

This is a postprint version of the following published document:

Güemes, A., Sanmiguel Vila, C., & Discetti, S. (2022). Super-resolution generative adversarial networks of randomly-seeded fields. *Nature Machine Intelligence*, 4(12), 1165-1173.

DOI: [10.1038/s42256-022-00572-7](https://doi.org/10.1038/s42256-022-00572-7)

© 2022, The Author(s), under exclusive licence to Springer Nature Limited

# Super-resolution GANs of randomly-seeded fields

Alejandro Güemes<sup>1</sup>, Carlos Sanmiguel Vila<sup>1,2</sup>, Stefano Discetti<sup>1</sup>

*1: Aerospace Engineering Research Group, Universidad Carlos III de Madrid, Leganés, Spain*

*2: Sub-Directorate General of Terrestrial Systems,*

*Spanish National Institute for Aerospace Technology (INTA), San Martín de la Vega, Spain*

Reconstruction of field quantities from sparse measurements is a problem arising in a broad spectrum of applications. This task is particularly challenging when the mapping between sparse measurements and field quantities is performed in an unsupervised manner. Further complexity is added for moving sensors and/or random on-off status. Under such conditions, the most straightforward solution is to interpolate the scattered data onto a regular grid. However, the spatial resolution achieved with this approach is ultimately limited by the mean spacing between the sparse measurements. In this work, we propose a super-resolution generative adversarial network (GAN) framework to estimate field quantities from random sparse sensors without needing any full-field high-resolution training. The algorithm exploits random sampling to provide incomplete views of the high-resolution underlying distributions. It is hereby referred to as RAndomly-SEEDed super-resolution GAN (RaSeedGAN). The proposed technique is tested on synthetic databases of fluid flow simulations, ocean surface temperature distributions measurements, and particle image velocimetry data of a zero-pressure-gradient turbulent boundary layer. The results show excellent performance even in cases with high sparsity or with levels of noise. To our knowledge, this is the first GAN algorithm for full-field high-resolution estimation from randomly-seeded fields with no need of full-field high-resolution representations.

Sparse observations are often the only option for geophysical [1], astrophysical [2], biological [3], or fluid mechanical systems characterization [4–6]. Meteorological or oceanographic flows are remarkable examples in which atmospheric pressures, temperatures, and wind are only measured at a limited number of stations [7]; cardiac blood or flow control applications [8, 9] suffer similar issues, being sensor intrusiveness and noise contamination particularly relevant.

Various approaches to obtain full-field description from incomplete sparsely-seeded measurements can be followed depending on the sparsity level and the number of the available samples. For limited gappyness and a sufficiently large number of samples, reconstruction techniques based on extracting flow features using methodologies such as the Proper Orthogonal Decomposition (POD) [10] or the Dynamic Mode Decomposition (DMD) [11] can be applied. For instance, the Gappy-POD filling procedures [12–14] estimate non-gappy POD modes and time coefficients by least-squares regression on the available sparse fields; the estimated modes are later used to calculate the reconstructed full fields. A recent approach, proposed in [5], merges Gappy-POD and ensemble-particle-averaging techniques, paving the way to the reconstruction of fields with significantly higher gappyness fraction. The main advantage is that a high-resolution dictionary is not required beforehand and is constructed directly using low-resolution spatially-averaged fields and sparse fields. On the downside, a relatively high number of samples is required to reach suitable convergence. For cases with a limited number of samples, compressed-sensing (CS) methods provide efficient reconstruction using sparse representations [4, 9, 15]. CS is based on solving an under-determined linear system, which relates the sparse measurements with the entire fields, with additional constraints typically enforcing the smoothness of the solution [4]. A significant drawback is that they rely on a linear mapping between the sparse measurements and the high-resolution flow field. Therefore, in the cases where the relationship is strongly nonlinear or when a high signal-to-noise ratio corrupts the available measurements, the use of these methods is limited to cases with relatively low gappyness fraction.

In the last years, neural-network-based approaches have emerged to overcome these limitations [6, 7, 16–21]. Neural networks have been proved to be a powerful non-linear mapping tool. Super-resolution architectures based on generative adversarial networks (GANs) have been successfully applied in turbulent flows [22–25] and climate data [26] to enhance the spatial resolution and/or fill gaps between sparse measurements. While the success of these works is undeniable, they are all limited by the need for a non-sparse target for training, i.e., a library of examples where the entire field is required to train the reconstruction algorithm. This requirement often limits their applicability in real cases, in which the impossibility to record high-resolution measurements for training is more the rule than the exception.

In response to this problem, a novel super-resolution GAN framework to estimate field quantities without needing full-resolution fields for training is presented. The main advantage of the presented methodology is that it only requires the information of randomly located sensors at a variable location from sample to sample. This situation is quite common in particle-based measurements, such as tracers in particle tracking velocimetry (PTV) or phosphor thermography, but also in the case of moving sensors, such as buoys or ships. Low-resolution snapshots created by a simple

binning procedure and high-resolution sparse snapshots are used for the training process, making it unnecessary to have complete high-resolution fields as the target for training. Since the proposed framework exploits the random sampling to provide incomplete views of the underlying high-resolution fields, the method is named RANdomly-SEEDed super-resolution GAN (RaSeedGAN). A set of challenging validation cases which include numerical simulations of a fluidic pinball flow, a turbulent channel flow and data from the NOAA sea surface temperature database and a zero-pressure-gradient (ZPG) turbulent boundary layer particle image velocimetry (PIV) experiment is employed to test the robustness of the methodology in complex cases where nonlinearities and noise are present.

### The Randomly-Seeded super-resolution GAN

Generative adversarial networks are composed of two competing networks: a generator that produces an artificial output that mimics reality, and a discriminator, which oversees distinguishing between reality and artificial outputs. For this work, the architecture proposed in [27] has been used as a baseline. The RaSeedGAN architecture is sketched in Figure 1, and briefly outlined here. For further details on the architecture and on the training settings see the Methods section.

The generator is fed with full-field information at low-resolution, directly obtained from the scattered data, and provides in output full high-resolution fields. Several options can be explored to generate the low-resolution fields, including interpolation, Gaussian process regression, or Voronoi tessellation, among others. In this work, we adopt a simple binning procedure. Each bin is a compact support function whose value is assigned as the average quantity measured by the sensors within it. The bin size is defined to have approximately 10 sensors per bin on average. This approach has several advantages, such as noise robustness, spatial resolution uniformity, and it generates an output similar to techniques used for field quantity estimation based on particle-image measurements, such as particle image velocimetry/thermometry.

Real and generated fields are then fed into the discriminator network. It must be noted that the real fields are sparsely sampled, thus the full-field information provided by RaSeedGAN is masked with 0s in the empty bins of the reference (we refer to it as “high-resolution”).

The target is built from the scattered distributions, using a binning approach similar to the proposed above. Reducing the bin size will progressively increase the probability of finding empty bins, i.e., bins without sensors. The error during training is weighted to consider only those bins in the target carrying data.

Random sampling in space, with different distributions among realizations, is required to obtain a complete mapping of the domain of interest, i.e. RaSeedGAN requires random moving sensors covering the region of interest to be reconstructed. A single RaSeedGAN trained model can handle different sparsity levels, but the sparsity level would influence the quality of the obtained reconstructions, similarly to other methods such as cubic interpolation. In particular, in regions where there is not enough information to calculate the values corresponding to a bin location for the equivalent Low Resolution image throughout the dataset, there will be a blanked region, as in methods such as Gappy POD. Furthermore, when stacking quantities in the input of RaSeedGAN using channels in the neural network, such as can be the different velocity components  $U$ ,  $V$ , their sparsity level should be similar. Better results are obtained by training a single RaSeedGAN model per quantity of interest in case of significant difference in sparsity level, although at the expense of the computational cost.

## RESULTS

### Test cases outline

In this section, the RaSeedGAN algorithm is applied on synthetic test cases generated from Direct Numerical Simulation (DNS) of fluid flows and on real experimental data. Two synthetic test cases with different challenges are considered:

- velocity data of the wake of a cluster of three cylinders with equal diameter  $D$ , referred as fluidic pinball [28]. The wake of the pinball includes the interaction of the wakes of three bodies, a region of development and the final merging in a large-scale shedding wake. The domain is discretized with  $25\text{pixels}/D$ . The fields are seeded with 0.01 vectors per pixels. The low-resolution input is built setting the binning region to  $32 \times 32$  pixels with 50% overlap, corresponding to approximately 10 vectors per bin. Beware that these parameters aim to simulate the typical processing of PIV, in which the measured velocity closely relates to the average velocity measured within the interrogation window (here referred to as bin).
- a turbulent channel flow, generated from a dataset available at the Johns Hopkins Turbulence Database (<http://turbulence.pha.jhu.edu/>). Sub-domains are taken of size  $2h \times h$  in the streamwise and wall-normal directions,

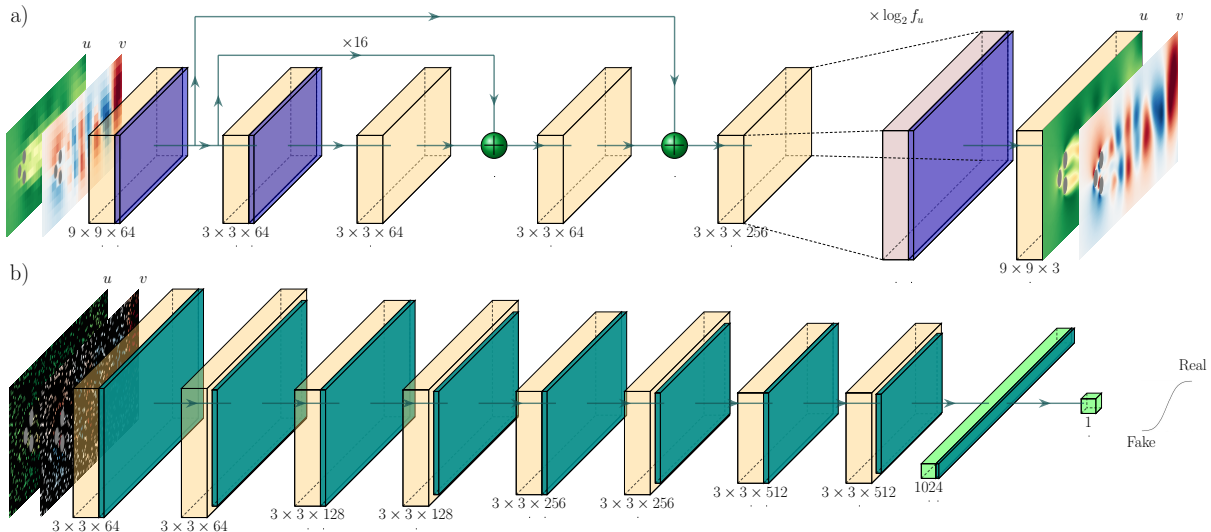


FIG. 1: Schematic illustration of RaSeedGAN architecture. a) Generator. b) Discriminator. The novel element is the use for training and testing of gappy high-resolution fields provided directly by binned scattered data.

where  $h$  is the half-channel height. The domains are discretized with 512 pixels/ $h$ . The same vector density and bin averaging size/procedure of the first test case are adopted.

In the synthetic test cases, the binned distributions for training are generated using directly the original positions on the particles (or point sensors) rather than performing particle tracking on virtual images. Noise in the measurement is introduced by adding a random perturbation on velocity vectors of realistic intensity. For the pinball test case, two Gaussian noise conditions are added to the displacement vectors to study the robustness of the method against noisy measurements, with standard deviation in pixels of 0.01 (moderate intensity, Test Case 1a) and 0.05 (high intensity, Test Case 1b), respectively. For the channel, a higher noise is introduced (0.1 pixels) to account for the higher freestream displacement if compared to the pinball case. In both cases, the error realistically represents noise perturbation of Particle Tracking Velocimetry data. With this approach, we attempt to isolate the errors due to spatial resolution limitations while maintaining a controlled random error level due to uncertainty in feature matching.

Additionally, two complex experimental test cases are chosen:

- temperature data from the NOAA sea surface temperature (<http://www.esrl.noaa.gov/psd/>), based on temperature data collected by satellites and ships. The temperature data are randomly seeded to provide artificial point sensors, while the original data are used for assessment.
- in-house experimental PIV data of a zero-pressure-gradient (ZPG) turbulent boundary layer (TBL). In order to have a ground truth available for comparison, the low-resolution field is estimated on a fictitiously large window of  $128 \times 128$  pixels with 50% overlap with a multi-step PIV interrogation [29]. Then an analysis using a window with  $32 \times 32$  pixels is used to generate a high-resolution reference. It must be remarked here that these high-resolution fields are not used in any form in the training process and are just exploited for performance assessment. The number of vectors is then artificially reduced by a factor of 10 by randomly dropping sensors, and fed to RaSeedGAN to test the super-resolution capabilities.

A more detailed description of the test cases is reported in the section Methods.

The RaSeedGAN seeks to achieve upsampling levels ranging from  $f_u = 2 - 8$ , with  $f_u$  being the upsampling factor. The window used to bin the vector fields has a size equal to  $b = D_I/f_u$ , with  $b, D_I$  being the bin size of the binned data and the bin size for the low-resolution data, respectively. Furthermore, the grid is refined by the same amount, thus maintaining the overlap between adjacent bins. In all cases,  $D_I$  is set in order to have on average 10 sensors per bin, thus ensuring negligible probability of empty bins within the low-resolution input. It is important to underline that selecting a smaller bin  $b$  increases the levels of gappyness of the distributions to be reconstructed.

It must be remarked that, in all test cases, the position of the sparse sensors on each individual snapshot is fixed during training, i.e. each individual field is sampled at a predetermined finite number of locations which does not change during training. The positions of the sensors, on the other hand, change from snapshot to snapshot. This

TABLE I: Root mean squared error for each test case, scaled with the spatially-averaged standard deviation of their corresponding fluctuating quantities. Cases correspond to: (1a) Pinball with 1% noise disturbance, (1b) Pinball with 5% noise disturbance, (2) Turbulent Channel Flow with 1% noise disturbance, (3) NOAA sea surface data and (4) Experimental Turbulent Boundary Layer. In bold are highlighted best-performing method for each test case.

Method	$f_u$	Test Case 1a		Test Case 1b		Test Case 2		Test Case 3	Test Case 4 <sup>a</sup>	
		$u$	$v$	$u$	$v$	$u$	$v$	$T$	$u$	$v$
RaSeedGAN		0.261	0.243	<b>0.217</b>	<b>0.200</b>	<b>0.117</b>	<b>0.303</b>	0.058	<b>0.081</b>	<b>0.316</b>
Cubic	2	<b>0.218</b>	<b>0.171</b>	0.241	0.232	0.126	0.337	<b>0.054</b>	0.176	0.605
Gappy POD		0.274	0.243	0.287	0.277	0.137	0.373	0.063	0.172	0.610
RaSeedGAN		0.169	0.182	<b>0.137</b>	<b>0.157</b>	<b>0.116</b>	<b>0.307</b>	0.055	<b>0.071</b>	<b>0.275</b>
Cubic	4	0.219	<b>0.177</b>	0.241	0.237	0.124	0.337	0.058	0.175	0.590
Gappy POD		<b>0.163</b>	0.191	0.180	0.224	0.163	0.529	<b>0.050</b>	0.171	0.649
RaSeedGAN		0.133	<b>0.165</b>	<b>0.098</b>	<b>0.139</b>	<b>0.117</b>	<b>0.315</b>	0.056	-	-
Cubic	8	0.220	0.178	0.242	0.237	0.124	0.337	0.055	-	-
Gappy POD		<b>0.125</b>	0.184	0.135	0.200	0.185	0.635	<b>0.049</b>	-	-

<sup>a</sup> In Test Case 4 Gappy POD has been trained on a smaller dataset (5000 images) for computational cost limitations. However, only marginal improvements have been observed when increasing the dataset size.

framework is intended to replicate realistic scenarios of optical particle-based measurements, or of field measurements with moving point sensors.

### A. Validation on synthetic test cases

For all validation cases, we explore the performances of RaSeedGAN with instantaneous visualisations of reconstructed fields. Figure 2 shows an instantaneous realisation for the synthetic test cases considered under study, including in columns: the input low-resolution (LR) field, the binned field, referred as ‘‘Sparse HR (high-resolution) Reference’’, the reconstruction achieved with the RaSeedGAN algorithm and the reference field used to generate the particle distributions. The complete high-resolution data is included in the last column of Figure 2 as a quality check, although it is not used at any moment during training. Furthermore, Table I reports the pixel-based root mean squared error for each test case and upsampling levels compared to the error obtained by a standard cubic interpolation method and the Gappy POD, following the implementation from Ref. [13]. The error is normalised to the standard deviation of the corresponding quantity.

In the first row of Figure 2, an instantaneous flow visualisation of the fluidic pinball case with upsampling factor  $f_u = 8$  is reported for both components of the velocity fields. The RaSeedGAN allows recovering correctly the developing regions around the cylinders in the near-wake and around the cylinders. For instance, the crosswise-velocity details in the front of the cylinders, which are barely seen in the low-resolution input fields, are nearly perfectly recovered. Similar performance is observed in the pinball wake, where the RaSeedGAN reconstructs the streamwise momentum defect underestimated by the low-resolution fields. The error level reported in Table I for cases with upsampling levels of  $f_u = 2 - 8$  supports the quality of the overall reconstructed fields.

Table I shows that the reconstruction errors for the Test Case 1a (1% noise level) for cases with upsampling levels of  $f_u = 4 - 8$  are significantly lower than the standard cubic interpolation, although the Gappy POD still appears to be superior due to the high energy content of the first POD modes of the pinball case. However, when the noise conditions are higher, as in Test Case 1b where a 5% noise level is included, we observe that the accuracy of the cubic interpolation and the Gappy POD is degraded while RaSeedGAN can improve even further its performance. RaSeedGAN, like other GAN architectures, appears to be able to explore better and generate from the underlying data distribution in presence of noise [30, 31]. This is relevant when applying RaSeedGAN to experimental datasets where the noise level cannot be easily controlled. Additionally, it has to be remarked that the current architecture can be further improved using more data or layers/filters with respect to the baseline RaSeedGAN, at expense of a higher computational cost. As an example of this capability, a test doubling the number of filters of the RaSeedGAN generator showed that for Test Case 1a and  $f_u = 2$  the pixel-based root mean squared error decreases up to  $u = 0.215$  and  $v = 0.186$ .

In the second row of Figure 2, the streamwise and wall-normal velocity contours for the turbulent channel flow simulation are represented. Alongside this, contours that identify extreme Reynolds-shear-stress events - the so-called *wusters* as defined in [32] - have been plotted. The *wusters*, together with the clusters of vortices, play a crucial role in the modelling of the dynamics of wall-bounded turbulence based on coherent structures [32]. The recovery of wall-attached vertical velocity fluctuations in the near-wall region is remarkable, where even the recovery of small

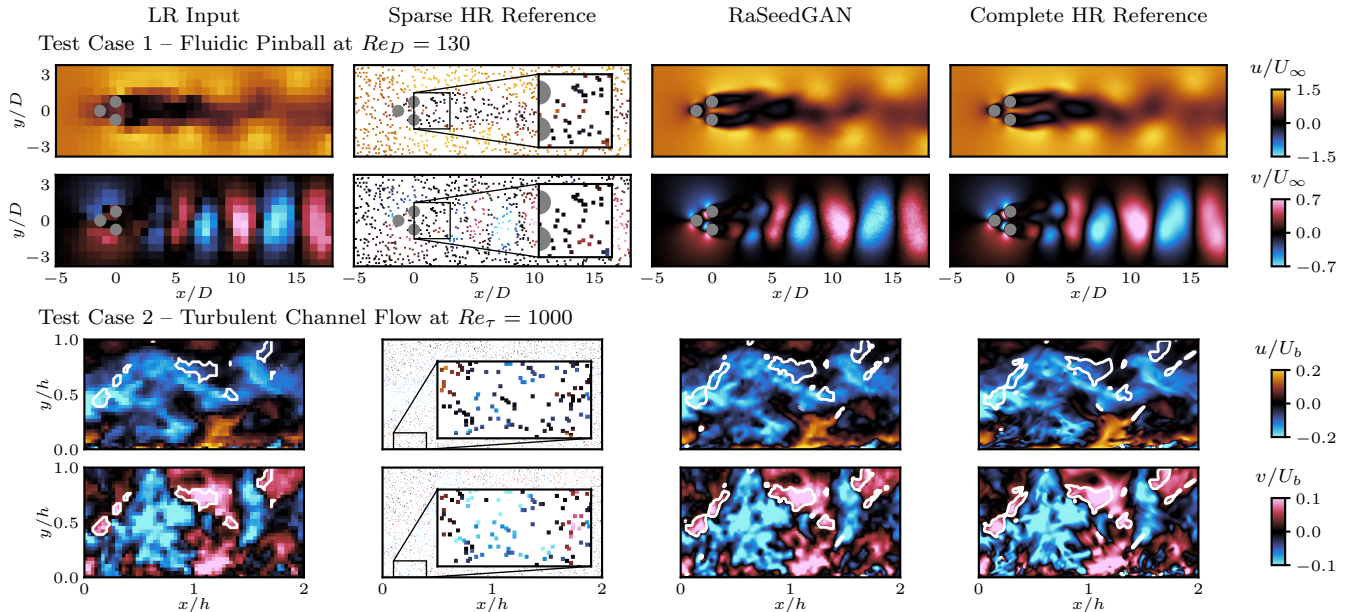


FIG. 2: Panel of synthetic test cases, with examples of instantaneous field reconstructions. Lower-case letters  $u$  and  $v$  indicate respectively the streamwise and crosswise velocity fluctuation components. The reference velocity  $U_\infty$  and  $U_b$  are the freestream and the bulk velocity for the pinball and the channel flow test case, respectively.  $Re_D$  and  $Re_\tau$  indicate the Reynolds number based on the cylinder diameter  $D$ , and the friction-based Reynolds number. The columns indicate, from left to right: low-resolution input field; binned distribution with  $f_u = 8$ ; reconstruction using the proposed sparsely-trained super-resolution GANs; the reference full field. From top to bottom: wake of the fluidic pinball from DNS, streamwise and crosswise velocity fluctuation components; turbulent channel flow from DNS, wall-parallel and wall-normal velocity component. White isolines in Test Case 2 indicate *wusters*. For clarity, an inset view has been included in the Sparse HR Reference to show how the sparse images are represented.

structures is appreciable. Regarding the outer part of the flow, the *wusters* are largely improved. Analysing the error reconstruction of Table I, the RaSeedGAN still outperforms the standard cubic interpolation, although in this case to a lower extent due to the richer range of spatial scales involved. In this complex case, Table I shows that RaSeedGAN outperforms the standard cubic interpolation and the Gappy POD for all the reported upsampling levels.

## B. Application on experimental test cases

In this section an assessment of RaSeedGAN on the two experimental test cases is conducted. In Figure 3, an instantaneous temperature field visualisation of the NOAA sea surface temperature with upsampling factor  $f_u = 8$  is reported. The RaSeedGAN allows recovering the temperature fields with a higher level of detail, as is evident from the reconstructed silhouette of the continents of the Earth. For this case the performances of RaSeedGAN seem in line with the cubic interpolation, and slightly inferior to Gappy POD. These can be ascribed to the seasonality of the data, which delivers a very compact POD eigenspectrum. However, from visual inspection and analysing the temporal spectra, it has been observed that the noise level of the RaSeedGAN is slightly lower compared with the Gappy POD and the standard cubic interpolation.

In Figure 4, the streamwise and wall-normal velocity contours for the experimental data of ZPG TBL flow are presented. As for Test Case 2, the contours that represent the *wusters* have also been plotted. It has to be remarked that this particular case is an actual experiment in which the noisy conditions are natural and not simulated nor imposed. Again, as in the turbulent channel flow case, the recovery of wall-attached vertical velocity fluctuations in the near-wall region is observed even though the quality of the reconstruction does not match the one reported in Test Case 2. This result is not surprising since the quality of the experimental data is lower than the synthetic data. This statement is supported by the performance of the standard cubic interpolation and the Gappy POD reported in Table I, being the difference of the RaSeedGAN and both methods more considerable than in Test Case 2.

As a further indicator of the quality of the reconstruction achieved with RaSeedGAN, the wall-normal profiles of first and second-order statistics of the streamwise velocity are reported in Figure 5a. The quantities are presented in

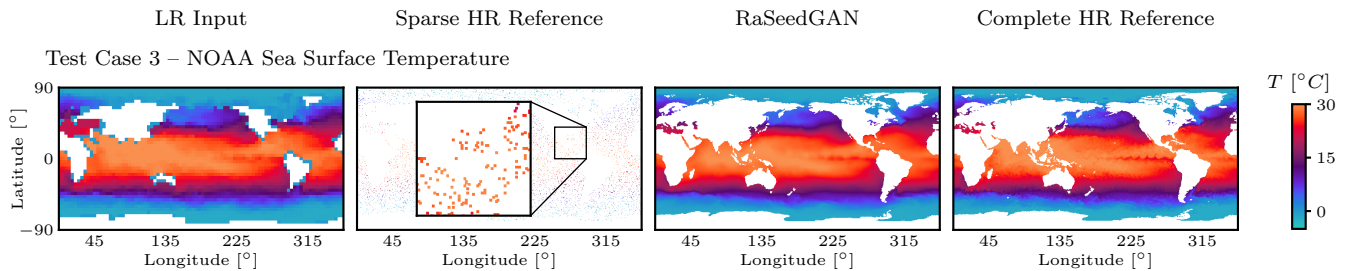


FIG. 3: Measurement of ocean surface temperature from NOAA sea surface data. From left to right: low-resolution input field; binned distribution with  $f_u = 8$ ; reconstruction using the proposed sparsely-trained super-resolution GANs; the reference full field. For clarity, an inset view has been included in the Sparse HR Reference to show how the sparse images are represented.

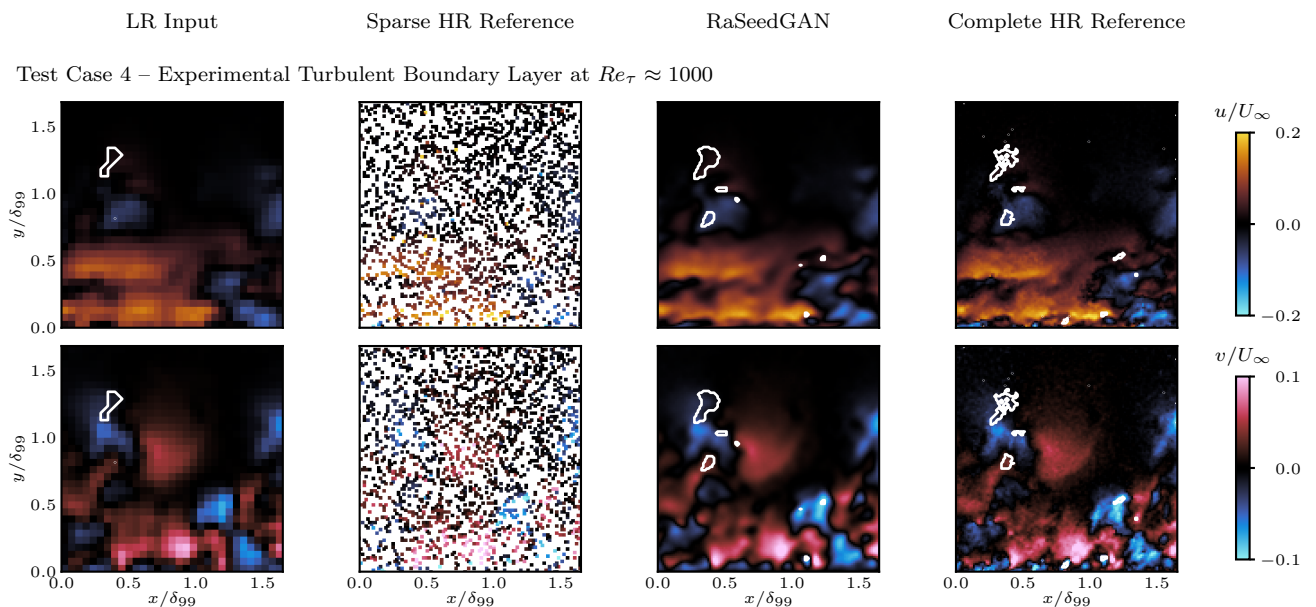


FIG. 4: The columns indicate, from left to right: low-resolution input field; binned distribution with  $f_u = 4$ ; reconstruction using the proposed sparsely-trained super-resolution GANs; the reference full field. The quantities  $\delta_{99}$  and  $U_\infty$  indicate respectively the boundary layer thickness and the freestream velocity. From top to bottom: experimental data of a turbulent boundary layer, wall-parallel and wall-normal velocity fluctuating components. Black isolines indicate *uvsters*.

inner scaling, i.e. the mean velocity  $U$  and the standard deviation of the fluctuation  $u'$  are scaled with the friction velocity  $u_\tau$ , while the wall-normal coordinate  $y$  is scaled with the viscous length scale  $\ell^*$ . Inner-scaled quantities are indicated with the superscript  $+$ . The results are obtained with ensemble averaging over the testing dataset and spatial averaging along the streamwise direction. For Test Case 2, the reference profiles are directly extracted from the simulation data in the Johns Hopkins Turbulence Database. For Test Case 4, the reference profile from the high-resolution PIV analysis is reported. For completeness, profiles from a friction-Reynolds-number matched simulation are also included [33]. It is important to remark that the network is not trained to recover flow statistics but rather to reconstruct field features. In this scenario, high-resolution statistics can be obtained with higher fidelity using directly ensemble averaging approaches on the scattered data. Furthermore, physics-based constraints are not included during training, i.e. RaSeedGAN in the present implementation is unaware of the underlying physics. Even though physics can be easily embedded in the training process (for instance as an additional content loss contribution based on residual from governing equations), this is not explicitly enforced in the analysis. Consequently, this assessment can

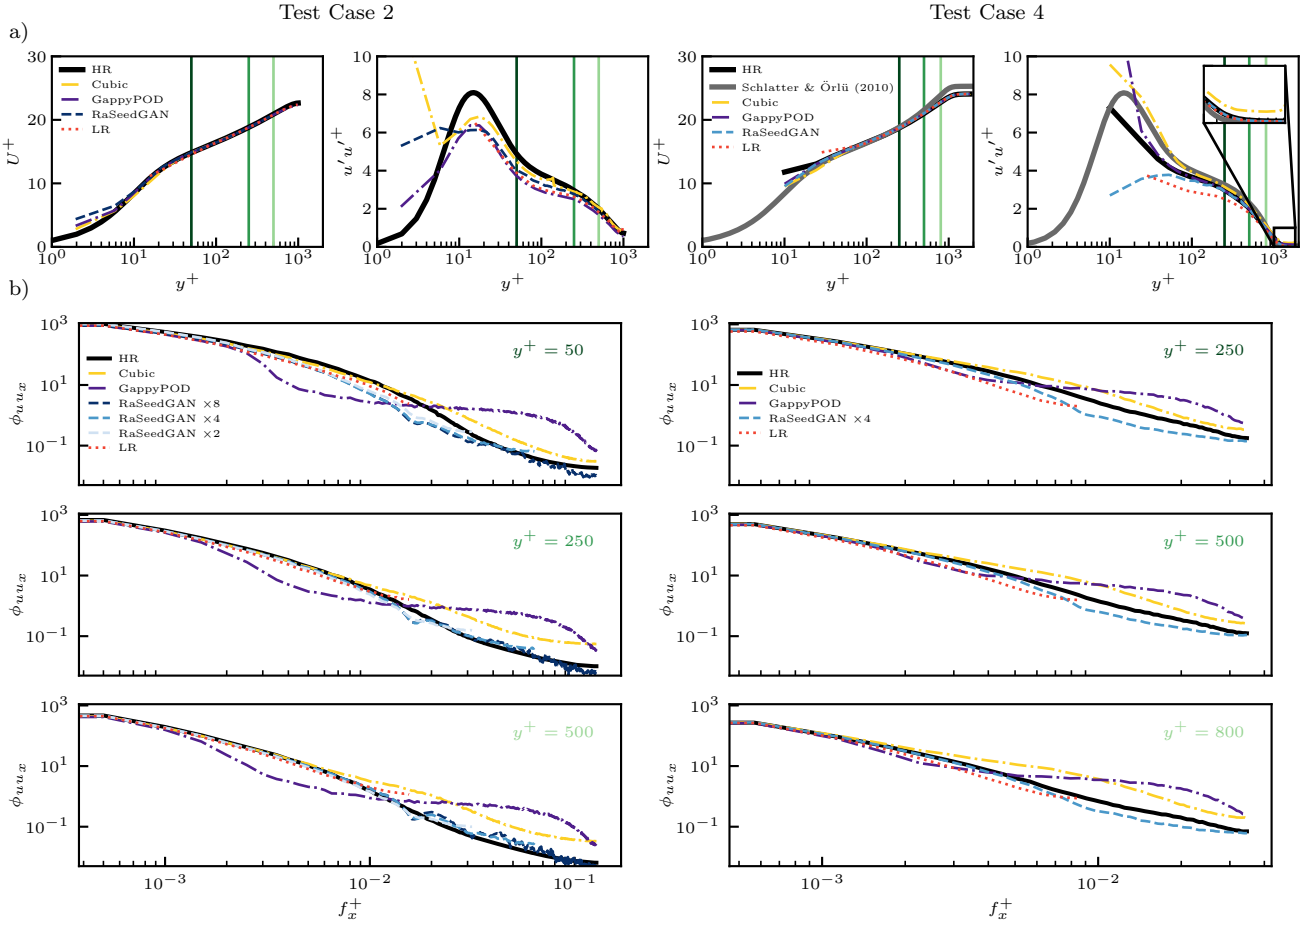


FIG. 5: a) Inner-scaled mean streamwise velocity and variance profiles for Test Case 2 (DNS turbulent channel flow) and Test Case 4 (experimental turbulent boundary layer). Lines represent the data used to calculate the quantity: high-resolution reference (solid black), low-resolution original measurements (dotted red), cubic interpolation (yellow dashed-dotted), Gappy POD (purple dashed-dotted) and RaSeedGAN reconstruction (dashed blue, from light to dark indicating  $f_u = 2, 4, 8$ ). The inset in the variance profile of Test Case 4 provides a magnified view of the region with wall distance above 1000 inner units. Vertical green lines (from dark to light for increasing distance from the wall) indicate stations where the streamwise turbulent spectra are computed. For Test Case 4, gray profiles refer to data from [33]. b) Streamwise turbulent spectra at different wall distances.

be considered a fair metric for faithful flow reconstruction.

Although RaSeedGAN seems to improve the turbulence statistics profiles of the initial low-resolution data, it is observed that the cubic interpolation seems to have better performances, especially in the inner region, i.e. below 15 wall units (which corresponds to about 8 pixels in Test Case 2, and 16 pixels in Test Case 4). This result, apparently at odds with the error statistics of Table I, is not surprising since it is indicative that RaSeedGAN might inherit part of the intensity modulation error from the low-resolution field, while cubic interpolation works directly on the pointwise data. On the other hand, fields obtained from cubic interpolation have visibly higher noise contamination. Thus, the variance estimation might be biased towards higher intensity due to noise, partly compensating for modulation issues and providing a misleading offset towards the reference profiles. This offset effect is well known in the fluid mechanics' community [34], and evidences of it can be observed at wall distances larger than 1000 wall-units for Test Case 4, where a plateau of the variance is observed for the cubic interpolation data to a higher value than the reference and the RaSeedGAN data (see inset of Figure 5a). To better assess the resolution in terms of scales of the flow and discern the noise contamination, the streamwise spectra of the velocity fluctuations are computed at different distances from the wall (see Figure 5b). It is observed that, except for small wall distances, RaSeedGAN can significantly increase the reach of the spectrum and follow with high fidelity the high-resolution reference. In all cases, the cubic interpolation



data deliver a spectrum that peels away from the reference already at a relatively large scale. If considering as cut-off the frequency at which the spectrum of the low-resolution data and RaSeedGAN reach 80% of the intensity of the reference spectrum, scales of at least 3 times smaller are recovered in all cases of Figure 5 except for  $y^+ = 50$  of Test Case 2.

## CONCLUSIONS

A novel deep-learning approach based on generative adversarial networks to perform super-resolution reconstruction of sparse measurements has been proposed and assessed. The method exploits the ability of the neural networks to learn mapping functions even when some layers have randomly-ignored neurons during the training process. We propose the RaSeedGAN method based on a GAN architecture that can obtain a non-linear mapping to high-resolution reconstructed images, using as inputs sparse low-resolution images obtained by binning the randomly-distributed sensor inputs, and as targets the original incomplete sparse measurements. The proposed technique requires randomly-moving sensors from snapshot to snapshot to establish a mapping in the domain of interest. This situation is common, for instance, in particle-based optical measurements.

RaSeedGAN allows the user great flexibility for application to experimental data in configurations where high-resolution benchmarks are not available. Its main novelty is that a full field is not required beforehand and is constructed directly using gappy binned images and the original sparse measurements. Different examples that include numerical and experimental data from velocity and temperature measurements were successfully tested, demonstrating the accuracy and robustness of the proposed method even in the presence of noisy data. Even though RaSeedGAN is not explicitly trained to fulfil physical constraints, evidence of consistency with physics is found from the analysis of derived quantities, such as turbulence statistics. Furthermore, the robustness of the proposed architecture is proved in a range of cases, although the user might still find room for improvement with the tuning of the model parameters and of the loss function to minimize during training. We foresee further performance improvements from the embedding of physics laws and/or models within the training process.

## Methods

### RaSeedGAN implementation details

The RaSeedGAN architecture, sketched in Figure 1, is based on the one proposed in [27]. However, the batch-normalization layers have been removed to ensure a proper fitting of the deep layers in the architecture as proposed by [35]. The generator is fed with full-field information at low-resolution, and is trained to provide in output full high-resolution fields with fixed upsampling.

The low-resolution fields are fed to the generator, which applies a convolutional layer with filter size  $9 \times 9$  and 64 feature maps, followed by a parametric ReLU (PReLU) activation function. After these initial layers, 16 residual blocks are applied with the layout proposed by [36], i.e., a convolutional layer followed by a PReLU activation function and a second convolutional layer, with 64 feature maps of filter size  $3 \times 3$  for both convolution operations. Before increasing the resolution, a skip-connection sum is performed between the residual blocks' output and the initialization layers' output. The subpixel convolution layer proposed by [37] is used to increase the resolution. Finally, a convolution operation with linear activation function and filter size  $9 \times 9$  is applied to recover the in-plane high-resolution measurements as output.

The discriminator network is then fed with real (sparse) and generated fields. The initialization of the network is carried out by a convolution operation of filter size  $3 \times 3$  and 64 feature maps, followed by the Leaky ReLU (LReLU) activation function. After this, 7 discriminator blocks of successive increasing feature maps are applied, where the discriminator blocks are composed of convolutional layers of  $3 \times 3$  size and LReLU activation function. Note that odd blocks applied a stride step larger than one to reduce the height and width of the feature maps. Finally, the feature-map tensor is flattened into a vector, followed by a fully-connected layer of 1024 neurons with LReLU activation function and output fully-connected layer with one neuron and sigmoid activation function. The output of the discriminator provides the probability of whether the input is real (0s) or fake (1s).

The network architectures are implemented in the open-source framework Tensorflow [38]. All models are randomly initialized. Networks' weights are updated with a learning rate of  $\epsilon = 0.0001$ .

The discriminator loss function is defined as:

$$\mathcal{L}_D = -\mathbb{E}[\log D(H_R)] - \mathbb{E}[\log(1 - D(F_V \otimes G(L_R)))], \quad (1)$$

where  $\mathbb{E}[\ ]$  represents the operation of taking average in the mini-batches,  $H_R$  indicates high-resolution images,  $L_R$  low-resolution images,  $D(\ )$  is the discriminator network that takes as input matrices with dimensionality equal to the final high-resolution grid and  $F_V$  is the matrix that indicates if a bin has a sensor (1) or not (0).

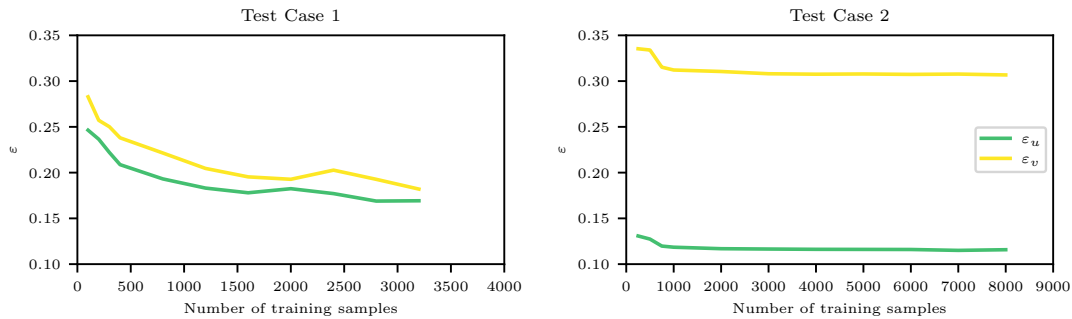


FIG. 6: Pixel-based root mean square error as a function of the number of training samples for Test Case 1 (left) and Test Case 2 (right) with  $f_u = 4$ . The error is scaled with the standard deviation of the corresponding fluctuating quantities. Green and yellow lines indicate respectively the error on the streamwise and crosswise velocity components.

The loss function of the generator network is defined as:

$$\mathcal{L}_G = \sum_{i=1}^{Nx} \sum_{j=1}^{Nz} |H_{R_{i,j}} - F_V \otimes G(L_R)_{i,j}|^2 + \lambda \mathcal{L}_A, \quad (2)$$

where  $G()$  is the generator network that takes as input matrices with dimensionality equal to the low-resolution data and  $\lambda$  is a scalar parameter set to 0.001 to weigh the contribution of the adversarial loss  $\mathcal{L}_A$ , defined as:

$$\mathcal{L}_A = -\mathbb{E}[\log(D(F_V \otimes G(L_R)))]. \quad (3)$$

The adversarial loss is defined as the binary cross-entropy of the predicted fields, i.e., it checks whether the discriminator has labelled as ‘fake’ the generated fields. The discriminator loss is the mean value between the binary cross-entropy of the predicted fields and the target ones. The mask  $F_V$  is formally equivalent to a dropout layer at the last stage of the networks with random activations.

The test cases included in this work have been trained during 100 epochs with 8 samples per batch. The generator and discriminator states have been stored every 5 epochs to find the optimum point before the GAN starts overfitting. The weights of both networks have been updated using Adam optimizer with a learning rate of 0.0001. All the cross-entropy errors have been perturbed with a random fluctuation of standard deviation 0.2 to increase the stability of the GAN training.

It must be remarked that we used for training a relatively large number of samples to eliminate doubts on the training convergence. Nonetheless, the performances seem to be affected only to a small extent when working on reduced-size datasets, as observed in Figure 6. Test Case 1 involves a narrower range of scales and a less complex behaviour; this allows a continuous improvement of the error as the number of samples increases. On the other hand, due to the more complex behaviour of moderate-Reynolds-number channel flow, there seems to be no significant gain in increasing the number of samples beyond 1000 for Test Case 2. It has to be remarked that the same neural network architecture has been employed for the training of every test case, and while there are complex cases such as the turbulent channel that may benefit from complex neural network architectures, it is left to the user to decide what level of complexity is required for each application.

### Synthetic test case settings

Test Case 1 consists of a *fluidic pinball* [28], a two-dimensional wake flow around a cluster of three equidistantly-spaced cylinders with equal radius  $R = D/2$ , whose centres form an equilateral triangle with a side length equal to  $3R$ . The triangle is oriented with an upstream vertex and the downstream side orthogonal to the freestream flow, located at  $x = 0$  and centred on the  $y$ -axis. The DNS data at  $Re_D = 130$  (i.e. within the chaotic regime [28], with  $Re_D$  being the Reynolds number) are used to generate synthetic PIV data. The details of the simulation settings and flow behaviour can be found in [28].

A set of 4737 fields with scattered vectors are generated. The simulation covers approximately 9000 convective times (defined as  $D/U_\infty$ , with  $U_\infty$  being the freestream velocity). This corresponds to about 1250 shedding cycles.

The snapshots have a fixed time separation of 1.9 convective times, i.e. approximately 0.27 shedding cycles. This allows covering correctly all phases of the shedding. The first 4000 fields are used for training, while the rest is used for testing. The domain spans from  $-5D$  to  $18.04D$  in the streamwise direction and  $-3.84D$  to  $3.84D$  in the spanwise direction. In order to simulate a typical PIV setup, we discretized the domain with 25 pixels/ $D$ , which results in images with  $576 \times 192$  pixels. Approximately 1100 vectors are generated for each sample, corresponding to 0.01 vectors per pixel. For this first test case, the RaSeedGAN seeks to achieve upsampling levels ranging from  $f_u = 2-8$ , starting from an initial bin size of  $D_I = 32$  pixels. The level of gappyness is approximately 11%, 57%, 87% for  $f_u = 2, 4, 8$ , respectively.

Test Case 2, generated from a dataset available at the Johns Hopkins Turbulence Database (<http://turbulence.pha.jhu.edu/>), is a turbulent channel flow with a dimension of 2 half-channel-heights  $h$  from wall to wall,  $3\pi h$  in the spanwise direction and  $8\pi h$  in the streamwise direction. The details of the simulation settings and flow behaviour can be found in [39]. Sub-domains of  $2h \times h$  in the streamwise and wall-normal directions have been discretized with 512 pixels/ $h$  and seeded with particles with a density of 0.01 particles per pixel. The particles have been randomly distributed in the subdomain and tracked for 10 time steps of the simulations to generate the position in the second frame. A time separation of 1 convective time between snapshots (based on the bulk velocity  $U_b$  and the half-channel height) is set to reduce statistical correlation between samples. In order to obtain a sufficient number of snapshots, data are extracted in subdomains at different locations in the streamwise and spanwise directions. The streamwise and spanwise separation between domains was equal to  $2h$  and  $0.25h$ , respectively. A total of 11856 snapshots have been generated, with 10000 used for training. As in the previous test case, the reference low-resolution dataset is built with a bin size of  $32 \times 32$  pixels with 50% overlap. Upsampling factors up to 8 have been tested. The gappyness level is 12%, 43%, 85% for  $f_u = 2, 4, 8$ , respectively.

### Experimental test case description

Test Case 3 is the NOAA sea surface temperature (<http://www.esrl.noaa.gov/psd/>), based on temperature data collected by satellites and ships. A total of 7305 temperature fields has been acquired, where the training data ranges from January 1<sup>st</sup> 2000 to June 4<sup>th</sup> 2016 and the testing set elapses from June 5<sup>th</sup> 2016 to December 31<sup>st</sup> 2019. The original dataset has been interpolated on a grid with  $720 \times 1440$  points. Point temperature measurements have been extracted at approximately 10300 locations per sample and used to generate the low-resolution training dataset and the sparse high-resolution target. For the low-resolution training dataset, a bin size of  $32 \times 32$  points has been selected, resulting in approximately 10 points on average per bin. Upsampling factors of  $f_u = 2, 4, 8$  have been explored. The gappyness level correspond to 36%, 69% and 90% for upsampling factors of  $f_u = 2, 4, 8$  respectively.

Test Case 4 is an experimental dataset of a turbulent boundary layer flow in a wind tunnel with Particle Image Velocimetry. The dataset comprises 39000 images, with a spatial resolution of 48.4 pixels/ $mm$ . The real density of identified vectors is 0.006 vector per pixel, which results in approximately 25000 vectors in an area of  $2048 \times 2048$  pixels. The flow fields have a domain size  $1.68 \times 1.68 \delta_{99}$  in the streamwise and wall-normal directions, respectively, where  $\delta_{99}$  is the boundary layer thickness. The input low-resolution field are estimated on a fictitiously large window of  $128 \times 128$  pixels with 50% overlap, while the reference high-resolution fields are obtained using a window with  $32 \times 32$  pixels. The number of vectors is then artificially reduced to only 2560 vectors per sample by randomly dropping sensors. This corresponds to 10 vectors in the original window of  $128 \times 128$  pixels, thus resulting in a mean vector spacing of 40 pixels. The sparse binned distributions for GANs training are obtained by binning with upsampling factors  $f_u = 2, 4$ , resulting in a gappyness level of 15% and 61%, respectively. Of the 39000 fields available, 30000 are used for training, while the rest are used for testing.

### Data availability

All datasets used in this work are openly available in Zenodo, accessible at [10.5281/zenodo.7191210](https://zenodo.org/record/105281)

### Code Availability

All codes developed in this work are openly available in GitHub, accessible through the link: <https://github.com/eaplab/RaSeedGAN>

### Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 949085). NOAA High Resolution SST data provided

by the NOAA/OAR/ESRL PSL, Boulder, Colorado, USA.

#### Author contributions

AG: Methodology, Software, Validation, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization. CSV: Conceptualization, Methodology, Writing - Original Draft, Writing - Review & Editing. SD: Methodology, Software, Resources, Data Curation, Supervision, Writing - Original Draft, Writing - Review & Editing, Funding acquisition

#### Competing interests

The authors declare no competing interests.

#### References

- 
- [1] T. Bolton and L. Zanna, Applications of deep learning to ocean data inference and subgrid parameterization, *J. Adv. Model. Earth Sy.* **11**, 376 (2019).
  - [2] K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, *et al.*, First M87 event horizon telescope results. III. Data processing and calibration, *Astrophys. J. Lett.* **875**, L3 (2019).
  - [3] A. Yakhot, T. Anor, and G. E. Karniadakis, A reconstruction method for gappy and noisy arterial flow data, *IEEE T. Med. Imaging* **26**, 1681 (2007).
  - [4] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns, *IEEE Contr. Syst. Mag.* **38**, 63 (2018).
  - [5] J. Cortina-Fernández, C. Sanmiguel Vila, A. Ianiro, and S. Discetti, From sparse data to high-resolution fields: ensemble particle modes as a basis for high-resolution flow characterization, *Exp. Therm. Fluid Sci.* **120**, 110178 (2020).
  - [6] 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).
  - [7] K. 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).
  - [8] H. Shen, X. Li, Q. Cheng, C. Zeng, G. Yang, H. Li, and L. Zhang, Missing information reconstruction of remote sensing data: A technical review, *IEEE Geosci. Remote Sens. Mag.* **3**, 61 (2015).
  - [9] J. L. Callahan, K. Maeda, and S. L. Brunton, Robust flow reconstruction from limited measurements via sparse representation, *Phys. Rev. Fluids* **4**, 103907 (2019).
  - [10] 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).
  - [11] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
  - [12] R. Everson and L. Sirovich, Karhunen–Loeve procedure for gappy data, *JOSA A* **12**, 1657 (1995).
  - [13] D. Venturi and G. E. Karniadakis, Gappy data and reconstruction procedures for flow past a cylinder, *J. Fluid Mech.* **519**, 315 (2004).
  - [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] X. Huang, Compressive sensing and reconstruction in measurements with an aerospace application, *AIAA J.* **51**, 1011 (2013).
  - [16] R. Maulik, K. Fukami, N. Ramachandra, K. Fukagata, and K. Taira, Probabilistic neural networks for fluid flow surrogate modeling and data recovery, *Phys. Rev. Fluids* **5**, 104401 (2020).
  - [17] 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).
  - [18] N. B. Erichson, L. Mathelin, Z. Yao, S. L. Brunton, M. W. Mahoney, and J. N. Kutz, Shallow neural networks for fluid flow reconstruction with limited sensors, *Proc. Roy. Soc. A* **476**, 20200097 (2020).
  - [19] 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).
  - [20] A. Arzani, J.-X. Wang, and R. M. D’Souza, Uncovering near-wall blood flow from sparse data with physics-informed neural networks, *Phys. Fluids* **33**, 071905 (2021).
  - [21] S. L. Brunton, B. R. Noack, and P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid Mech.* **52**, 477 (2020).
  - [22] K. Fukami, K. Fukagata, and K. Taira, Super-resolution reconstruction of turbulent flows with machine learning, *J. Fluid Mech.* **870**, 106 (2019).

- [23] H. Kim, J. Kim, S. Won, and C. Lee, Unsupervised deep learning for super-resolution reconstruction of turbulence, *J. Fluid Mech.* **910** (2021).
- [24] A. Güemes, S. Discetti, A. Ianiro, B. Sirmacek, H. Azizpour, and R. Vinuesa, From coarse wall measurements to turbulent velocity fields through deep learning, *Phys. Fluids* **33**, 075121 (2021).
- [25] 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).
- [26] K. Stengel, A. Glaws, D. Hettinger, and R. N. King, Adversarial super-resolution of climatological wind and solar data, *Proc. Natl Acad. Sci.* **117**, 16805 (2020).
- [27] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al., Photo-realistic single image super-resolution using a generative adversarial network, in *Proc. IEEE CVPR* (2017) pp. 4681–4690.
- [28] N. Deng, B. R. Noack, M. Morzyński, and L. R. Pastur, Low-order model for successive bifurcations of the fluidic pinball, *J. Fluid Mech.* **884**, A37 (2020).
- [29] F. Scarano, Iterative image deformation methods in PIV, *Meas. Sci. Technol.* **13**, R1 (2001).
- [30] A. Radford, L. Metz, and S. Chintala, Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434 (2015).
- [31] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, Amortised map inference for image super-resolution, arXiv preprint arXiv:1610.04490 (2016).
- [32] A. Lozano-Durán, O. Flores, and J. Jiménez, The three-dimensional structure of momentum transfer in turbulent channels, *J. Fluid Mech.* **694**, 100 (2012).
- [33] P. Schlatter and R. Örlü, Assessment of direct numerical simulation data of turbulent boundary layers, *J. Fluid Mech.* **659**, 116 (2010).
- [34] C. Atkinson, N. A. Buchmann, O. Amili, and J. Soria, On the appropriate filtering of PIV measurements of turbulent shear flows, *Exp. Fluids* **55**, 1 (2014).
- [35] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. C. Loy, ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, in *Proc. Eur. Conf. Comput. Vis. Workshop (2018)*, edited by L. Leal-Taixé and S. Roth (2018) pp. 63–79.
- [36] S. Gross and M. Wilber, Training and investigating residual nets, *Facebook AI Research* **6**, 3 (2016).
- [37] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network, in *Proc. IEEE CVPR* (2016) pp. 1874–1883.
- [38] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: A system for large-scale machine learning, in *12th USENIX symposium on operating systems design and implementation (OSDI 16)* (2016) pp. 265–283.
- [39] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. Eyink, A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence, *J. Turbul.* **9**, N31 (2008).