



Application of Video-to-Video Translation Networks to Computational Fluid Dynamics

Hikomitsu Kigure*

Independent Researcher, Kanagawa, Japan

In recent years, the evolution of artificial intelligence, especially deep learning, has been remarkable, and its application to various fields has been growing rapidly. In this paper, I report the results of the application of generative adversarial networks (GANs), specifically video-to-video translation networks, to computational fluid dynamics (CFD) simulations. The purpose of this research is to reduce the computational cost of CFD simulations with GANs. The architecture of GANs in this research is a combination of the image-to-image translation networks (the so-called “pix2pix”) and Long Short-Term Memory (LSTM). It is shown that the results of high-cost and high-accuracy simulations (with high-resolution computational grids) can be estimated from those of low-cost and low-accuracy simulations (with low-resolution grids). In particular, the time evolution of density distributions in the cases of a high-resolution grid is reproduced from that in the cases of a low-resolution grid through GANs, and the density inhomogeneity estimated from the image generated by GANs recovers the ground truth with good accuracy. Qualitative and quantitative comparisons of the results of the proposed method with those of several super-resolution algorithms are also presented.

Keywords: deep learning, generative adversarial networks (GANs), image-to-image translation networks (pix2pix), long short-term memory (LSTM), computational fluid dynamics (CFD)

OPEN ACCESS

Edited by:

Alex Hansen,
Norwegian University of Science and
Technology, Norway

Reviewed by:

Bernhard C. Geiger,
Know Center, Austria
Ajey Kumar,
Symbiosis International (Deemed
University), India

*Correspondence:

Hikomitsu Kigure
hiro.kig@gmail.com

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

Received: 20 February 2021

Accepted: 17 August 2021

Published: 10 September 2021

Citation:

Kigure H (2021) Application of Video-
to-Video Translation Networks to
Computational Fluid Dynamics.
Front. Artif. Intell. 4:670208.
doi: 10.3389/frai.2021.670208

1 INTRODUCTION

Artificial intelligence is advancing rapidly and has come to be comparable to or outperform humans in several tasks. In generic object recognition, deep convolutional neural networks have surpassed human-level performance (e.g. He et al., 2015; He et al., 2016; Ioffe and Szegedy, 2015). The agent trained by reinforcement learning is capable of reaching a level comparable to professional human game testers (Mnih et al., 2015). In the case of machine translation, Google’s neural machine translation system, using Long Short-Term Memory (LSTM) recurrent neural networks [Hochreiter and Schmidhuber (1997), Gers et al. (2000)], is a typical and famous example and its translation quality is becoming comparable to that of humans (Wu et al., 2016).

One of the hottest research topics in artificial intelligence is generative models and one approach to implementing a generative model is generative adversarial networks (GANs) proposed by Goodfellow et al. (2014). GANs consist of two models trained with conflicting objectives. Radford et al. (2016) applied deep convolutional neural networks to those two models, whose architecture is called deep convolutional GANs (DCGAN). DCGAN can generate realistic synthesis images from vectors in the latent space. Isola et al. (2017) proposed the network learning the mapping from an input image to an output image to enable the translation between two images. This network, the so-called pix2pix, can convert black-and-white images into color images, line drawings into photo-realistic images, and so on.

The combination of deep learning and simulation has been recently researched. One of such applications is to use simulation results for improving the prediction performance of deep learning. Since deep learning requires a lot of data for training, numerical simulations that can generate various data by changing physical parameters could help compensate for the lack of training data. Another application is to speed up the solver of computational fluid dynamics (CFD). Guo et al. (2016) used a convolutional neural network (CNN) to predict velocity fields approximately but fast from the geometric representation of the object. Another example is that velocity fields are predicted from parameters such as source position, inflow speed, and time by CNN (Kim et al. (2019)). Their method is feasible to generate velocity fields up to 700 times faster than simulations. As a more general method, not limited to CFD problems, Raissi et al. (2019) proposed the physics-informed neural network (PINN), which utilizes a relatively simple deep neural network to find solutions to various types of nonlinear partial differential equations.

GANs also have been combined with numerical simulations to enable a new type of solution method. Farimani et al. (2017) used the conditional GAN (cGAN) to generate the solution of steady-state heat conduction and incompressible flow from boundary conditions and calculation domain shape/size. Xie et al. (2018) proposed a method for super-resolution fluid flow by a temporally coherent generative model (tempoGAN). They showed that tempoGAN can infer high-resolution, temporal, and volumetric physical quantities from those of low-resolution data.

The above-mentioned studies about the combination of GANs and simulations show that GANs can generate the three-dimensional data of the solution of physical equations. The main topic in this research is the translation of images (distributions of the physical quantity) by GANs. In the case that the accuracy of the simulation is particularly important, a large number of computational grids are needed. Additionally, the number of simulation cases for design optimization is typically numerous. It means that the computational cost (machine power and time) becomes large. In such a case, it is important to reduce the computational cost, and one way to do so is to make effective use of low-cost simulations. Based on such an idea, I investigated the feasibility of time-series image-to-image translation: translation from time-series distribution plots in the case of low-resolution computational grids to those in the case of high-resolution grids. A quantitative evaluation of the quality of generated images was also performed.

The method proposed in this paper is the video (sequential images)-to-video translation in which the difference of solutions between the high- and low-resolution grid simulations is learned. Meanwhile, the PINN constructs universal function approximators of physical laws by minimizing the loss function composed of a mismatch of state variables including the initial and boundary conditions and the residual for the partial differential equations (Meng et al., 2020). In other words, the PINN is an alternative to CFD, while the proposed method is a complement to CFD.

The paper is organized as follows. In **section 2**, I describe the outline of the simulations whose results are input to GANs and

TABLE 1 | The initial conditions of simulations.

Physical quantity	Description	Value
ρ	Density	$25\pi/36$
v_x	x-component of velocity	$-\sin(2\pi y)$
v_y	y-component of velocity	$\sin(2\pi x)$
v_z	z-component of velocity	0
B_x	x-component of magnetic field	$-B_0 \sin(2\pi y)$
B_y	y-component of magnetic field	$B_0 \sin(2\pi x)$
B_z	z-component of magnetic field	0
p	Pressure	$5\pi/12$

the details of the network architecture. In **section 3**, I give the results of time-series image-to-image translation (in other words, video-to-video translation) of physical quantity distribution and a discussion mainly about the quality of generated images. Conclusions are presented in **section 4**.

2 METHODS

2.1 Numerical Simulations

I solved the following ideal magnetohydrodynamic (MHD) equations numerically in two dimensions to prepare input images to GANs:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \quad (1)$$

$$\frac{\partial}{\partial t} (\rho \mathbf{v}) + \nabla \cdot (\rho \mathbf{v} \mathbf{v} + p_T \mathbf{I} - \mathbf{B} \mathbf{B}) = 0 \quad (2)$$

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla \cdot (\mathbf{v} \mathbf{B} - \mathbf{B} \mathbf{v}) = 0 \quad (3)$$

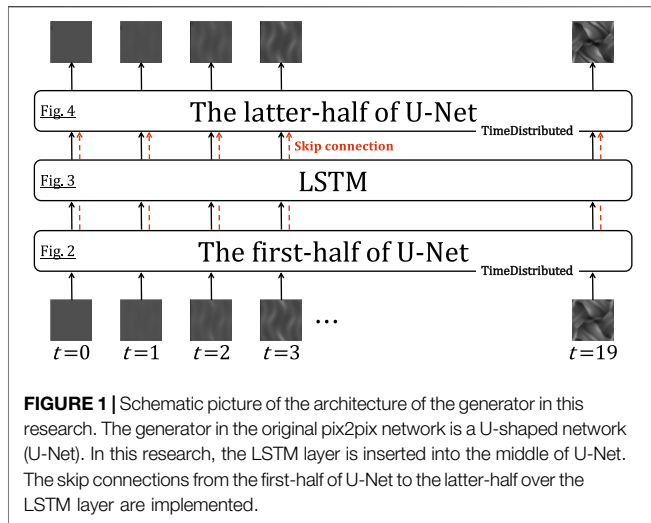
$$\frac{\partial e}{\partial t} + \nabla \cdot ((e + p_T) \mathbf{v} - \mathbf{B} (\mathbf{v} \cdot \mathbf{B})) = 0 \quad (4)$$

$$p_T = p + \frac{|\mathbf{B}|^2}{2} \quad (5)$$

$$e = \frac{p}{\gamma - 1} + \frac{\rho |\mathbf{v}|^2}{2} + \frac{|\mathbf{B}|^2}{2} \quad (6)$$

where ρ , p , and \mathbf{v} are the density, pressure, and velocity of the gas; \mathbf{B} is the magnetic field; γ represents the heat capacity ratio and is equal to 5/3 in this paper; p_T and e represent the total pressure and the internal energy density; \mathbf{I} is the unit matrix.

One of typical test problems for MHD, the so-called Orszag-Tang vortex problem (Orszag and Tang, 1979), was solved by the Roe scheme (Roe 1981) with MUSCL [monotonic upstream-centered scheme for conservation laws; (van Leer 1979)]. The initial conditions are summarized in **Table 1**. B_0 is a parameter for controlling the magnetic field strength. The computational domain is $0 \leq x \leq 1$ and $0 \leq y \leq 1$. The periodic boundary condition is applied in both x - and y -directions. Simulations for each condition were performed twice on computational grids with different resolutions. The number of grid points is $(N_x \times N_y) = (51 \times 51)$ or (251×251) . In the case of $(N_x \times N_y) = (251 \times 251)$, the calculation time is more than 70 times longer than in the other case though the obtained solution is expected to be close to the true solution.



2.2 Generative Adversarial Network Architecture

After the original concept of GANs was proposed by Goodfellow et al. (2014), various GANs have been researched. Among such networks, I focused on pix2pix, which is a type of conditional GAN and a network for learning the relationship between the input and output images. The feasibility of translating from the results of low-resolution grid simulations to those of high-resolution grid simulations has been investigated in this research. Furthermore, in order to enable the translation across two time-series, the architecture combined pix2pix and LSTM has been constructed.

Figure 1 shows the schematic picture of the architecture of the generator in this research. The role of the LSTM layer is to adjust the image translation dependent on the physical time of the simulation; for the initial state of the simulation ($T = 0$), no translation is needed at all, but as physical time passes, progressively larger translations are needed. Note that the weights of the encoder (decoder) before (after) the LSTM layer are the same in the time direction.

Plots of the time evolution of the density in the low-resolution simulations are input to the generator (plots are read as single-channel images). The input images are converted to vectors by the first-half of a U-shaped network (U-Net). In **Figure 2**, I denote the architecture of the first-half of U-Net in detail. It consists of eight convolutional blocks with a kernel size of (4×4) or (2×2) . The instance normalization (Ulyanov et al. (2017)) is applied except for the first and last blocks. The activation function is a leaky rectified linear unit [leaky ReLU; Maas et al. (2013)] with a slope of 0.2 for all blocks. A 512-dimensional vector is generated at the end of this architecture.

A series of 512-dimensional vectors converted from the time-series plots is input to the LSTM layer. An input vector \mathbf{x}_t originated from the plot at time t is calculated with the hidden state \mathbf{h}_{t-1} and memory cell \mathbf{c}_{t-1} . A forget gate (f), an input gate (i), an output gate (o), and part of the term to be added to the memory cell (z) in **Figure 3** are calculated as follows:

$$f = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{R}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (7)$$

$$i = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{R}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (8)$$

$$o = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{R}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (9)$$

$$z = \tanh(\mathbf{W}_z \mathbf{x}_t + \mathbf{R}_z \mathbf{h}_{t-1} + \mathbf{b}_z) \quad (10)$$

where σ is the sigmoid function and \tanh is the hyperbolic tangent function; \mathbf{W} and \mathbf{R} are the input-to-hidden weight matrices and the recurrent weight matrices; \mathbf{b} are bias vectors. The hidden state and memory cell are updated by:

$$\mathbf{c}_t = f \odot \mathbf{c}_{t-1} + i \odot \mathbf{z} \quad (11)$$

$$\mathbf{h}_t = o \odot \tanh(\mathbf{c}_t) \quad (12)$$

The hidden state \mathbf{h}_t is reshaped as $(1, 1, 512)$. The reshaped hidden state \mathbf{h}_t' is passed to the latter-half of U-Net and is decoded to the image data (**Figure 4**). This part consists of eight deconvolutional blocks with an upsampling of the feature map, convolution with a kernel size of (2×2) or (4×4) (the size of the feature map does not change because the stride

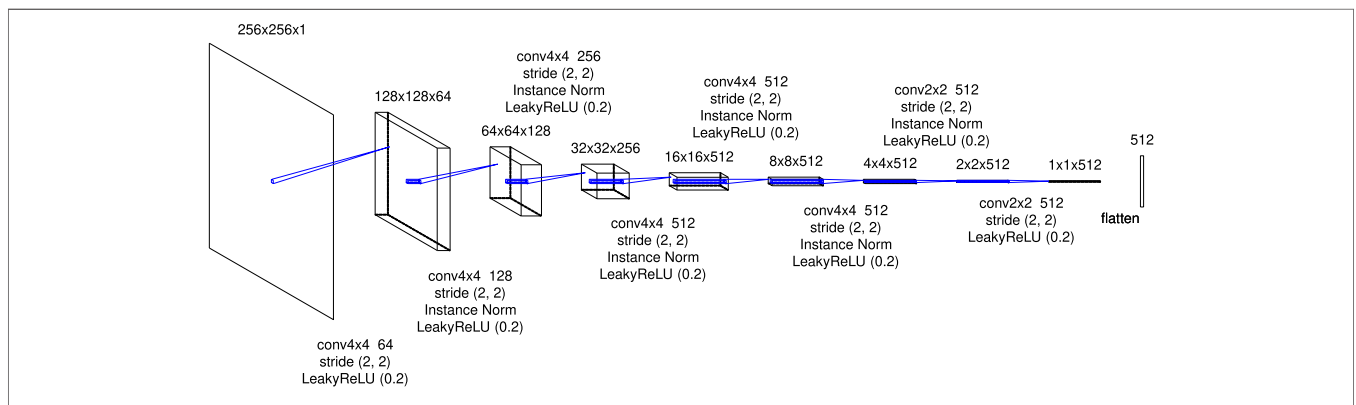


FIGURE 2 | The details of the first-half of U-Net. The expression “conv4x4 64” refers to a convolutional layer with a kernel size of (4×4) and 64 channels. Each feature map is copied and is concatenated to the feature map of the corresponding block in the latter-half of U-Net denoted in **Figure 4**.

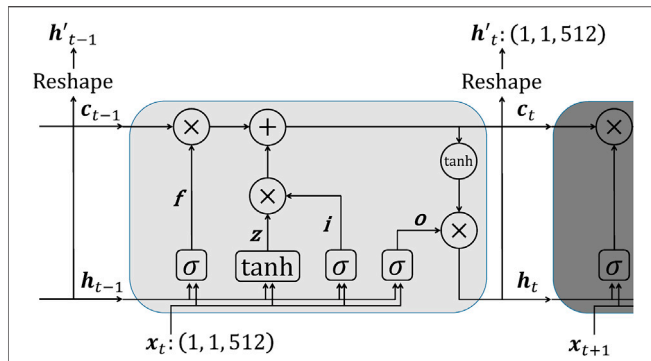


FIGURE 3 | The architecture of LSTM. The input to LSTM (x_t) is the vector transformed from an image of density distribution, and the output is the reshaped hidden state vector (h'_t) resulting from several operations. The vector c is the memory cell, and f , i , o , and z are a forget gate, an input gate, an output gate, and part of the term to be added to the memory cell (see equations (7) to (10) for details).

of convolution is 1), the instance normalization and activation by ReLU function except the last block. As seen in **Figure 1**, the generator outputs synthetic time-series plots of density distribution.

The authenticity of the images is judged by the discriminator. **Figure 5** shows the details of the architecture of the discriminator in this research. A real image (plot of the density distribution in a high-resolution simulation) or a synthesis image is input to the discriminator. It consists of five convolutional blocks with a kernel size of (4×4) . The instance normalization is applied except for the first and last blocks. Except for the last block, the leaky ReLU function with a slope of 0.2 is applied as the activation function. The 16×16 patch is eventually output. The discriminator classifies each patch into real or synthetic. We call its architecture the patchGAN (Isola et al., 2017).

The objective of the network is the same as the regular pix2pix as follows:

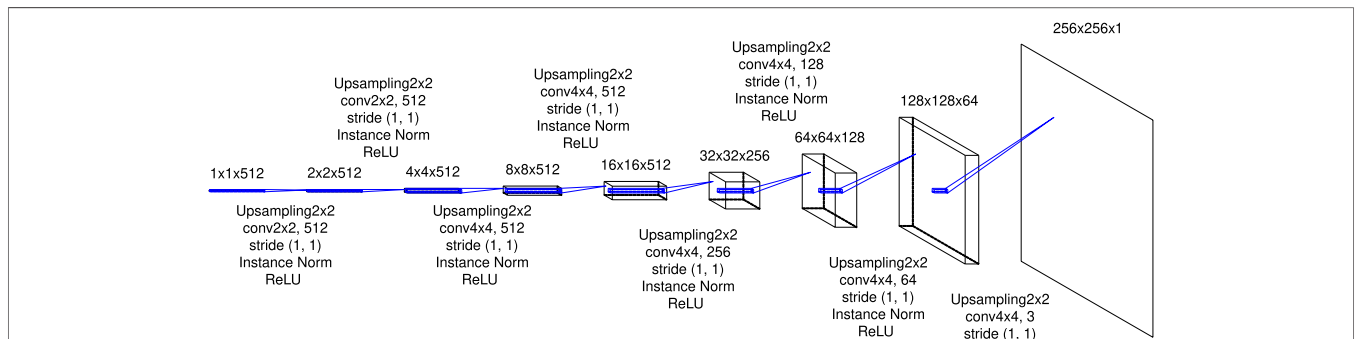


FIGURE 4 | The details of the latter-half of U-Net. The expression “Upsampling2x2” refers to an upsampling layer that doubles the size of input by copying one value twice horizontally and vertically, respectively. From the first-half of U-Net displayed in **Figure 2**, feature maps are passed to corresponding blocks and are concatenated to the feature maps output from the previous blocks.

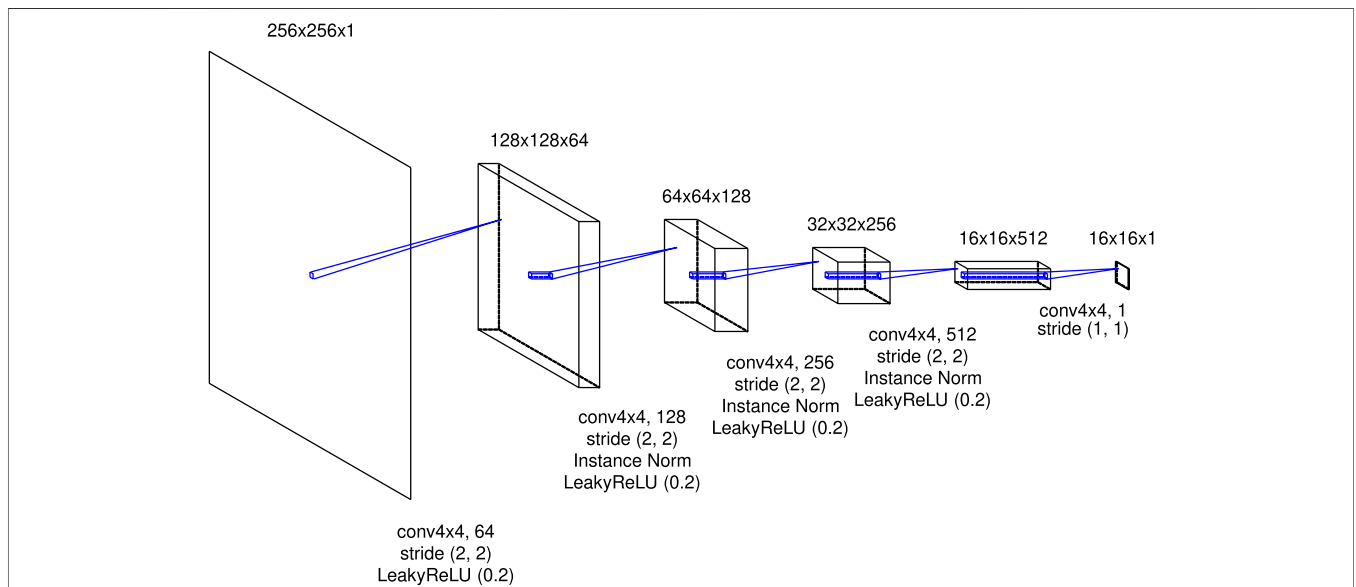
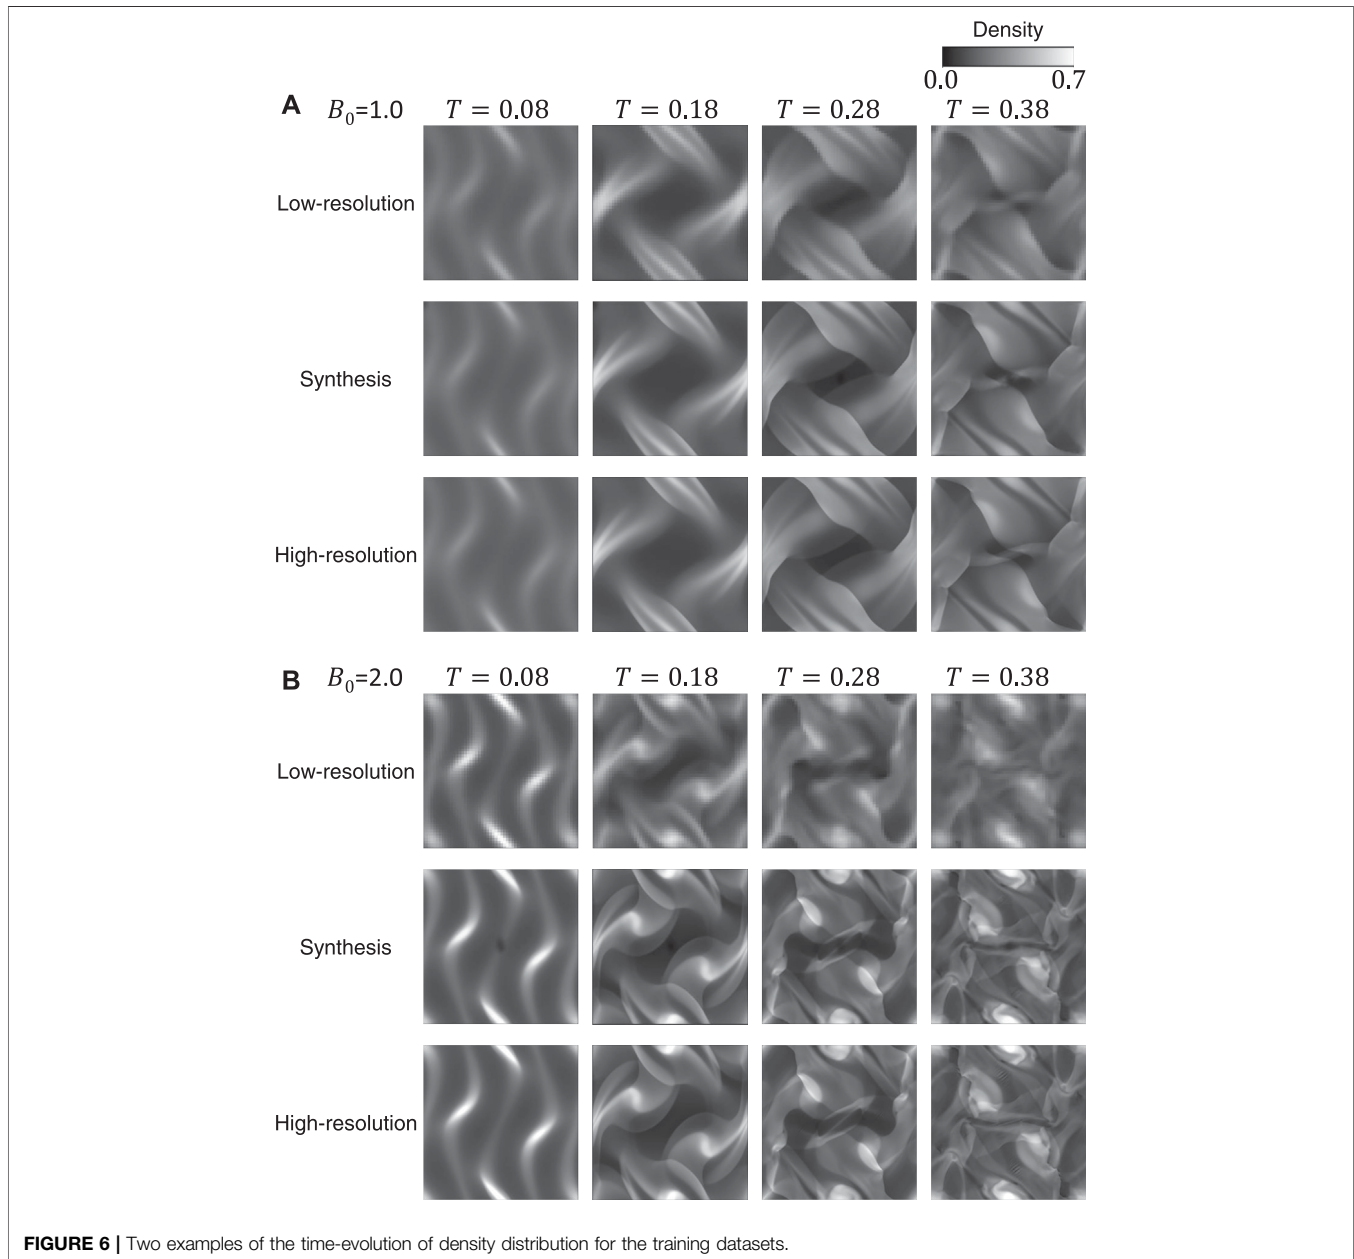


FIGURE 5 | The details of the architecture of the discriminator in this research.

TABLE 2 | The details of the training and testing datasets.

Training/Testing	The number of cases	The total number of images	The pixel size of images	The value of B_0
Training	16	320	256 × 256	0.1–1.5 with the interval of 0.1, and 2.0
Testing	19	380	256 × 256	0.15–1.55 and 1.6–1.9 with the interval of 0.1

**FIGURE 6** | Two examples of the time-evolution of density distribution for the training datasets.

$$G^* = \arg \min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G) \quad (13)$$

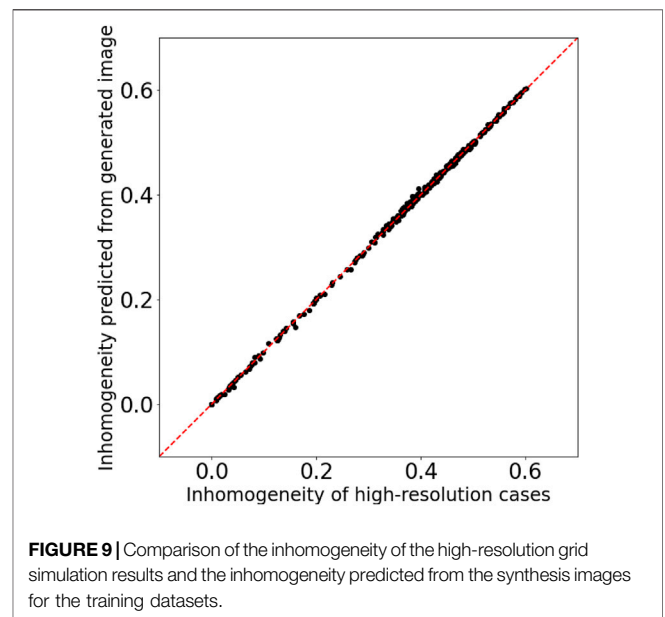
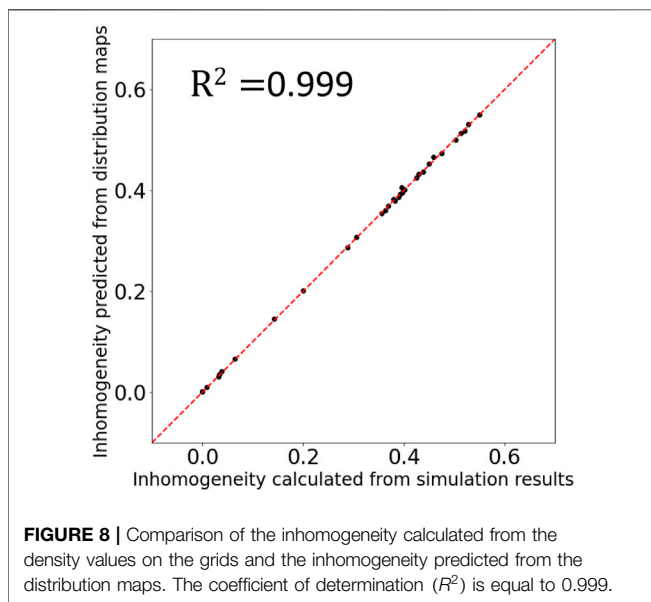
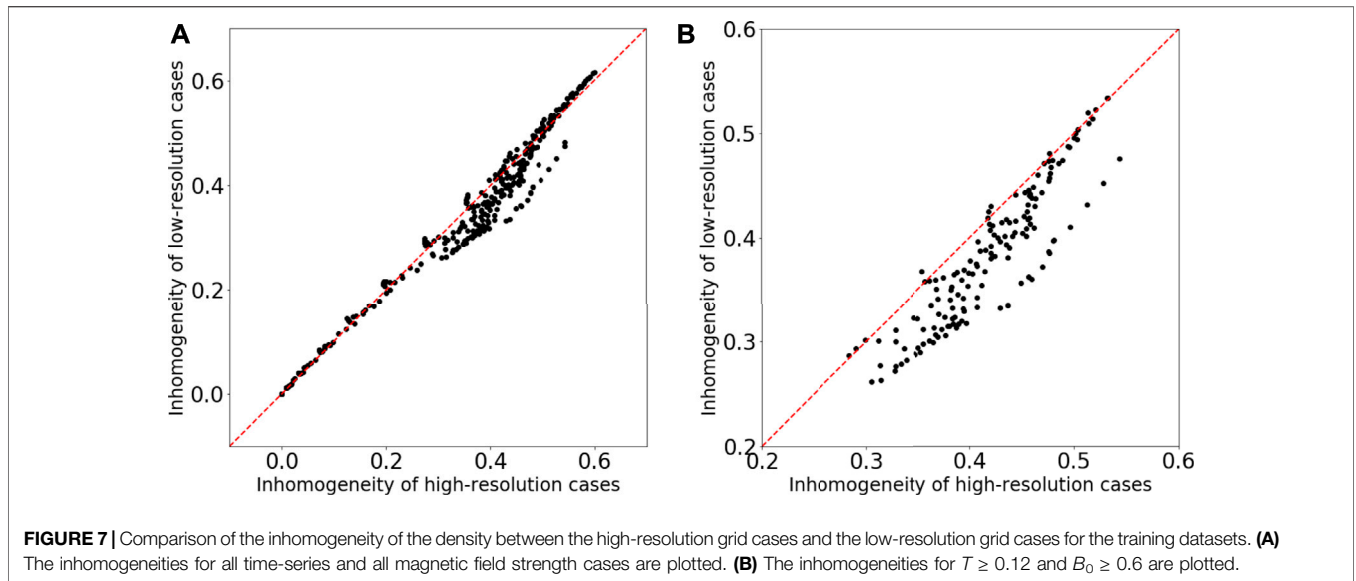
$$\mathcal{L}_{cGAN}(G, D) = \mathbb{E}[\log D(x, y)] + \mathbb{E}[\log(1 - D(x, G(x)))] \quad (14)$$

$$\mathcal{L}_{L1}(G) = \mathbb{E}[\|y - G(x)\|] \quad (15)$$

where G and D denote the generator and discriminator, λ is the weighted sum parameter and equal to 100 in this research, and x and y

mean the source and target images. $G(x)$ returns a synthesis image and $D(x, y)$ or $D(x, G(x))$ returns the probability that y or $G(x)$ is a real target image. $\mathcal{L}_{L1}(G)$ is the mean absolute error (L1 loss) calculated from the pixel-wise comparison between the real image and the synthetic image. The optimizer is Adam with a learning rate of 0.0002.

The architecture is implemented using Keras 2.5.0 and TensorFlow as a backend. The model was trained on Google



Colaboratory with Tesla P100-PCIe GPU. For applying the convolution and deconvolution to the sequential data, sets of operations as shown in **Figures 2,4,5** are passed to the TimeDistributed layer. The skip connections are implemented by feeding the outputs of the previous upsampling block in the latter-half of U-Net and the same-level (it means that the size of the feature map is the same) convolutional block in the first-half of U-Net to the Concatenate layer. The “return_sequences” and “stateful” parameters in the LSTM layer are set to True and False, respectively.

3 RESULTS AND DISCUSSION

In this chapter, I show the results of time-series image-to-image translation for the training datasets first and then explain the way

to evaluate the quality of the synthesis images quantitatively. The evaluation result of the synthesis images for the training datasets is presented next. Then, I show the results for the testing datasets. Finally, the quality of the synthesis images is compared with those of images upsampled by conventional super-resolution algorithms. The conditions (the magnetic field strength) of the simulations are shown in **Table 2** that summarizes the details of the training and testing datasets. The sixteen cases were performed to prepare the training datasets, and the nineteen cases were performed to prepare the testing datasets. For each case, two simulations were run with the high-resolution and the low-resolution grids.

3.1 Results for the Training Datasets

Figure 6 shows two examples of the time-evolution of density distribution for the training datasets. The top and bottom images

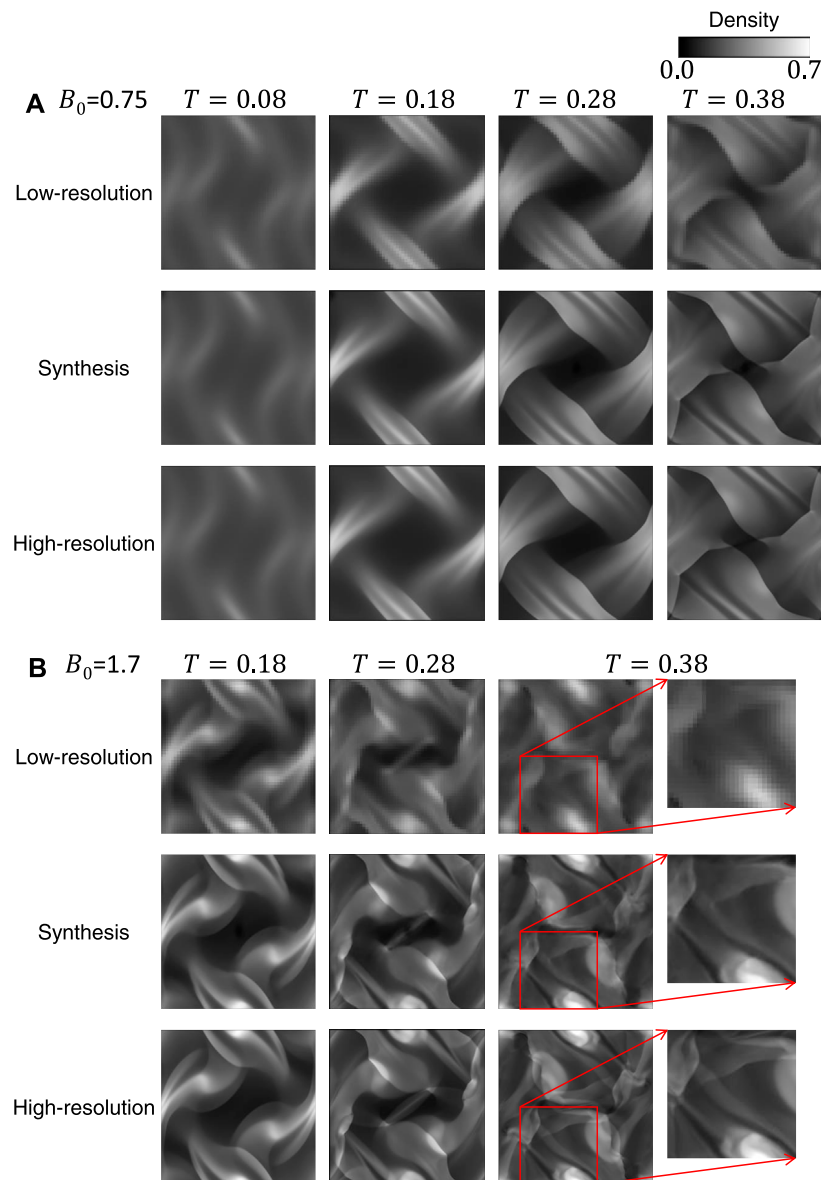
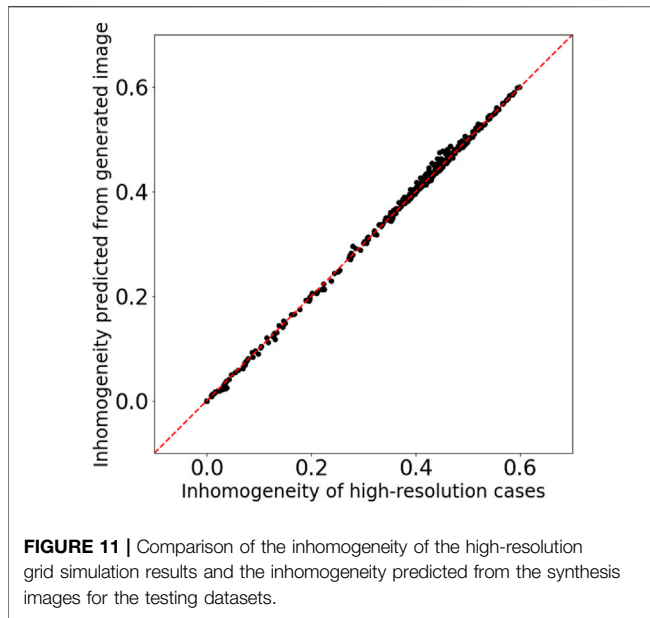


FIGURE 10 | Examples of the time-evolution of density distribution for the testing datasets.

of **Figures 6A,B** show the simulation results, and the middle images are synthesis ones generated from the top ones (the results of low-resolution grid simulations) through the generator. Compared to the high-resolution grid cases, the density distributions in the low-resolution grid cases show less fine structure and become closer to the uniform. **Figure 7** displays the comparison of the inhomogeneity of the density between the high-resolution grid cases and the low-resolution grid cases. The inhomogeneity is defined by $\alpha = \sigma_\rho / \bar{\rho}$, where σ_ρ and $\bar{\rho}$ are the standard deviation and the average of the density. In the low-resolution grid, the numerical diffusion is larger than in the high-resolution grid, and therefore the inhomogeneity of the density tends to be smaller especially from the middle stage of the vortex development and in the relatively strong magnetic field (**Figure**

7B). The synthesis images reproduce the fine structures of the density distributions and appear to be well consistent with the high-resolution grid results.

To quantitatively evaluate the quality of the synthesis images, I estimated the density inhomogeneity from the distribution map. When calculating the density inhomogeneity from the simulation result, we can use the value of the density on each grid; however, the density distribution maps (including synthesis images in this research) have only the information of the RGB values. Therefore, to estimate the density inhomogeneity from the distribution map, I trained a three-layer fully connected neural network with 196,608 ($256\text{pixel} \times 256\text{pixel} \times 3$) inputs, two hidden layers of 1,024 and 128 neurons and one output layer. **Figure 8** shows the result of the inhomogeneity prediction from the density



distribution maps. The horizontal axis is the inhomogeneity calculated from the density values on the grids, and the vertical axis is the inhomogeneity predicted from the distribution maps by the trained neural network. The coefficient of determination (R^2) is equal to 0.999. Thus we conclude that the trained neural network provides an accurate estimation of the density inhomogeneity from the distribution maps and the synthesis images.

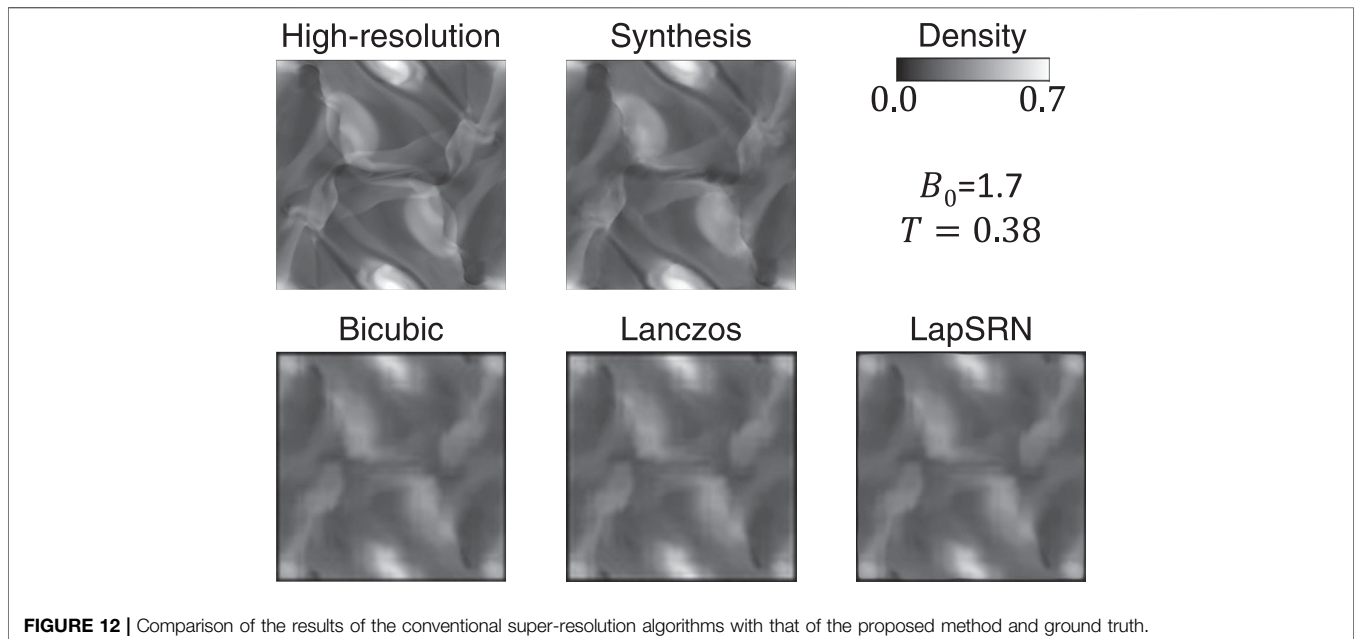
We can quantitatively evaluate the quality of the synthesis images by inputting those into the neural network and comparing the outputted inhomogeneity with the inhomogeneity calculated from the high-resolution grid simulation results. **Figure 9** shows

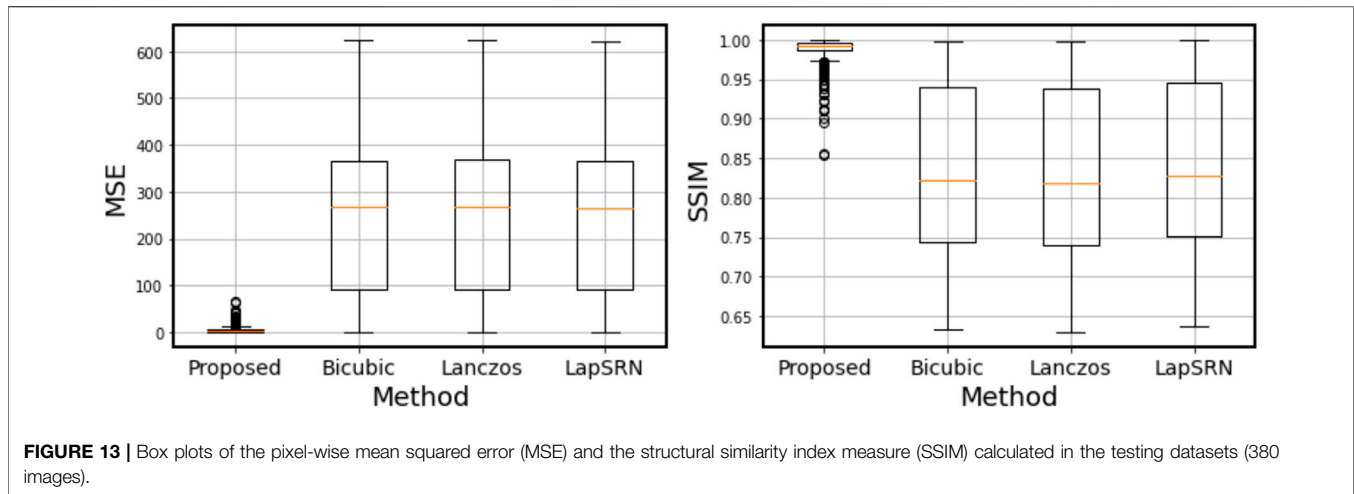
that the inhomogeneity predicted from the synthesis images matches that calculated from the high-resolution grid simulation results with good accuracy; therefore the quality of the synthesis images is definitely good for the training datasets.

3.2 Results for the Testing Datasets

In the previous subsection, I have shown that the results for the training datasets are pretty good. However, the generalization ability needs to be investigated for practical use. The testing datasets (the magnetic field strength is different from the training datasets as shown in **Table 2**) that were not used for training are input to the trained model, and the synthesis images are output from the generator. **Figures 10A,B** show the comparison of the simulation results and the synthesis images for two example cases. From the 19 cases in the testing datasets, the results for the cases with $B_0 = 0.75$ and 1.7 were selected for presentation. The $B_0 = 1.7$ case is especially suitable for verifying the generalization ability because there is no training data between $B_0 = 1.5$ and 2.0 . The top images show the time evolution of the density distribution of low-resolution grid simulations, which are input for the generator; the bottom images show that of high-resolution grid simulations, which are compared with the synthesis images; the middle images are synthesis ones generated through the generator. As with the cases for the training datasets, the synthesis images qualitatively reproduce the density distributions of the high-resolution grid simulations. Even in the $B_0 = 1.7$ case, the synthesis images show the fine structure of the density distribution very similar to that in the ground truth images, as shown in the zoomed-in image in **Figure 10B**.

Figure 11 is almost the same as **Figure 9** but for the testing datasets. The density inhomogeneity predicted from the synthesis images through the fully connected neural network (explained in the previous subsection) is in good agreement with the inhomogeneity calculated from the results





of high-resolution grid simulations. This result indicates that the method in this research is capable of obtaining high generalization ability.

3.3 Comparison With Conventional Super-resolution Algorithms

To demonstrate the effectiveness of the proposed method and the quality of the generated images, I compare the results with those obtained by conventional super-resolution algorithms. The algorithms investigated here are a bicubic interpolation, a Lanczos interpolation, and Laplacian Pyramid Super-Resolution Network [LapSRN; Lai et al. (2017)]. The pixel size of the image to be used as the basis of the super-resolution is 64×64 , and each algorithm quadruples the pixel size. These results were compared qualitatively and quantitatively with the result of high-resolution grid simulation and the image generated by the proposed method. Plots of the density distribution in high-resolution simulations in the training datasets were used to train LapSRN.

I performed the super-resolution algorithms to the testing datasets (380 images). As an example, the results for the $B_0 = 1.7$ and $T = 0.38$ case are compared in **Figure 12**. In this case, none of the three conventional super-resolution algorithms can work with a quality comparable to the method proposed in this research. To compare the proposed method with the others quantitatively, the pixel-wise mean squared error (MSE) and the structural similarity index measure [SSIM; Wang et al. (2004)] are calculated between the ground truth image and the synthesis image or the result of super-resolution. **Figure 13** shows that the quality of the synthesis images by the proposed method is significantly high compared to that of the results by the conventional super-resolution algorithms.

3.4 Application of This Research

In this subsection, I discuss an application of this research. As mentioned above, results of high computational cost simulations can be estimated from those of low-cost simulations by the method in this paper. However, it is important to note that simulation

results of quite a few cases are needed to train the network¹. Therefore, it is not beneficial for a small number of simulations. The more simulations are required, the greater the benefits arise. One such case is optimization based on CFD simulations. As the number of objective variables to be optimized increases, the number of calculations required to obtain the desired performance is expected to increase; in some cases, it takes several thousand cases to evaluate. In such multi-objective optimization simulations, for example, the first dozens to several hundred cases are simulated on both high- and low-resolution grids, and the results are used to train the GANs. After the GANs are trained, low-resolution grid simulations are run, the results are input to the GANs to reproduce the results of high-resolution grid simulations, and objective variables are estimated from synthesis images by, for example, a neural network.

I demonstrate the estimation of computational cost reduction. If the number of simulations required originally and that to train the GANs are N (several thousands in some cases) and N_t ($N > N_t$), the calculation times of the high- and low-resolution grid simulations are T_h and T_l ($T_h > T_l$), and the computational cost to train the GANs is T_b , the computational cost reduction is roughly equal to

$$N \times T_h - (N_t \times T_h + T_l + N \times T_l) \quad (16)$$

where the first term corresponds to the computational cost in the case that all simulations are run on the high-resolution grid, and the second term corresponds to that in the case that the method in this research is applied (the cost to reproduce the results of high-resolution grid simulations by the GANs is negligible compare to performing the simulations). In this way, by substituting low-resolution grid simulations and the result conversion by the GANs for quite a part of high-resolution grid simulations, a great reduction of the computational cost should be achieved.

¹In this research, simulation results of 16 cases were used as the training datasets; the training was successful with a relatively small number of data, probably due to the simple situation. If the target is a simulation of a realistic engineering situation, it is expected that much more data will be needed for the training.

4 CONCLUSION

In this paper, I validated an idea to use GANs for reducing the computational cost of CFD simulations. I studied the idea of reproducing the results of high-resolution grid simulations with a high computational cost from those of low-resolution grid simulations with a low computational cost. More specifically speaking, distribution maps of a physical quantity in time series were reproduced using pix2pix and LSTM. The quality of the reproduced synthesis images was good for both the training and testing datasets. The conditions treated in this paper are simple; the computational region is a square with a constant grid interval, the boundary conditions are cyclic, and the governing equations are the ideal MHD

equations. In the next step, I need to examine the idea in more realistic conditions.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

AUTHOR CONTRIBUTIONS

HK: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Visualization, Writing.

REFERENCES

- Farimani, A. B., Gomes, J., and Pande, V. S. (2017). *Deep Learning the Physics of Transport Phenomena* arXiv.
- Gers, F. A., Schmidhuber, J., and Cummins, F. (2000). Learning to Forget: Continual Prediction with Lstm. *Neural Comput.* 12, 2451–2471. doi:10.1162/089976600300015015
- Goodfellow, I., Pouget-Abadie, J., and Mirza, M., (2014). “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems*, 27, 2672–2680.
- Guo, X., Li, W., and Iorio, F. (2016). “Convolutional Neural Networks for Steady Flow Approximation,” in KDD '16: The 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, California (New York: Association for Computing Machinery), 481–490. doi:10.1145/2939672.2939738
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep Residual Learning for Image Recognition,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, Nevada (Manhattan, New York: IEEE). doi:10.1109/cvpr.2016.90
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). “Delving Deep into Rectifiers: Surpassing Human-Level Performance on Imagenet Classification,” in Proceedings of the IEEE International Conference on Computer Vision (ICCV). doi:10.1109/iccv.2015.123
- Hochreiter, S., and Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Comput.* 9, 1735–1780. doi:10.1162/neco.1997.9.8.1735
- Ioffe, S., and Szegedy, C. (2015). “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” in *Proceedings of Machine Learning Research* PMLR (Proceedings of Machine Learning Research).
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). “Image-to-image Translation with Conditional Adversarial Networks,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, Hawaii (Manhattan, New York: IEEE). doi:10.1109/cvpr.2017.632
- Kim, B., Azevedo, V. C., Thuerey, N., Kim, T., Gross, M., and Solenthaler, B. (2019). Deep Fluids: A Generative Network for Parameterized Fluid Simulations. *Comp. Graphics Forum* 38, 59–70. doi:10.1111/cgf.13619
- Lai, W.-S., Huang, J.-B., Ahuja, N., and Yang, M.-H. (2017). “Deep Laplacian Pyramid Networks for Fast and Accurate Super-resolution,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, Hawaii (Manhattan, New York: IEEE). doi:10.1109/cvpr.2017.618
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. (2013). “Rectifier Nonlinearities Improve Neural Network Acoustic Models,” in *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*. Editors S. Dasgupta and D. McAllester JMLR.
- Meng, X., Li, Z., Zhang, D., and Karniadakis, G. E. (2020). Ppinn: Parareal Physics-Informed Neural Network for Time-dependent Pdes. *Comp. Methods Appl. Mech. Eng.* 370, 113250. doi:10.1016/j.cma.2020.113250
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level Control through Deep Reinforcement Learning. *Nature* 518, 529–533. doi:10.1038/nature14236
- Orszag, S. A., and Tang, C.-M. (1979). Small-scale Structure of Two-Dimensional Magnetohydrodynamic Turbulence. *J. Fluid Mech.* 90, 129–143. doi:10.1017/s002211207900210x
- Radford, A., Metz, L., and Chintala, S. (2016). “Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks,” in 4th International Conference on Learning Representations, San Juan, Puerto Rico (arXiv).
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). 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. doi:10.1016/j.jcp.2018.10.045
- Roe, P. L. (1981). Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes. *J. Comput. Phys.* 43, 357–372. doi:10.1016/0021-9991(81)90128-5
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2017). “Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, Hawaii (Manhattan, New York: IEEE). doi:10.1109/cvpr.2017.437
- van Leer, B. (1979). Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov’s Method. *J. Comput. Phys.* 32, 101–136. doi:10.1016/0021-9991(79)90145-1
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image Quality Assessment: from Error Visibility to Structural Similarity. *IEEE Trans. Image Process.* 13, 600–612. doi:10.1109/TIP.2003.819861
- Wu, Y., Schuster, M., and Chen, Z., (2016). *Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation* arXiv. e-prints , arXiv:1609.08144.
- Xie, Y., Franz, E., Chu, M., and Thuerey, N. (2018). Tempogan: A Temporally Coherent, Volumetric gan for Super-resolution Fluid Flow. *ACM Trans. Graph.* 37, 1–15. doi:10.1145/3197517.3201304

Conflict of Interest: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Kigure. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.