfully

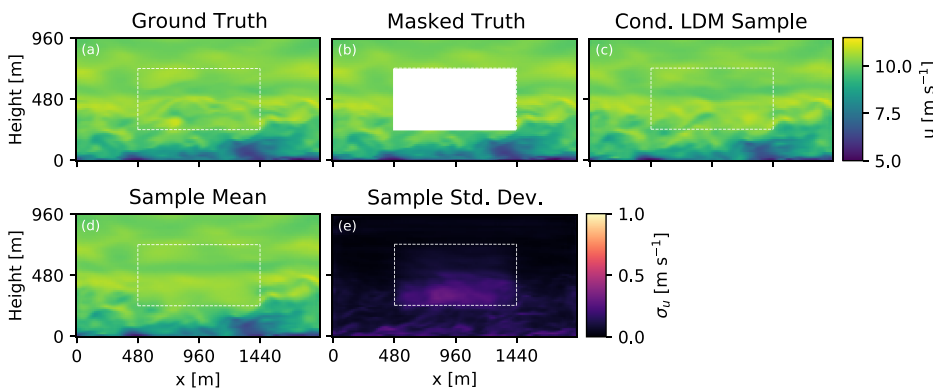


FIG. 5. Streamwise cross sections of u at $y = 960$ m for (a) a ground truth sample, (b) the box observation created from the ground truth sample, and (c) one sample from the conditional LDM, given the observation. (d) The mean from $n = 100$ conditional LDM samples. (e) The standard deviation from $n = 100$ conditional LDM samples. All subpanels display a white dashed line as a reference to the mask.

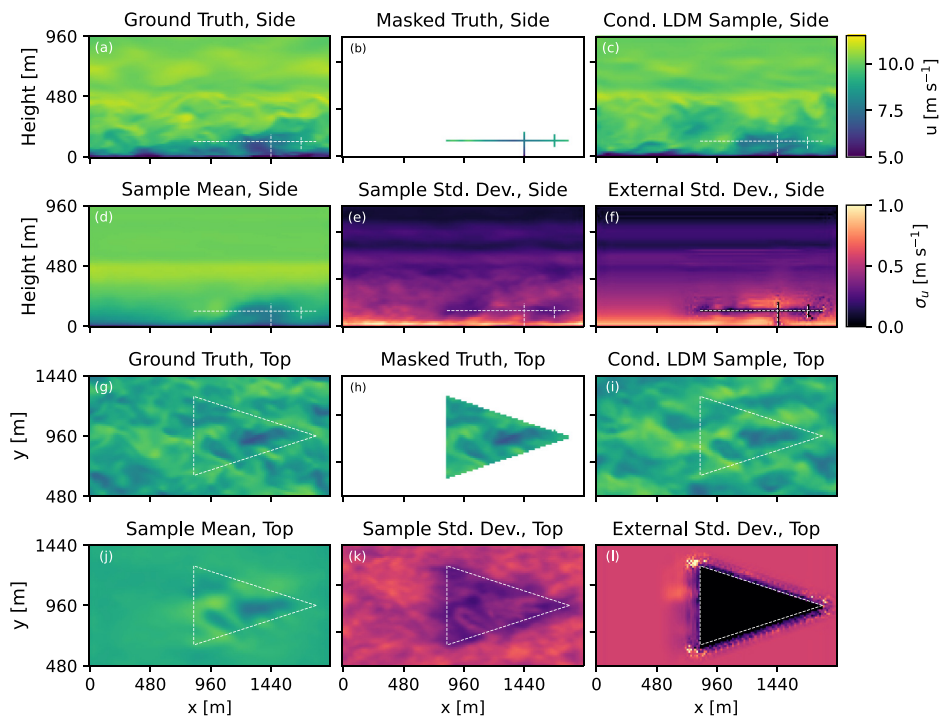


FIG. 6. (a)–(e) Same as Fig. 5, except for the FC mask. (f) An external estimate of the conditional standard deviation for the observation in panel (b), which is discussed in greater detail in Sec. IV C. (g)–(l) Corresponding top-down views at the height of the scanning lidar for the data in panels (a)–(f).

act initial conditions for our LES code. We generate a FC mask sample using the observation in Fig. 6 and then use it as the initial condition in a 1 h LES. We create a video that visualizes a time history of a stream-wise cross section and a top-down view of the simulation at <https://zenodo.org/doi/10.5281/zenodo.10206880>. The simulation successfully runs and does not display any obvious numerical artifacts such as discontinuities in the flow field around the masked region or ringing in the form of waves or instabilities.

B. Quantitative assessment of LDM samples

We next quantitatively assess the performance of the LDMs against physical aspects of the LES: vertical profiles of velocity components and kinematic fluxes, probability density functions of velocity components, turbulence spectra, and mass conservation. For completeness, we compare LDM sample statistics to those from both the training dataset and testing dataset. The training and testing datasets have similar but distinct statistics, as they come from distinct simulation periods. However, as we demonstrate below, their statistics match within our region of interest (rotor disk heights), thereby improving the probability of success for our flow reconstruction problem. To facilitate these quantitative comparisons, we examine characteristics between heights of 0 and 480 m, namely, below the capping inversion.

The LDM accurately captures average velocity profiles [Figs. 7(a)–7(c)], while slightly underperforming on kinematic flux profiles [Fig. 7(d)]. We compare horizontally averaged profiles from the training dataset, testing dataset, $n = 100$ unconditional samples, and $n = 100$ conditional samples from the FC mask. The mean and standard deviation of LDM velocity profiles agree well with training data at all heights and with testing data beneath the capping inversion. Thus, the LDM performs well on first-order statistics. However, the LDMs

do show some deviations when comparing kinematic fluxes, a second-order statistic that is more challenging to capture than the mean. In particular, the unconditional LDM underestimates the downward flux $\overline{u'w'}$ by as much as 20% between heights of 15 and 150 m. The conditional LDM performs better in this region with only an 8% discrepancy, possibly because of the near-surface conditioning information.

Similarly, the probability distribution functions (pdfs) of velocity components show broad agreement, while missing fine-scale details (Fig. 8). We calculate velocity distributions using the same three-dimensional data that were used to calculate average vertical profiles. The pdfs from LDM samples match the pdfs from training data well in all bins, only showing minor discrepancies.

While the power spectra of LDM samples match ground truth at large spatial scales, the LDM spectra have too much energy at the smallest spatial scales (Fig. 9). We calculate one-dimensional spectra of u at a height of 90 m. The LDM and ground truth spectra have similar values for wavenumbers smaller than 0.01 m^{-1} , and the LDM accurately captures the inertial scale of turbulence, as visualized by the $-5/3$ slope. However, for wavenumbers larger than 0.01 m^{-1} , the LDM has too much energy. This behavior has been observed in other machine learning reconstruction studies^{16,18} and is consistent with the “spectral bias” problem⁶¹ inherent to many neural networks—the autoencoder part of our network seeks to minimize the L_1 of the samples, thereby prioritizing reconstructing fields on the largest spatial scales. As such, the smallest scales have a relatively small impact on the loss function.

Finally, we assess the LDM’s ability to satisfy mass conservation by examining distributions of velocity divergence $\frac{\partial u_i}{\partial x_i}$ (Fig. 10). If continuity were exactly satisfied, the distribution of divergence would be shaped like a Dirichlet function at $\frac{\partial u_i}{\partial x_i} = 0$. The testing and training




FIG. 7. (a)–(c) Vertical profiles of horizontally averaged velocity components from training data, testing data, unconditional LDM samples, and FC-mask conditional LDM samples. Averages are shown as solid lines, and ± 1 standard deviations are shown as faded lines. (d) Kinematic fluxes of the same data.

data approximately satisfy this behavior, showing small deviations from 0 due to numerical artifacts. While distributions for both the unconditional and conditional LDM data show a spike centered on 0, their spikes are not nearly as sharp. Thus, the LDMs appear to have some awareness of mass conservation, but they do not satisfy it here as well as the ground truth data.

C. Assessment of ensembles

1. Rank histogram analysis

For ensemble-based flow reconstruction, it is important to have a well-calibrated ensemble.³⁴ Ensemble members are drawn from some pdf, and in an ideal scenario, the ground truth would be seen as just another draw from that pdf, such that the ground truth is statistically indistinguishable from the ensemble members, a condition termed “ensemble consistency.” Ideally, the simulated ensemble would show zero bias relative to the ground truth, and the ensemble should have an appropriate but not excessive amount of spread.

One common approach to assess ensemble consistency in the geosciences is the rank histogram.^{60,62} When given a group of predictions, the rank of the ground truth is calculated by sorting the ensemble members from smallest to largest by some scalar quantity (e.g., u at a particular grid cell) and then identifying the position of the ground truth relative to the ensemble members. When repeated for several reconstructions, rank calculations can be collected into a rank histogram, a diagram that can diagnose ensemble bias as well as overconfidence or underconfidence.

We assess ensemble quality by compiling a rank histogram using the first 6 h (360 observations) from the test set and generating 10 ensemble members per observation (Fig. 11). We calculate the rank of each ground truth by examining u at the unobserved grid cell immediately downwind of the first scanning lidar observation, a cell that can be thought of as the turbine nacelle. An ideal rank histogram would look like a uniform distribution. However, the rank histogram here

looks like a combination of a U-shape and a downward linear slope. The U-shape shows that LDM samples are underly diverse—the ground truth disproportionately falls at one extreme of the ensemble. The downward slope indicates that the ground truth tends to be disproportionately smaller than the rest of the ensemble, or in other words, the ensemble often overpredicts. We verify this by calculating the bias of the LDM samples relative to their respective ground truths, finding that u is on average 0.02 m s^{-1} higher. Thus, while the downward slope suggests a tendency to overpredict u at turbine hub height, the near-zero bias is balanced out by the right arm of the U-shaped distribution. In the end, the LDM ensembles deviate from ideal behavior, which is in practice the case even for skillful ensemble forecasts.^{63,64} These deviations may be small in the end, and future work will use these ensembles in a data assimilation process and assess if the unoptimal behavior is problematic in practice.

2. External estimate of conditional standard deviation

While rank histograms assess ensemble consistency and are used primarily in the ensemble weather forecasting field, we can alternatively assess the ensemble quality by examining “sample diversity.” This concept is commonly examined in the deep generative modeling field, and when applied to computer vision problems, metrics such as Inception Score³⁵ are commonly used. Here, we assess sample diversity by simply examining standard deviation, which is a valuable measure for physics-based problems. In Sec. IV A, we assessed LDM sample diversity by taking a single observation, e.g., the one in Fig. 6(a), and calculating the standard deviation of u across 100 conditionally generated samples [Fig. 6(e)]. This conditional standard deviation was calculated using samples from the LDM for a given observation. Alternatively, by applying tools such as fully connected neural networks or stochastic estimation²³ to our training dataset, it is possible to estimate what the conditional standard deviation should be for a given observation, without the need to generate any samples. This function




FIG. 8. The probability distribution function for training data, testing data, unconditional LDM samples, and conditional LDM samples for (a) u , (b) v , and (c) w .

has no awareness of the LDM or the samples it generates, and as such, it serves as an external check on the sample-calculated conditional standard deviation.

Following Hassanaly *et al.*,²³ we obtain an external estimate [Figs. 6(f) and 6(l)] of the conditional standard deviation for the observation in Figs. 6(a) and 6(g) by training a fully connected neural network. We provide more details on the external estimate methodology in Appendix C. From the side [Fig. 6(f)], this external estimate agrees well with the LDM sample standard deviation [Fig. 6(e)], but it fills in weaknesses that arise from calculating the sample standard deviation with a finite number of samples. From the side, the two standard deviations agree well in magnitude everywhere: the lowest grid cells, the observation network, and above the capping inversion. Both estimates




FIG. 9. The turbulence spectra of u at $z=90$ m for training data, testing data, unconditional LDM samples, and conditional LDM samples.

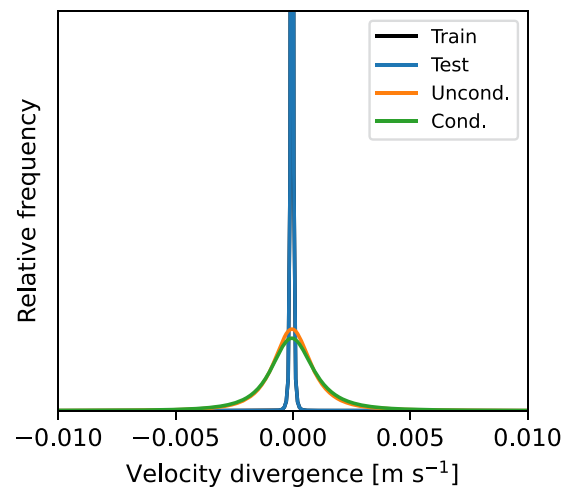


FIG. 10. An assessment of mass conservation, as visualized by distributions of velocity divergence for training data, testing data, unconditional LDM samples, and conditional LDM samples. Training and testing distributions sit on top of one another.

also show larger variance just upwind of the meteorological mast and above the scanning lidar, a behavior that arises due to the presence of a coherent structure. When viewed from the top, however [Fig. 6(l)], the external estimate disagrees with the sample standard deviation. This discrepancy arises because the LDM here treats observations as a soft constraint, whereas the external estimate treats observations more like a hard constraint. As such, the external estimate shows a near-zero standard deviation in pixels where observation is available. From both the side and the top, the external estimate clearly shows that the sensing network has an impact on the flow reconstructions close to the observations, up to a maximum reach of roughly 150m in any

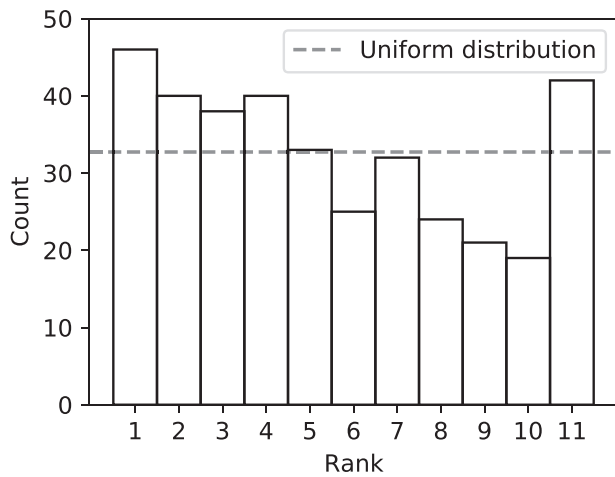


FIG. 11. Rank histogram for u at the turbine nacelle, given 10-member LDM FC ensembles for 360 observations.

direction. By comparing the two standard deviations, we build confidence that the LDM does not show major deficiencies in terms of sample diversity. This result complements the rank histogram analysis, suggesting that indeed the LDM is only slightly underly diverse.

3. Influence of FC observations on reconstruction accuracy

Finally, we quantify the impact of FC observations on the reconstruction accuracy at three locations (Fig. 12). We calculate vertical profiles of root mean squared error (RMSE) using the dataset that was used in the rank histogram analysis: 360 ground truth samples from which observations are sampled, and 10 corresponding reconstructions for each ground truth. We calculate RMSE along the centerline of observations ($y = 960$ m) at the location of the meteorological mast ($x = 1440$ m), the upstream edge of scanning lidar measurements ($x = 795$ m), and an upstream location where conditional standard deviation analysis revealed that observations no longer had an influence on reconstructions ($x = 480$ m).




FIG. 12. Vertical profiles of root mean squared error for (a) u , (b) v , and (c) w along the centerline of the domain at the meteorological mast ($x = 1440$ m), the upstream edge of scanning lidar measurements ($x = 795$ m), and a location upstream of all observations ($x = 480$ m).

The RMSE profiles show that observations improve reconstruction accuracy predominantly as would be expected. Far away from observations at $x = 480$ m, the RMSE values for all three velocity components are generally largest when compared to the other locations. The shape of these RMSE profiles correlates well with velocity standard deviations (Fig. 7)—heights with larger standard deviations are also associated with larger RMSE for the uncorrelated location. Within the perimeter of the synthetic field campaign at $x = 795$ m and $x = 1440$ m, observations have their strongest impacts at rotor disk heights, with relatively small impacts above and below the rotor disk. Relative to the RMSE outside of the measurement perimeter, the RMSE of u is decreased most substantially (approximately 0.40 m s^{-1}) at the height of the scanning lidar measurements, as observations of u are abundant at this elevation. While the field campaign measurements predominantly measure u , the observations still marginally impact RMSE profiles of v and w at rotor disk heights, decreasing v RMSE by approximately 0.05 – 0.01 m s^{-1} at most and w RMSE by approximately 0.1 m s^{-1} . The meteorological mast at $x = 1440$ m has the only observations of v and w , and correspondingly, RMSE for these variables slightly outperforms RMSE at $x = 795$ m.

V. CONCLUSION

In this paper, we investigate turbulent flow reconstruction in the context of a synthetic field campaign in the atmospheric boundary layer in which only spatially limited observations are available. We demonstrate that latent diffusion models (LDMs) create diverse, three-dimensional turbulent fields that are convincingly similar to true LES fields. These reconstructions match many physical characteristics of LES, such vertical profiles and the largest spatial scales of the spectra. However, LDMs struggle at the smallest spatial scales, diverging from LES spectra in this region and failing to preserve continuity.

Our LDM work extends machine learning turbulent flow reconstruction literature in several key manners. In contrast to many of the deterministic approaches, LDMs generate diverse samples, which is important for turbulent environments. Our LDM reconstructed samples with minimal observations ($< 1\%$ of the domain is observed), though the algorithm also works well with abundant observations. Notably, LDM samples can be used as initial conditions for computationally stable LESs.

19 December 2023 15:29:10

This study suggests several lines of possible future inquiry for turbulent flow reconstruction. In upcoming work, we will explore applying LDMs to noisy, real-world measurements in the RAAW field campaign. The LDM architecture could be modified to improve performance at the smallest spatial scales, perhaps through the use of physics-informed losses.⁶⁵ Finally, diffusion model researchers are investigating methods for quicker sampling,⁶⁶ which could potentially open the road to real-time flow reconstruction. In summary, the results indicate that diffusion models are a powerful class of machine learning algorithm that is worthy of further exploration in turbulence research.

ACKNOWLEDGMENTS

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding was provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

The research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. This work would not have been possible without several open-source diffusion modeling resources, including those by HuggingFace (<https://huggingface.co/blog/annotated-diffusion>), Janspiry (<https://github.com/Janspiry/Palette-Image-to-Image-Diffusion-Models>), Walter H. L. Pinaya (https://huggingface.co/spaces/Warvito/diffusion_brain), and CompVis (<https://github.com/CompVis/latent-diffusion>). The analysis was performed using a number of open-source Python libraries, including numpy,⁶⁷ Xarray,⁶⁸ and matplotlib.⁶⁹ Thank you to Hristo Chipilski for helpful discussions regarding data assimilation, and Michael Kuhn for AMR-Wind support. Additionally, thanks go to the editor and the two reviewers in helping craft this manuscript.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Alex Rybchuk: Investigation (lead); Methodology (lead); Writing – original draft (lead). **Malik Hassanalay:** Formal analysis (supporting); Writing – review & editing (supporting). **Nicholas Hamilton:** Investigation (supporting); Project administration (supporting); Writing – review & editing (supporting). **Paula Doubrawa:** Funding acquisition (lead); Project administration (equal); Writing – review & editing (supporting). **Mitchell J. Fulton:** Investigation (supporting); Writing – review & editing (supporting). **Luis A. Martinez-Tossas:** Project administration (equal); Writing – review & editing (supporting).

DATA AVAILABILITY

The code for our modified latent diffusion model as well as the AMR-Wind configuration used to generate data are available at <https://zenodo.org/doi/10.5281/zenodo.10206880>. While the entire dataset is too large to store on a public repository, we have included a truncated version of the dataset at <https://doi.org/10.5281/zenodo.10206885>. This link also contains the trained model weights of the autoencoder and the diffusion model used to generate samples in this manuscript.

APPENDIX A: ADDITIONAL VISUALIZATIONS OF TEST DATA AND UNCONDITIONAL SAMPLES

In this section, we provide supplemental visualizations for samples generated by an unconditional LDM (Figs. 13–15).

APPENDIX B: ADDITIONAL VISUALIZATIONS OF BOX AND FC CONDITIONAL SAMPLES

In this section, we provide supplemental visualizations for samples generated by conditional LDMs trained on the box mask (Figs. 16,17) and the FC mask (Figs. 18–20).




FIG. 13. (a)–(f) Same as Fig. 4, but for v .




FIG. 14. (a)–(f) Same as Fig. 4, but for w .

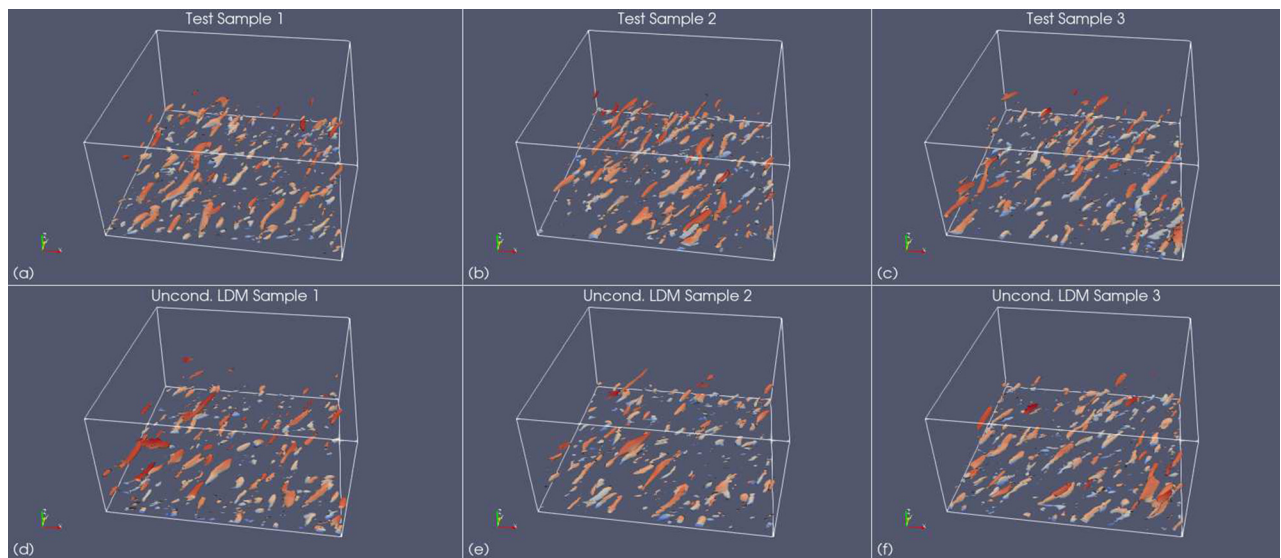


FIG. 15. (a)–(f) Isometric views of vortices identified as contours of the Q-criterion⁷⁰ with a value of 0.0003 m s^{-1} .




FIG. 16. A distinct (a) ground truth and (b) box observation from the one shown in Fig. 5. The (c) mean prediction and (d) standard deviation of prediction from 100 LDM samples are shown. (e)–(h) Four distinct LDM samples that are conditioned on the observation.

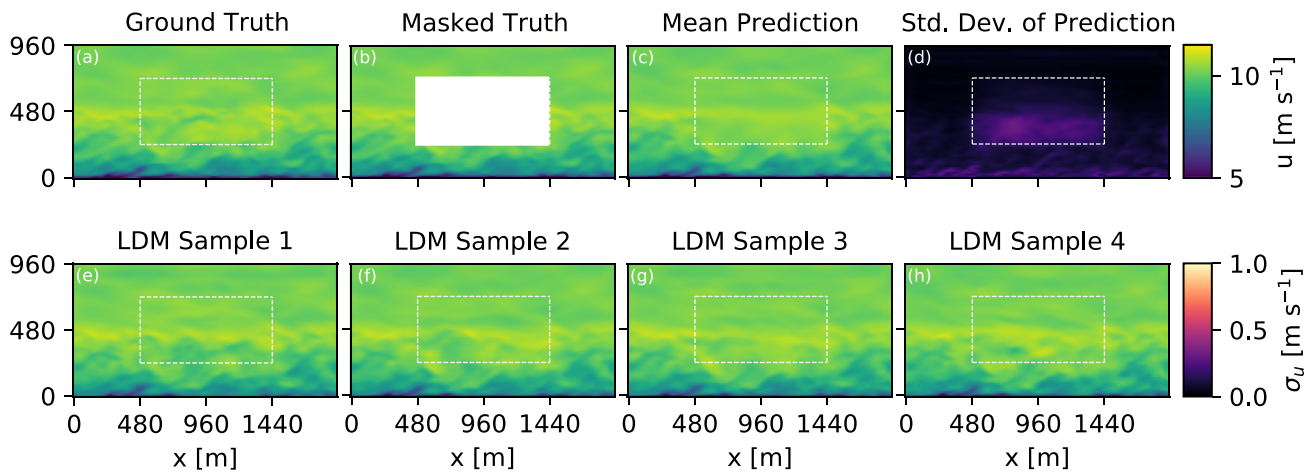


FIG. 17. (a)–(h) Same as Fig. 16, except for a distinct observation.

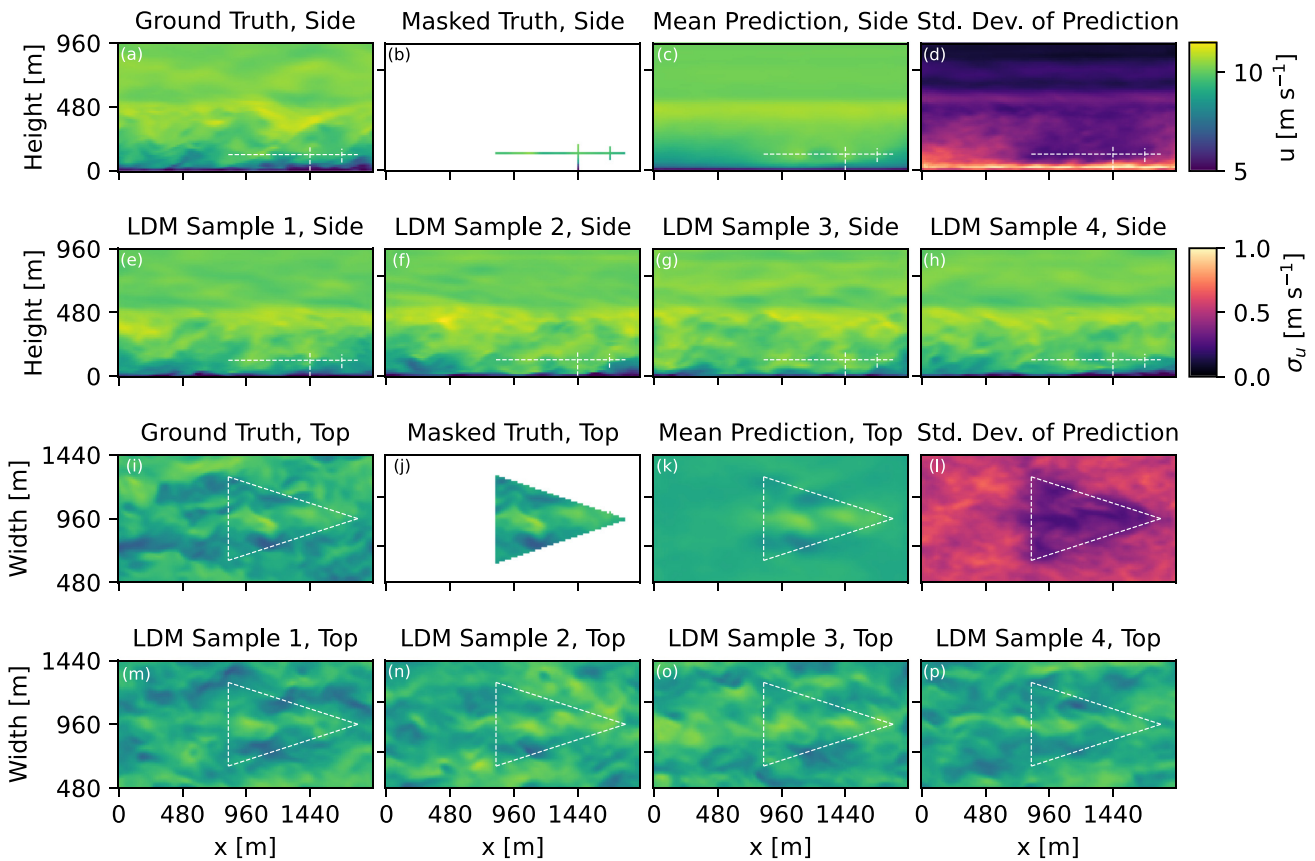


FIG. 18. A distinct (a) ground truth and (b) FC observation of u from the one shown in Fig. 6. The (c) mean prediction and (d) standard deviation of prediction from 100 LDM samples are shown. (e)–(h) Four distinct LDM samples that are conditioned on the observation.

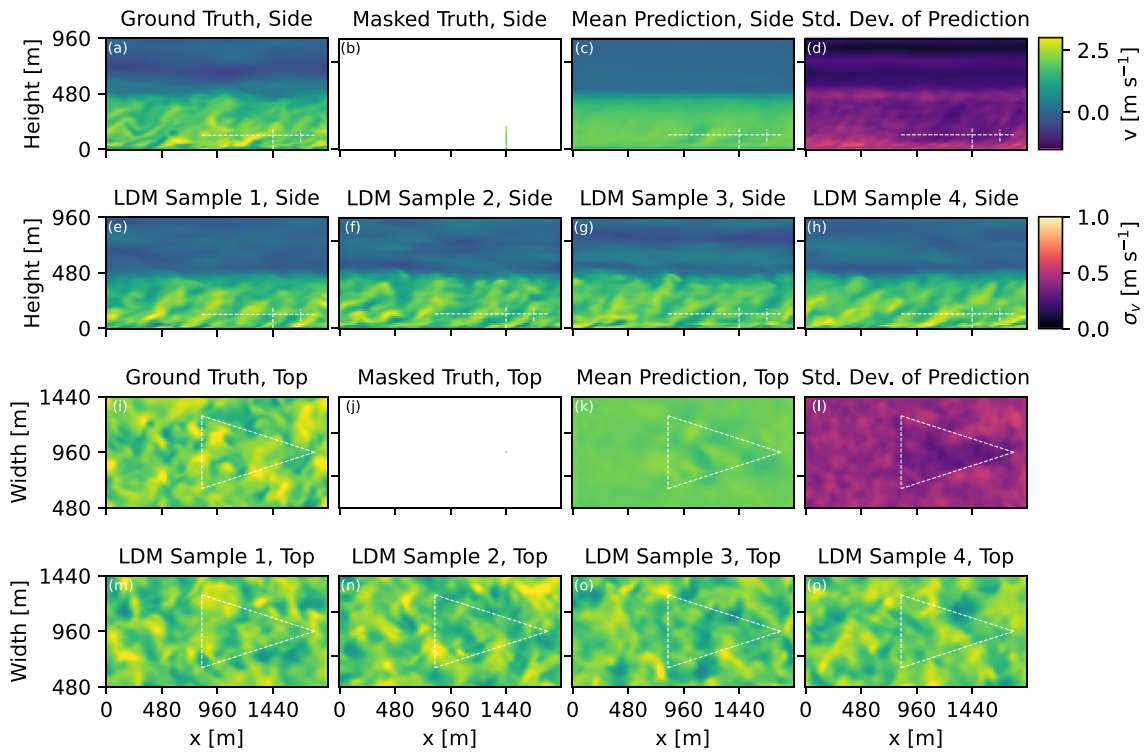


FIG. 19. (a)–(p) The same as Fig. 18, except for v .

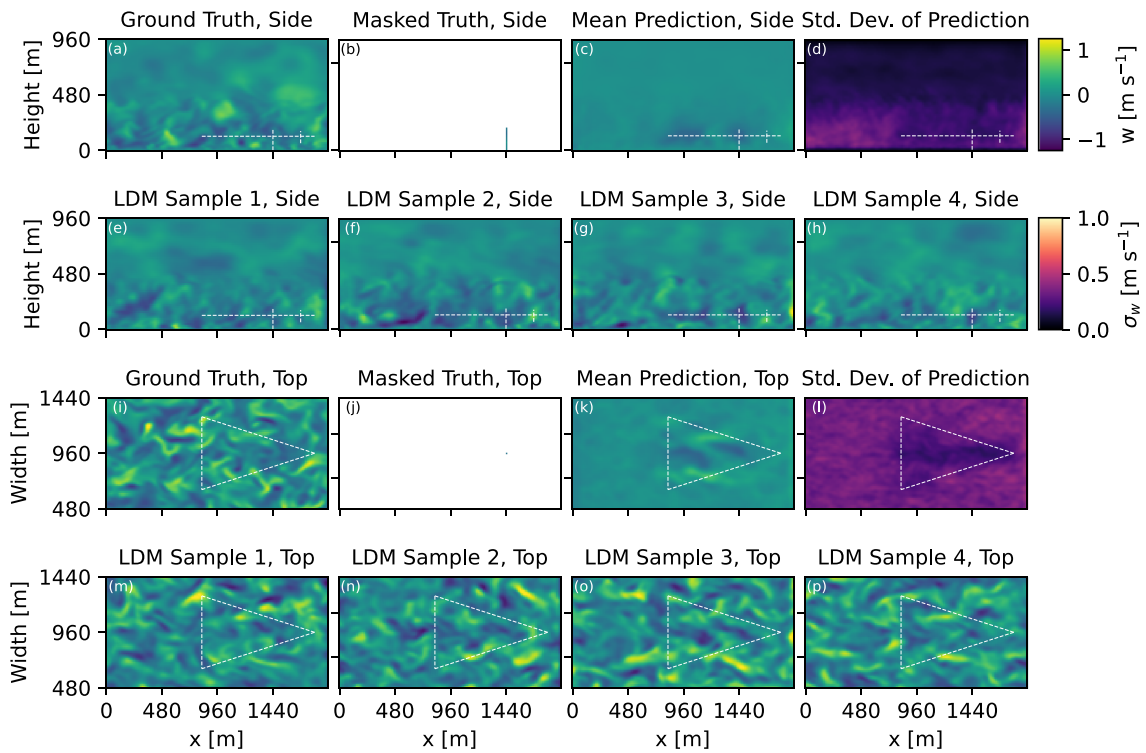


FIG. 20. (a)–(p) The same as Fig. 18, except for w .

19 December 2023 15:29:10

APPENDIX C: TECHNICAL DETAILS FOR EXTERNAL ESTIMATES OF THE CONDITIONAL STANDARD DEVIATION

The *a priori* moment estimation method is detailed here. Consider an observation $y \in \mathbb{R}^d$ where d is the length of the observation vector, and $x \in \mathbb{R}^m$ where m is the length of the state vector. The function f that best approximates x given y , i.e., that minimizes $\|f(y) - x\|_2$, is called the optimal estimator of x given y . In addition, the optimal estimator is the conditional expectation of x given y .⁷¹ Therefore, by seeking the optimal estimator of x given y , one can approximate the conditional moments of x given y . Likewise, by seeking the conditional moments of x^2 given y , one can approximate the diagonal of the covariance matrix of the distribution sampled (i.e., the pointwise variance).

We obtain *a priori* estimates of conditional standard deviation by using a neural network. We construct the neural network to approximate the optimal estimator of the atmospheric state (for the first moment) and the square of the atmospheric state (for the second moment), given the observed field as the input. We use the architecture of Ledig *et al.*,⁷² a fully convolutional network with residual blocks and skip connections.

To ensure that a reasonable approximation of the optimal estimator has been reached, we train multiple neural nets with increasing complexity, until the mean square-error (MSE) loss stops improving. Similar to Hassanaly *et al.*,²³ we adjust two model architecture hyperparameters to converge to the optimal estimator: the number of filters per convolutional layers varied in the set $\{1, 2, 4, 8\}$ and the number of residual blocks varied in the set $\{2, 4, 8, 16\}$. Unlike in Hassanaly *et al.*,²³ the input and output of the network have the same dimension. The input is altered by setting the masked pixels to a zero value. The loss function is calculated by computing the MSE with the masked pixels only, thereby enforcing that unmasked pixels do not contribute to evaluating the conditional estimator. Since 3D convolutions are used and to accommodate memory requirements, the batch size is typically set to 8 but changed to 4 for the biggest network. A fixed learning rate of 1×10^{-3} was used. Similar to Hassanaly *et al.*,²³ the second moment is estimated by outputting the square of the state subtracted with the first moment. All networks considered here were sufficiently trained to reach convergence (Fig. 21).

Once fully trained, the final training loss is examined to decide if further increase in network complexity is warranted to appropriately approximate the optimal estimator (Fig. 22). We find that as

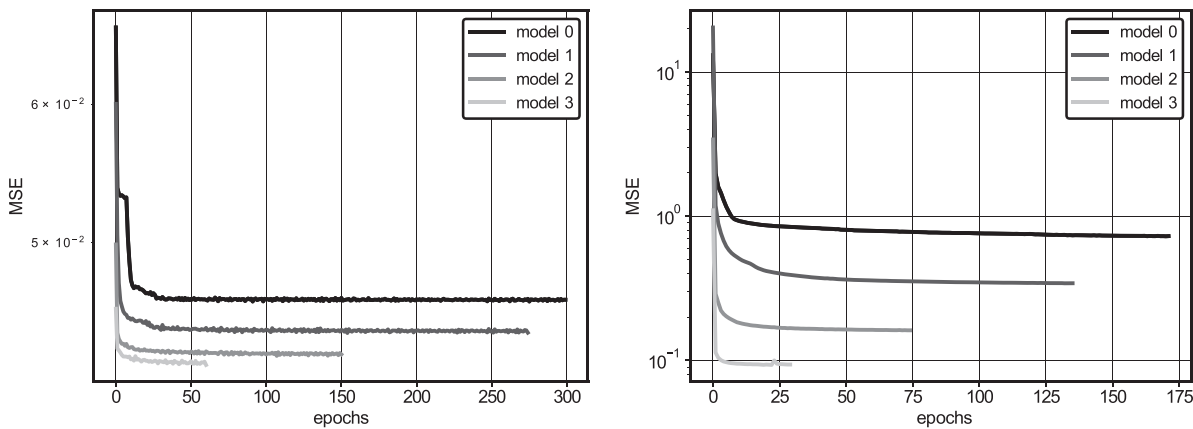


FIG. 21. MSE training loss vs epoch number for all the model architectures for FC mask for the first moment estimate (left) and the second moment estimate (right).




FIG. 22. Final MSE training loss vs number of trainable parameters for the FC mask for the first moment estimate (left) and the second moment estimate (right).

the number of trainable parameters increases beyond approximately 9000 to approximately 60 000, the MSE decreases until a point where adding additional trainable parameters only marginally improves the loss. As such we stop increasing the complexity of the estimator network and produce the figure shown in Fig. 6(f) using the model with the highest number of parameters.

REFERENCES

- ¹P. Moriarty, N. Hamilton, M. Debnath, T. Herges, B. Isom, J. K. Lundquist, D. Maniaci, B. Naughton, R. Pauly, J. Roadman, W. Shaw, J. van Dam, and S. Wharton, "American WAKE experiment (AWAKEN)," Technical Report No. NREL/TP-5000-75789, National Renewable Energy Lab. (NREL), Golden, CO, USA, 2020.
- ²C. L. Kelley, P. Doubrawa, and J. W. Naughton (2023) "Rotor, Aeroelastics, Aerodynamics, and Wake (RAAW) Project," [U.S. Department of Energy](#).
- ³M. L. Barad, *Project Prairie Grass, A Field Program in Diffusion: Volume 1* (Air Force Cambridge Research Center, 1958), Vol. 58.
- ⁴K. J. Allwine and J. E. Flaherty, "Joint Urban 2003: Study overview and instrument locations," Technical Report No. PNNL-15967, Pacific Northwest National Lab. (PNNL), Richland, WA, USA, 2006.
- ⁵C. Warneke, J. P. Schwarz, J. Dibb, O. Kalashnikova, G. Frost, J. Al-Saad, S. S. Brown, W. A. Brewer, A. Soja, F. C. Seidel, R. A. Washenfelder, E. B. Wiggins, R. H. Moore, B. E. Anderson, C. Jordan, T. I. Yacovitch, S. C. Herndon, S. Liu, T. Kuwayama, D. Jaffe, N. Johnston, V. Selimovic, R. Yokelson, D. M. Giles, B. N. Holben, P. Goloub, I. Popovici, M. Trainer, A. Kumar, R. B. Pierce, D. Fahey, J. Roberts, E. M. Gargulinski, D. A. Peterson, X. Ye, L. H. Thapa, P. E. Saide, C. H. Fite, C. D. Holmes, S. Wang, M. M. Coggon, Z. C. J. Decker, C. E. Stockwell, L. Xu, G. Gkatzelis, K. Aikin, B. Lefer, J. Kaspari, D. Griffin, L. Zeng, R. Weber, M. Hastings, J. Chai, G. M. Wolfe, T. F. Hanisco, J. Liao, P. C. Jost, H. Guo, J. L. Jimenez, J. Crawford, and T. F.-A. S. Team, "Fire influence on regional to global environments and air quality (FIREX-AQ)," *J. Geophys. Res.: Atmos.* **128**, e2022JD037758, <https://doi.org/10.1029/2022JD037758> (2023).
- ⁶M. Buziccotti, "Data reconstruction for complex flows using AI: Recent progress, obstacles, and perspectives," *Europhys. Lett.* **142**, 23001 (2023).
- ⁷C. Yu, X. Bi, and Y. Fan, "Deep learning for fluid velocity field estimation: A review," *Ocean Eng.* **271**, 113693 (2023).
- ⁸M. Meldi and A. Poux, "A reduced order model based on Kalman filtering for sequential data assimilation of turbulent flows," *J. Comput. Phys.* **347**, 207–234 (2017).
- ⁹P. Chandramouli, E. Memin, and D. Heitz, "4D large scale variational data assimilation of a turbulent flow with a dynamics error model," *J. Comput. Phys.* **412**, 109446 (2020).
- ¹⁰V. Mons, Y. Du, and T. A. Zaki, "Ensemble-variational assimilation of statistical data in large-eddy simulation," *Phys. Rev. Fluids* **6**, 104607 (2021).
- ¹¹M. Wang and T. A. Zaki, "State estimation in turbulent channel flow from limited observations," *J. Fluid Mech.* **917**, A9 (2021).
- ¹²T. A. Zaki and M. Wang, "From limited observations to the state of turbulence: Fundamental difficulties of flow reconstruction," *Phys. Rev. Fluids* **6**, 100501 (2021).
- ¹³P. C. Di Leoni, A. Mazzino, and L. Biferale, "Synchronization to big data: Nudging the Navier–Stokes equations for data assimilation of turbulent flows," *Phys. Rev. X* **10**, 011023 (2020).
- ¹⁴M. Zauner, V. Mons, O. Marquet, and B. Leclaire, "Nudging-based data assimilation of the turbulent flow around a square cylinder," *J. Fluid Mech.* **937**, A38 (2022).
- ¹⁵K. Fukami, K. Fukagata, and K. Taira, "Super-resolution analysis via machine learning: A survey for fluid flows," *Theor. Comput. Fluid Dyn.* **37**, 421 (2023).
- ¹⁶K. Fukami, K. Fukagata, and K. Taira, "Super-resolution reconstruction of turbulent flows with machine learning," *J. Fluid Mech.* **870**, 106–120 (2019).
- ¹⁷B. Liu, J. Tang, H. Huang, and X.-Y. Lu, "Deep learning methods for super-resolution reconstruction of turbulent flows," *Phys. Fluids* **32**, 025105 (2020).
- ¹⁸K. Fukami, K. Fukagata, and K. Taira, "Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows," *J. Fluid Mech.* **909**, A9 (2021).
- ¹⁹Y. Xie, E. Franz, M. Chu, and N. Thuerey, "tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow," *ACM Trans. Graphics* **37**, 1–15 (2018).
- ²⁰Z. 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).
- ²¹K. Stengel, A. Glaws, D. Hettinger, and R. N. King, "Adversarial super-resolution of climatological wind and solar data," *Proc. Natl. Acad. Sci. U. S. A.* **117**, 16805–16815 (2020).
- ²²M. 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).
- ²³M. Hassanal, A. Glaws, K. Stengel, and R. N. King, "Adversarial sampling of unknown and high-dimensional conditional distributions," *J. Comput. Phys.* **450**, 110853 (2022).
- ²⁴Z. Wang, X. Li, L. Liu, X. Wu, P. Hao, X. Zhang, and F. He, "Deep-learning-based super-resolution reconstruction of high-speed imaging in fluids," *Phys. Fluids* **34**, 037107 (2022).
- ²⁵M. Buziccotti, 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).
- ²⁶K. 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**, 945–951 (2021).
- ²⁷K. Gundersen, A. Oleynik, N. Blaser, and G. Alendal, "Semi-conditional variational auto-encoder for flow reconstruction and uncertainty quantification from limited observations," *Phys. Fluids* **33**, 017119 (2021).
- ²⁸S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis, "Flow over an espresso cup: Inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks," *J. Fluid Mech.* **915**, A102 (2021).
- ²⁹P. C. Di Leoni, K. Agarwal, T. A. Zaki, C. Meneveau, and J. Katz, "Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks," *Exp. Fluids* **64**, 95 (2023).
- ³⁰Y. Du, M. Wang, and T. A. Zaki, "State estimation in minimal turbulent channel flow: A comparative study of 4DVar and PINN," [arXiv:2210.09424](https://arxiv.org/abs/2210.09424) (2022).
- ³¹P. C. Di Leoni, L. Lu, C. Meneveau, G. E. Karniadakis, and T. A. Zaki, "Neural operator prediction of linear instability waves in high-speed boundary layers," *J. Comput. Phys.* **474**, 111793 (2023).
- ³²M. 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).
- ³³Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, and A. Anandkumar, "Physics-informed neural operator for learning partial differential equations," [arXiv:2111.03794](https://arxiv.org/abs/2111.03794) (2023).
- ³⁴D. S. Wilks, "Statistical methods in the atmospheric sciences," in *Statistical Methods in the Atmospheric Sciences*, 4th ed. (Elsevier, 2019), pp. i–ii.
- ³⁵T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen, "Improved techniques for training GANs," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2016), Vol. 29.
- ³⁶J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning* (PMLR, 2015), pp. 2256–2265.
- ³⁷Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," [arXiv:1907.05600](https://arxiv.org/abs/1907.05600) (2020).
- ³⁸P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2021), Vol. 34, pp. 8780–8794.
- ³⁹R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," [arXiv:2112.10752](https://arxiv.org/abs/2112.10752) (2022).
- ⁴⁰W. H. L. Pinaya, P.-D. Tudosiu, J. Dafflon, P. F. da Costa, V. Fernandez, P. Nachev, S. Ourselin, and M. J. Cardoso, "Brain imaging generation with latent diffusion models," [arXiv:2209.07162](https://arxiv.org/abs/2209.07162) (2022).
- ⁴¹A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "RePaint: Inpainting using denoising diffusion probabilistic models," in

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2022), pp. 11461–11471.
- ⁴²C. Saharia, W. Chan, H. Chang, C. A. Lee, J. Ho, T. Salimans, D. J. Fleet, and M. Norouzi, “Palette: Image-to-image diffusion models,” [arXiv:2111.05826](https://arxiv.org/abs/2111.05826) (2022).
- ⁴³P. Courtier, J.-N. Thépaut, and A. Hollingsworth, “A strategy for operational implementation of 4D-Var, using an incremental approach,” *Q. J. R. Meteorol. Soc.* **120**, 1367–1387 (1994).
- ⁴⁴G. Evensen, “Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics,” *J. Geophys. Res.: Oceans* **99**, 10143–10162, <https://doi.org/10.1029/94JC00572> (1994).
- ⁴⁵M. A. Sprague, “ExaWind: Predictive wind energy simulations,” Technical Report No. NREL/PR-5000-80401, National Renewable Energy Lab. (NREL), Golden, CO, USA, 2021.
- ⁴⁶A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, and M. L. Welcome, “A conservative adaptive projection method for the variable density incompressible Navier–Stokes equations,” *J. Comput. Phys.* **142**, 1–46 (1998).
- ⁴⁷J. Smagorinsky, “General circulation experiments with the primitive equations: I. The basic experiment,” *Mon. Weather Rev.* **91**, 99–164 (1963).
- ⁴⁸C. Luo, “Understanding diffusion models: A unified perspective,” [arXiv:2208.11970](https://arxiv.org/abs/2208.11970) (2022).
- ⁴⁹Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” [arXiv:2011.13456](https://arxiv.org/abs/2011.13456) (2021).
- ⁵⁰B. Kawar, M. Elad, S. Ermon, and J. Song, “Denoising diffusion restoration models,” [arXiv:2201.11793](https://arxiv.org/abs/2201.11793) (2022).
- ⁵¹D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013).
- ⁵²R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” [arXiv:1801.03924](https://arxiv.org/abs/1801.03924) (2018).
- ⁵³P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” [arXiv:1611.07004](https://arxiv.org/abs/1611.07004) (2018).
- ⁵⁴A. Brock, J. Donahue, and K. Simonyan, “Large scale GAN training for high fidelity natural image synthesis,” [arXiv:1809.11096](https://arxiv.org/abs/1809.11096) (2019).
- ⁵⁵Y. Wu and K. He, “Group normalization,” [arXiv:1803.08494](https://arxiv.org/abs/1803.08494) (2018).
- ⁵⁶D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Instance normalization: The missing ingredient for fast stylization,” [arXiv:1607.08022](https://arxiv.org/abs/1607.08022) (2017).
- ⁵⁷A. Carrassi, M. Bocquet, L. Bertino, and G. Evensen, “Data assimilation in the geosciences: An overview of methods, issues, and perspectives,” *WIREs Clim. Change* **9**, e535 (2018).
- ⁵⁸K. 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, 2016), pp. 770–778.
- ⁵⁹O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” [arXiv:1505.04597](https://arxiv.org/abs/1505.04597) (2015).
- ⁶⁰J. Leinonen, U. Hamann, D. Nerini, U. Germann, and G. Franch, “Latent diffusion models for generative precipitation nowcasting with accurate uncertainty quantification,” [arXiv:2304.12891](https://arxiv.org/abs/2304.12891) (2023).
- ⁶¹N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville, “On the spectral bias of neural networks,” in *Proceedings of the 36th International Conference on Machine Learning* (PMLR, 2019), pp. 5301–5310.
- ⁶²T. M. Hamill, “Interpretation of rank histograms for verifying ensemble forecasts,” *Mon. Weather Rev.* **129**, 550–560 (2001).
- ⁶³L. D. Monache, F. A. Eckel, D. L. Rife, B. Nagarajan, and K. Searight, “Probabilistic weather prediction with an analog ensemble,” *Mon. Weather Rev.* **141**, 3498–3516 (2013).
- ⁶⁴S. Ravuri, K. Lenc, M. Willson, D. Kangin, R. Lam, P. Mirowski, M. Fitzsimons, M. Athanassiadou, S. Kashem, S. Madge, R. Prudden, A. Mandhane, A. Clark, A. Brock, K. Simonyan, R. Hadsell, N. Robinson, E. Clancy, A. Arribas, and S. Mohamed, “Skilful precipitation nowcasting using deep generative models of radar,” *Nature* **597**, 672–677 (2021).
- ⁶⁵M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *J. Comput. Phys.* **378**, 686–707 (2019).
- ⁶⁶C. Meng, R. Rombach, R. Gao, D. P. Kingma, S. Ermon, J. Ho, and T. Salimans, “On distillation of guided diffusion models,” [arXiv:2210.03142](https://arxiv.org/abs/2210.03142) (2022).
- ⁶⁷C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, “Array programming with NumPy,” *Nature* **585**, 357–362 (2020).
- ⁶⁸S. Hoyer and J. Hamman, “Xarray: N-D labeled arrays and datasets in Python,” *J. Open Res. Software* **5**, 10 (2017).
- ⁶⁹J. D. Hunter, “Matplotlib: A 2D graphics environment,” *Comput. Sci. Eng.* **9**, 90–95 (2007).
- ⁷⁰J. Jeong and F. Hussain, “On the identification of a vortex,” *J. Fluid Mech.* **285**, 69–94 (1995).
- ⁷¹A. Papoulis and S. Unnikrishna Pillai, *Probability, Random Variables and Stochastic Processes* (McGraw-Hill, 2002).
- ⁷²C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2017), p. 4681–4690.