12th International Symposium on Turbulence and Shear Flow Phenomena (TSFP12)
Osaka, Japan, July 19–22, 2022

A STUDY ON SPATIAL ORGANIZATION OF COHERENT STRUCTURES IN CHANNEL FLOW USING UNSUPERVISED DEEP LEARNING

Mohammad Javad Sayyari
School of Mechanical Engineering
Pusan National University
Busan 46241, Republic of Korea
sayyari@pusan.ac.kr

Jinyul Hwang
School of Mechanical Engineering
Pusan National University
Busan 46241, Republic of Korea
jhwang@pusan.ac.kr

Kyung Chun Kim
School of Mechanical Engineering
Pusan National University
Busan 46241, Republic of Korea
kckim@pusan.ac.kr

ABSTRACT

This study aims to evaluate the capability of a GAN-based unsupervised deep learning network for generating inflow turbulent fields in a channel with the focus on the spatial organization of the coherent structures. For this purpose, a DNS is conducted at $Re_\tau \approx 178$ to collect the training dataset at the cross-stream plane of the channel at a fixed streamwise location. First, it is shown that the one-point and two-point turbulence statistics are well matched between the DNS and GAN data. Then, we extracted intense streamwise velocity clusters from the generated flow fields and compared them with those of the DNS in terms of their size distribution. Moreover, the two-point correlation coefficients computed from our GAN data are compared with those of the DNS data. Interestingly, the GAN successfully predicts the structural characteristics hidden in the training data such as the relation between the height and width of the identified clusters for negative, positive, and superposed clusters, the mean spanwise distance between the low-speed streaks, and the wall-normal variation of the conditional and unconditional length-scales in the wall-normal and spanwise directions. The successful performance of the network in this study suggests its potential to be utilized more in turbulence researches.

Introduction

As coherent structures play a critical role for understanding and controlling turbulent flows, a proposed method for predicting turbulent flow fields should be able to capture them accurately. In this regard, the direct numerical simulation (DNS) has shown promising results (Jiménez & Moin, 1991). However, the DNS is time-consuming especially at high Reynolds numbers. Alternatively, the computational costs of DNS related to a required large computational domain at high Reynolds numbers can be decreased by an accurate inflow turbulence generation method. To generate the inlet boundary condition, several methods are summarized by Wu (2017). Recently, Kim & Lee (2020) pointed out the disadvantages of these methods, suggesting machine learning (ML) techniques as an alternative. In this regard, supervised (Fukami *et al.*, 2019) and unsupervised (Kim & Lee, 2020) deep learning (DL) are applied to generate inflow for fully developed turbulent channel flows. These studies along with others have shown that the ML techniques have the potential to learn the underlying physics of turbulent flows. However, the ability of these methods for prediction of coherent structures has not been well examined yet. This study assesses the capability of an unsupervised DL technique for predicting the turbulence statistics as well as the spatial features of coherent structures in a fully developed turbulent channel flow. For this purpose, the training dataset is provided from a DNS of the channel flow at $Re_\tau \approx 178$. Then, the training dataset is fed to an unsupervised DL network based on the generative adversarial network (GAN) to generate instantaneous flow fields. Finally, the generated flow fields are compared to those of DNS in terms of the turbulence statistics and coherent structures with the focus on the two-point correlation. The results infer the potential of the unsupervised learning of turbulence to be enhanced further and become a powerful tool in the turbulence researches. As the shape and behavior of the coherent structures are mainly preserved when they grow and travel in the streamwise direction in a boundary layer flow (Robinson 1991; Adrian *et al.* 2000), a DL network can be trained at a low Reynolds number and be asked to generate instantaneous inflow boundary conditions at a high Reynolds number in a turbulent boundary layer to reduce the computational costs of DNS related to the required large computational domain at high Reynolds numbers. Recently, Kim & Lee (2020) showed successful application of a deep unsupervised network to predict turbulent flow fields in a channel flow at a higher Reynolds number ($Re_\tau = 720$) than the training dataset, but their study is limited to the one-point and two-point turbulence statistics without structural analysis.

Computational setup

DNS

To collect the training dataset, a DNS of fully-developed turbulent channel flow is performed using a fractional step method (see Kim *et al.* 2002). The periodic boundary con-

dition is applied in the streamwise (x) and spanwise (z) directions while the no-slip boundary condition is imposed on the walls. Also, the mesh is non-uniform in the wall-normal direction (y). Table 1 provides the parameters of the DNS data. Here, h denotes the channel half-height. Throughout this paper, U, V, and W indicate the streamwise, wall-normal and spanwise velolcity components. The corresponding velocity fluctuations are u, v, and w, respectively. The superscript + denotes the inner unit normalized by the friction velocity and the viscous length scales. The angled brackets $<>$ indicate the ensemble average.

Table 1: Parameters of the DNS data

Re_τ	$(L_x/h, L_z/h)$	(N_y, N_z)	Δy^+_{min}	Δy^+_{max}	Δz^+	Δt^+
178	$(4\pi, 2\pi)$	(192,256)	0.11	4.87	4.37	0.075

We used a total of 20,000 y–z planes from the DNS data at a fixed streamwise location and resized them to have a resolution of 192×192. Then, they are divided into a number of batches, having batch size of BS, and used as the training dataset, or real images.

Deep unsupervised learning network

The original GAN, first introduced by Goodfellow *et al.* (2014), consists of two main networks; namely, Generator (G) and Discriminator (D). The generator takes a random noise vector ($\vec{Z}$) in the latent space, having dimension of LS, and produces a fake image. Then, the fake and real images are fed to the discriminator which gives a number that determines the similarity between the two images (see Fig.1). The GAN network can be formulated as (Goodfellow *et al.* 2014):

$$\min_{\theta_G} \max_{\theta_D} E_X \left[\log\left(D_{\theta_D}(X)\right)\right] + E_{\vec{Z}} \left[\log\left(1 - D_{\theta_D}\left(G_{\theta_G}\left(\vec{Z}\right)\right)\right)\right] \tag{1}$$

where E denotes the expectation, and X, θ_G, and θ_D are the real image, trainable parameters of G, and trainable parameters of D, respectively. The real image is a three-channel image that includes the U^+, V^+, and W^+ information.

Gulrajani *et al.* (2017) suggested WGAN-GP which modifies the loss function of D as Eq.2:

$$L_D = E_X\left[-D(X)\right] + E_{\widetilde{X}}\left[D\left(\widetilde{X}\right)\right] + \lambda_D E_{\hat{X}}\left[\left(\left\|\nabla_{\hat{X}} D\left(\hat{X}\right)\right\|_2 - 1\right)^2\right] \tag{2}$$

where $||.||_2$ denotes the Euclidean norm, $\widetilde{X} = G(\vec{Z})$ is the generated fake image, and $\hat{X} = \sigma X + (1-\sigma)\widetilde{X}$ with σ being a random number. In this study, the WGAN-GP is implemented with the techniques proposed by Karras *et al.* (2017) such as pixelwise normalization (pn). The loss function of the generator is also constrained as Eq.3(Wu *et al.* (2020)):

$$L_{G,new} = L_G + \lambda_G ||S(X) - S(\widetilde{X})||_2^2 \tag{3}$$

where S, L_G and λ_G are the energy spectra, the loss function of the original GAN and the strength of the constraint, respectively. The statistical constraint term converges to a low value after enough number of epochs (EP). Hereafter, we call the DL network used in this study as GAN for simplicity.

As seen in Fig.1, the generator and discriminator are both networks which contain blocks of mainly convolution (conv) and fully connected (fc) layers. The network utilizes the Leaky-ReLU function with parameter a, and the ADAM optimizer with learning rate of lr. Additionally, the number of real images is increased at each iteration by mirror and spectral augmentations. The fixed parameters for the network are given in Table 2 if not mentioned. See Kim & Lee (2020) for more details about the DL network.

Table 2: Hyperparameters of the network (KS is the kernel size)

λ_D	λ_G	a	BS	lr	EP	KS	LS
100	1e-3	0.2	4	1e-3	120	3	256

Identification of the coherent structures

In this study, the coherent structures are defined as the connected points in a given instantaneous flow field that have streamwise velocities less or larger than a specific value. The connectivity rule is defined as the four orthogonal neighbors of each node in Cartesian coordinates. Thus, the positive and negative coherent structures (clusters) are defined as (Hwang & Sung 2018; Lozano-Duran & Jiménez 2014):

$$\begin{aligned} &\text{Positive clusters: } u > \alpha u_{rms} \\ &\text{Negative clusters: } u < -\alpha u_{rms} \end{aligned} \tag{4}$$

Based on percolation analysis, we found $\alpha = 1$ for both the DNS and GAN data. Although not shown here, the GAN and DNS results are well matched in the vicinity of the chosen threshold value. This is consistent with the work of Schmekel *et al.* (2022) who showed that the number and volume of clusters connected with the Reynolds stresses are preserved by a DL convolutional neural network. It should be mentioned that the clusters having area of less than $(30)^2$ in wall units are filtered out in this study.

Results and discussion
Instantaneous flow fields

A sample of the generated instantaneous velocity fields in the cross-stream plane is displayed in Fig.2. As illustrated, the spatial velocity fields are similar to the flow fields of DNS qualitatively, showing that the network is capable of emulating the training dataset. We examine these spatial fields statistically in terms of their structural characteristics. Fig.3 shows the joint probability density function (JPDF) of height and width of the identified structures in the instantaneous flow fields. Based on the physical geometry of the structures, their height and width are defined as $l_y = y_{max} - y_{min}$ and $l_z = z_{max} - z_{min}$, respectively. Here, y_{min} and y_{max} are the minimum and maximum wall-normal locations while z_{min} and z_{max} denote the minimum and maximum spanwise locations of a given structure, respectively. As seen, the identified structures are widely distributed with different sizes. As shown, the GAN lines follow the same pattern as the DNS contour levels. This agreement is also held when the structures are divided to the positive and negative clusters based on Eq.4 (see Fig.3(b) and 3(c)). This figure shows that the size and the relation between the height and width of the identified structures in the instantaneous flow fields are well learned by GAN.

Turbulence statistics

To validate our GAN results, one-point and two-point turbulence statistics obtained from the generated flow fields are compared with those of DNS. The mean statistics are ensemble-averaged over 10,000 generated fields whereas those of DNS are averaged over the whole training dataset.

Fig.4(a) compares the mean streamwise velocity, nondimensionalized by the friction velocity, for the current results

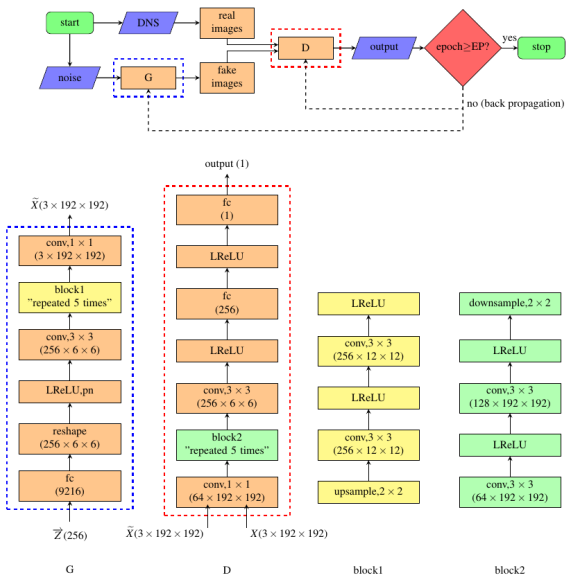

Figure 1: The GAN model used in this study with the detailed architectures of the generator (G) and discriminator (D).

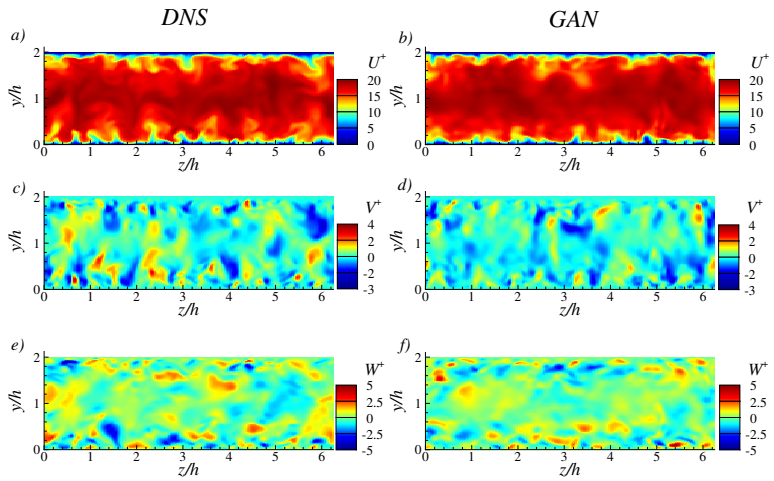

Figure 2: A sample of the generated velocity fields using the current GAN compared with DNS at $Re_\tau \approx 178$.

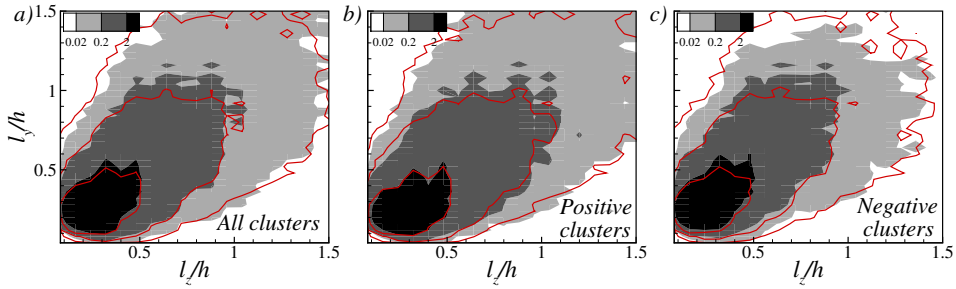

Figure 3: JPDF of height and width of *a*) all the identified clusters *b*) positive clusters, and *c*) negative clusters. The gray contour shows the DNS data while the red line indicates the GAN data. The contour levels are logarithmically distributed.

and DNS along the outer coordinates. As seen, except for the small differences mainly in the core region, where GAN predicts slightly higher values, very good agreement is obtained for the mean streamwise velocity profile. Also, as shown in Fig.4(*b*), the turbulent intensity profiles (i.e. $u^{+2}_{i,rms}(y)$, where i denotes the streamwise ($i = 1$), wall-normal ($i = 2$) and the spanwise ($i = 3$) directions) are in a good agreement when compared with those of DNS. However, GAN shows slightly lower turbulence intensities for the near-wall peaks. As seen in this figure, the GAN predicts lower values for the v and w components at all wall-normal locations which can be attributed to the problem that the mean profiles of these velocity fluctuations show higher alternations around zero compared to those of DNS (not shown here). The same issue can be seen in the results of Fukami *et al.* (2019), which is related to a lack of zero-mean constraint in the network.

Fig.4(*c*) demonstrates premultiplied one-dimensional spanwise spectra of the velocity fluctuations ($\Phi_{u_i^+ u_i^+} =$

$k_z S_{u_i^+ u_i^+}$) at two wall-normal locations; $y/\delta = 0.08$ and 0.3, with $k_z = 2\pi/\lambda_z$ being the spanwise wavenumber corresponding to the spanwise wavelength λ_z. In the near-wall region, the energy carried by large scales, which represent the superimposed large scales (or footprint), are well captured, and the premultiplied energy spectra for the near-wall and outer regions follow the same trend over a wide range of scales when compared with the DNS data. More importantly, the location of the peaks are very well captured in this figure, which shows that the GAN is capable of predicting the spanwise length-scale of the most energetic motions at a given wall-normal location. This result implies that the spatial organizations of near-wall streaks as well as outer large-scale structures are well predicted in the generated flow fields. However, the peak values predicted by GAN are lower than those of DNS, especially for the streamwise velocity.

To further explore the spatial organization of coherent structures, the two-dimensional two-point correlation coefficients of the velocity fluctuations having spanwise distance of Δz are calculated as Eq.5:

$$R_{u_i^+ u_i^+}(y,\Delta z) = \frac{< u_i^+(y,z,t) u_i^+(y_{ref}, z+\Delta z, t) >}{u_{i,rms}^+(y) u_{i,rms}^+(y_{ref})} \qquad (5)$$

where y_{ref} indicates the reference height. First, we simplify Eq.5 by focusing on the one-dimensional correlations at $y = y_{ref}$. Fig.4(d) shows the one-dimensional two-point correlation coefficients at two reference heights. As seen, the distant points are correctly uncorrelated in both DNS and GAN. As the reference height is increased from the near-wall to the logarithmic layer, the separation distance between the anti-correlated regions increases, indicating that the average spanwise length-scale of the structures are growing with height. This is in agreement with other studies including Lee & Sung (2013) and Monty *et al.* (2007). This behavior is discussed more in the next section. More importantly, the location of the minimum two-point correlation coefficient for the streamwise velocity in the logarithmic region, which gives an estimate of the mean spanwise distance between the low-speed streaks, is in good agreement with that of DNS for both wall-normal locations. As observed, the minimum location for the correlations of the spanwise and wall-normal velocity components are also in excellent agreement with DNS. These observations suggest the potential of the DL network to learn some underlying physics hidden in the turbulent flow fields such as the mean spanwise spacing between the low-speed streaks. At $y/h = 0.3$, in particular, the positive correlations for the streamwise velocity with spanwise separation of $\Delta z/h \approx 0.5$ are centered at the reference position and are flanked both sides by negative values with a distance of approximately h (see inset (b) of Fig.4(d)). This reveals that the spatial organization of the large-scale structures in the predicted flow fields are well captured at this specific height.

Wall-normal variation of the two-point correlation coefficient

To discuss the wall-normal variation of the two-point correlation coefficients, we track the integral length-scales at different reference heights. The integral length-scales are defined as Eq.6:

$$\Gamma_{i,j} = \int_a^b R_{u_i^+ u_i^+}(y,\Delta z) dr_j \qquad (6)$$

where the wall-normal and spanwise directions of integration correspond to $j = 1$ and $j = 2$, respectively, so that $r_1 = y$ and $r_2 = \Delta z$. The integration is performed on the lines passing through the reference point ($y = y_{ref}$ or $\Delta z = 0$) limited to within the zone where the correlation curve intersects with $R_{u_i^+ u_i^+} = 0.05$ (Sillero *et al.* 2014) (see inset (a) of Fig.4(d)). The first and second intersection points are shown with a and b, respectively. The integral length-scales are computed only for the lower half of the channel, due to the wall-normal symmetry of this figure with respect to the channel centerline, as the plots at the left-hand side of Fig.5(a). This figure clearly admits that the wall-normal ($j = 1$) and spanwise ($j = 2$) integral length-scales obtained from the GAN data follow the same trend as the DNS data for all the velocity perturbations. In particular, the spanwise integral length-scales for all the velocity components increase through out the boundary layer thickness, while the wall-normal integral length-scale of the streamwise velocity increases sharper near the wall and decreases near the center of the channel, which is in agreement with other studies (Monty *et al.*, 2007; Sillero *et al.*, 2014). This means that the coherent structures are persisted well even beyond the logarithmic layer in the channel flow (Lee & Sung, 2013), and the success of GAN to predict their length-scales provides evidence that this unique characteristic is reflected in the generated flow fields. However, the length-scales are underestimated by GAN compared to the DNS values. To explore the GAN capability for learning the length-scales of specific events, we further compute the length-scales based on the correlations conditioned on the negative and positive clusters.

The conditional two-point correlation coefficients are defined as Eq.7 (Sillero *et al.*, 2014):

$$R_{u_i^+ u_i^+,C}(y,\Delta z) = \frac{< u_i^+(y,z,t) u_i^+(y_{ref}, z+\Delta z, t) >|_C}{u_{i,rms}^+(y) u_{i,rms}^+(y_{ref})|_C} \qquad (7)$$

where the subscript C defines the condition being considered; that is, $C = P$ and $C = N$ mean conditioned on the positive and negative clusters formulated in Eq.4, respectively. Here, we only constrained the reference point to be within the identified clusters whereas the moving point is freely positioned on any node in the computational domain. Note that at $y = y_{ref}$ and $\Delta z = 0$, the conditional correlation gives $R_{u_i^+ u_i^+,C} = \frac{u_{i,rms}^+(y_{ref})|_C}{u_{i,rms}^+(y_{ref})}$ while the unconditional one reaches $R_{u_i^+ u_i^+} = 1$ (see Fig.5(b)). The length-scale of the conditional correlation is defined as the distance between the two intersection points of the correlation curve and the line $R_{u_i^+ u_i^+,C} = 0.3$ (Lee & Sung, 2013; Sillero *et al.*, 2014). The wall-normal and spanwise length-scales of the conditional correlation for the streamwise velocity component ($i = 1$) are shown at right-hand side of Fig.5(a) as a function of y_{ref}/h. Although not shown here, this figure also has an almost symmetric behavior with respect to the centerline of the channel, and thus is only plotted for the upper half of the channel. Surprisingly, trends of the GAN results for the positive and negative clusters are in a good agreement when compared to those of the DNS. However, the lengths obtained from the GAN data are shorter, especially for the wall-normal length-scale near the channel center. Fig.5(a) confirms that the GAN not only learns the characteristic length-scales of the superposed structures at different wall-normal locations, but also preserves the length-scales of their constitutive events such as the positive and negative clusters in a deeper level. Similar to our results, Lee & Sung (2013) reported the same conditional analysis for the streamwise two-point correlation and showed that the spanwise length-scales for the positive and negative events increases from near the wall to $y/\delta = 0.5$ and 0.6, respectively, where δ is the boundary layer thickness in their study, beyond which the curves start to decay. These locations are in excellent agreement with our GAN and DNS results.

Note that the streamwise turbulent intensities averaged separately over the positive and negative structures are in a

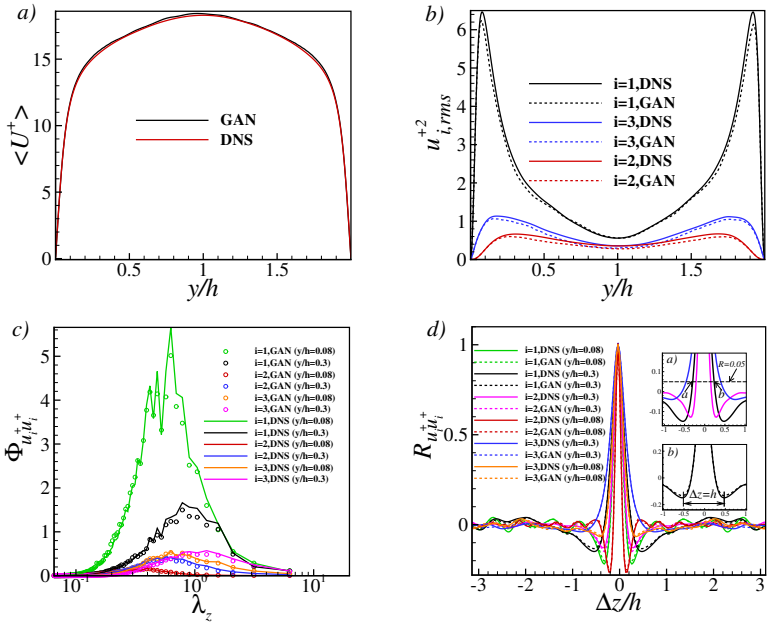

Figure 4: One-point and two-point statistics obtained from the GAN data compared with the DNS data: a) Mean streamwise velocity profile, b) Turbulence intensities, c) Pre-multiplied spanwise energy spectrum obtained at $y/h = 0.08$ and $y/h = 0.3$, and d) One-dimensional spanwise two-point correlation coefficient obtained at $y/h = 0.08$ and $y/h = 0.3$. In the inset (a) of figure (d), the dashed line shows $R_{u_i^+ u_i^+} = 0.05$ while a and b points indicate the intersections of this line with the curve of $R_{u^+u^+}$ at $y/h = 0.3$.

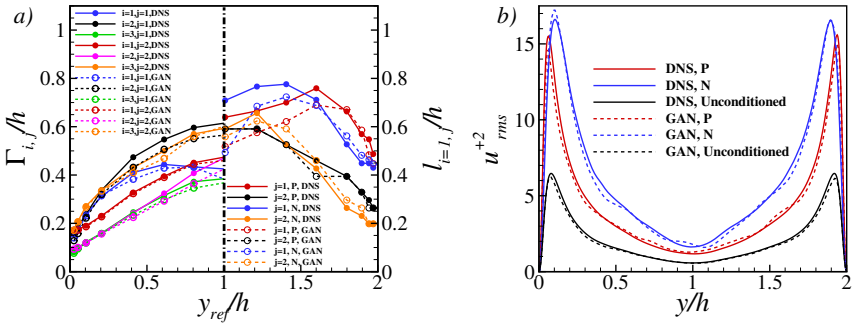

Figure 5: a) Wall-normal variation of the streamwise and wall-normal integral length-scales for all the velocity components (on the left), and length-scales for the streamwise velocity component corresponding to the line $R_{u^+u^+,C} = 0.3$ obtained based on the correlations conditioned on the positive (P) and negative clusters (N) (on the right). b) Conditional and unconditional turbulent intensities for the streamwise velocity component. The dash-dot line is the channel centerline.

very good agreement with the DNS results as given in Fig.5(b). This figure reveals that the turbulence intensities carried by the negative clusters are higher than those of the positive clusters which in turn are higher than the unconditional intensities in a large portion of the channel height. Also, the near-wall peak locations of the negative clusters, accurately predicted by GAN, occur at a larger distance from the wall compared to the other two curves.

A better understanding of the correlations dimensions can be achieved by Fig.6 which shows the 2D shape of the correlations at different wall-normal locations, covering from near the wall to the logarithmic and outer regions. The thick lines in this figure indicate the DNS data while the GAN results are shown with the thin lines. As depicted in Figs.6(a–c), the size of the objects for the three velocity fluctuations increases with the wall-normal location. In addition, the objects of u, v, and w fluctuations get closer to circular, elliptical, and squarish shape, respectively, as they move away from the wall, which is in agreement with the integral length-scales presented in Fig.5(a). Also, the distance between the anti-correlated regions (dashed lines) increases with the wall distance. These regions are located symmetrically at the left and right sides of the u and v correlations whereas the negative lobes of the w component are positioned above and below the reference point, suggesting the existence of a quasi-streamwise roll mode. In all cases, these regions disappear completely from the spanwise direction in the DNS data away from the wall. The findings discussed above are in agreement with the cross-stream analysis of Sillero *et al.* (2014) for a channel flow. Interestingly, these features are well preserved in the results obtained from the generated flow fields for a large portion of the flow thickness as shown with the thin lines in Fig.6. The GAN data shows almost identical correlation maps for the transversal velocities at all wall-normal locations as presented in Fig.6(b) and 6(c). While the correlations of the u fluctuations show similar patterns near the wall between the GAN and DNS data,

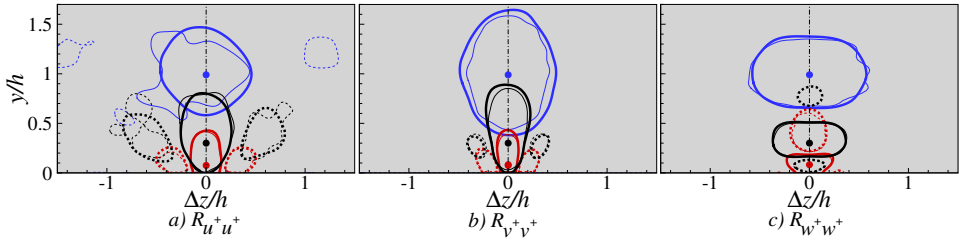

Figure 6: Two-dimensional two-point correlation coefficients of the *a*) streamwise, *b*) wall-normal, and *c*) spanwise velocity components at different reference points shown with colored circles ($y_{ref}/h = 0.08, 0.30, 1.01$). The thick lines indicate the DNS data while the thin lines are for the GAN data. At each wall-normal location, the contour level with $R = 0.1$ is shown with a solid line while the dashed line indicates $R = -0.1$. The dash-dot line indicates the spanwise symmetry.

they deviate near the edge of the boundary layer. In fact, the fluctuations surrounding the reference point in the GAN data maintain their coherence in a larger wavy pattern compared to that of the DNS with some negatively correlated regions which are not seen in the DNS correlation. As discussed by Sillero *et al.* (2014), flux of the velocity covariance in the plane normal to the velocity component is equal to zero theoretically. In our study, this integral is roughly of order 10^{-6} and 10^{-2} for the DNS and GAN data shown in Fig.6(*a*), respectively. As a result, the positive and negative correlations are strictly distributed differently in the DNS and GAN data, but as shown in this study, the important characteristics are still present in the GAN data.

Summary and conclusions

In this study, we successfully trained a GAN-based network with the cross-stream DNS data of a turbulent channel flow to generate similar flow fields that resemble those of the DNS both qualitatively and quantitatively. In particular, our results reveal that not only the flow field patterns, one-point and two-point statistics are quite matched between the DNS and GAN data, but also the intense velocity coherent structures obtained from the generated flow fields follow the same behavior as the DNS results in terms of the relation between their height and width. Comparing the two-point correlation coefficients for the DNS and GAN data in the cross-stream plane, it turned out that the wall-normal variation of the length-scales obtained from the conditional and unconditional correlations are well reflected in the GAN data. Overall, a satisfactory agreement is obtained between the DNS and GAN flow fields in this study, with a general underprediction in the GAN results. The evidences gathered in this study clearly show the potential of the unsupervised deep learning network to learn structural features hidden in the training dataset.

REFERENCES

Adrian, R. J., Meinhart, C. D. & Tomkins, C. D. 2000 Vortex organization in the outer region of the turbulent boundary layer. *Journal of Fluid Mechanics* **422**, 1–54.

Fukami, K., Nabae, Y., Kawai, K. & Fukagata, K. 2019 Synthetic turbulent inflow generator using machine learning. *Phys. Rev. Fluids* **4**, 064603.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. 2014 Generative adversarial networks.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. & Courville, A. C. 2017 Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, , vol. 30, pp. 5767–5777.

Hwang, J., Lee, J., Sung, H. J. & Zaki, T. A. 2016 Inner–outer interactions of large-scale structures in turbulent channel flow. *Journal of Fluid Mechanics* **790**, 128–157.

Hwang, J. & Sung, H. J. 2018 Wall-attached structures of velocity fluctuations in a turbulent boundary layer. *Journal of Fluid Mechanics* **856**, 958–983.

Jiménez, J. & Moin, P. 1991 The minimal flow unit in near-wall turbulence. *Journal of Fluid Mechanics* **225**, 213–240.

Karras, T., Aila, T., Laine, S. & Lehtinen, J. 2017 Progressive growing of gans for improved quality, stability, and variation.

Kim, J. & Lee, C. 2020 Deep unsupervised learning of turbulence for inflow generation at various reynolds numbers. *Journal of Computational Physics* **406**, 109216.

Kim, K., Baek, S.-J. & Sung, H. J. 2002 An implicit velocity decoupling procedure for the incompressible navier-stokes equations. *International Journal for Numerical Methods in Fluids* **38**, 125–138.

Lee, J. H. & Sung, H. J. 2013 Comparison of very-large-scale motions of turbulent pipe and boundary layer simulations. *Physics of Fluids* **25** (4), 045103.

Lozano-Duran, A. & Jiménez, J. 2014 Time-resolved evolution of coherent structures in turbulent channels: characterization of eddies and cascades. *Journal of Fluid Mechanics* **759**, 432–471.

Monty, J. P., Stewart, J. A., Williams, R. C. & Chong, M. S. 2007 Large-scale features in turbulent pipe and channel flows. *Journal of Fluid Mechanics* **589**, 147–156.

Robinson, S. K. 1991 Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics* **23** (1), 601–639.

Schmekel, D., Alcántara-Ávila, F., Hoyas, S. & Vinuesa, R. 2022 Predicting coherent turbulent structures via deep learning. *Frontiers in Physics* **10**.

Sillero, J. A., Jiménez, J. & Moser, R. D. 2014 Two-point statistics for turbulent boundary layers and channels at Reynolds numbers up to $\delta^+ \approx 2000$. *Physics of Fluids* **26** (10), 105109.

Wu, J.-L., Kashinath, K., Albert, A., Chirila Prabhat, D. & Xiao, H. 2020 Enforcing statistical constraints in generative adversarial networks for modeling chaotic dynamical systems. *Journal of Computational Physics* **406**, 109209.

Wu, X. 2017 Inflow turbulence generation methods. *Annual Review of Fluid Mechanics* **49**, 23–49.