Evaluation of Image-to-Image Surrogate Models for Film-Thickness Distribution Prediction in Organic-Film Deposition for OLED Lighting

Yuta Yamaguchi*†, Takaaki Miura†, and Masashi Nishiyama*
*Graduate School of Engineering, Tottori University, Tottori, Japan
†Kaneka Corporation, Tokyo, Japan
Email: d25t5001z@edu.tottori-u.ac.jp

***Abstract*—To our knowledge, this is the first study to evaluate neural operators alongside generative adversarial networks and diffusion models as image-to-image surrogates for predicting relative film-thickness distributions in organic light-emitting diode (OLED) lighting fabrication. Although fast surrogate prediction could reduce the cost of evaluating diffuser-plate candidates, the relative suitability of models developed for image translation and physical-field prediction remains unclear for this task. Using 1,650 simulation-generated image pairs covering three types of hole arrangements, we evaluated direct mapping from diffuser-plate images to thickness-distribution images normalized independently to represent relative in-plane variations. The evaluation combined five-fold cross-validation of nine methods using mean squared error (MSE), inspection of representative predictions for identical inputs, and a timing comparison between the most accurate surrogate and fluid-flow simulation. The factorized Fourier neural operator (F-FNO) achieved the lowest mean MSE of 124 ± 48 (mean $\pm$ standard deviation across folds), while its trained model required an average of 0.032 s per image with a batch size of 8, compared with $11{,}880$ s for fluid-flow simulation under the respective measurement conditions.**

***Index Terms*—Surrogate model, Diffuser plate, Thickness distribution, OLED lighting**

I. Introduction

Organic light-emitting diode (OLED) lighting is an area light source that emits light across its entire surface. The spatial distribution of indoor lighting is related to perceived brightness and discomfort glare [1], while maintaining uniform luminance over a large OLED surface remains challenging [2]. Because organic-film thickness affects the emission characteristics of OLED lighting, suppressing its in-plane variation is important.

Organic films for OLED lighting are formed by vacuum deposition, in which vaporized organic material is transported as a gas and deposited on a glass substrate. A perforated metal plate, called a diffuser plate, is used in the deposition system to control the material-gas flow. The hole diameter and arrangement of the diffuser plate affect the resulting film-thickness distribution. In this paper, an image representing these characteristics is referred to as a diffuser-plate image.

Fluid-flow simulation is used to evaluate candidate diffuser plates but requires the construction of a geometry-specific fluid model, mesh generation, and lengthy numerical computation. These processes must be repeated whenever the hole diameter or arrangement is changed, making it difficult to efficiently evaluate numerous candidates. Faster evaluation could improve diffuser-plate design efficiency for consumer OLED lighting and reduce the computational resources required for evaluation.

In this study, a thickness-distribution image is defined as a grayscale image obtained by normalizing the film-thickness values separately for each image. It therefore represents relative rather than absolute in-plane variations in film thickness. The contributions of this study are threefold. First, we formulate the task of predicting a thickness-distribution image from a diffuser-plate image and construct a dataset of 1,650 image pairs using fluid-flow simulation. Second, we construct image-to-image surrogate models that directly predict a thickness-distribution image from a diffuser-plate image without performing fluid-flow simulation during inference. Third, we compare the prediction accuracy of nine methods based on generative adversarial networks (GANs), diffusion models (DMs), and neural operators (NOs), analyze representative predictions, and compare the computation times of the best-performing method and fluid-flow simulation.

II. Related Work

A. Surrogate Models

Data-driven surrogate models approximate the input–output relationships of computationally expensive simulations and have been applied to fluid-flow problems. Bhatnagar et al. [3] proposed a convolutional-neural-network-based surrogate model for predicting flow fields around airfoils.

For deposition equipment, Jin et al. [4] divided a showerhead into 20 regions and predicted a flow nonuniformity index from the opening ratio of each region and two fillet radii. These inputs represent aggregated spatial information about the openings but do not directly represent the hole diameter or the positions of individual holes. In contrast, our model takes a diffuser-plate image that directly represents the hole diameter and individual-hole positions and predicts a thickness-distribution image over the glass substrate.

B. Methods Used for Image-to-Image Surrogate Models

For mapping diffuser-plate images to thickness-distribution images, we compare methods developed for image-to-image

translation with methods developed to learn mappings between physical fields. Methods based on generative adversarial networks (GANs) and diffusion models (DMs) are used for image-to-image translation. Methods based on neural operators (NOs) are used to learn mappings between physical fields.

GAN-based image-to-image translation methods include Pix2Pix [5], BicycleGAN [6], and CycleGAN [7]. Pix2Pix learns a mapping from an input image to an output image using paired training images and a conditional GAN. The Pix2Pix generator employs a U-Net architecture. Its skip connections transfer spatial information from the input image to the output image. BicycleGAN introduces latent variables to generate diverse output images from a single input image. CycleGAN uses cycle consistency to learn mappings between two unpaired image domains.

BBDM [8] represents the transformation between paired input and output images as a Brownian bridge diffusion process. I2SB [9] uses a Schrödinger bridge to represent the transformation between the input- and output-image distributions. DiffI2I [10] pretrains features required for image-to-image translation and learns a diffusion model that estimates output-image features from an input image.

Neural operators are designed to learn operators that map input functions to output functions. In this study, diffuser-plate images and thickness-distribution images are represented as arrays on a fixed grid, and NOs are applied to learn the mapping between them. The Fourier neural operator (FNO) [11] parameterizes an integral kernel in Fourier space to capture global relationships between input and output fields. U-FNO [12] adds a U-Net path to each U-Fourier layer to complement local features and high-frequency components. The factorized Fourier neural operator (F-FNO) [13] factorizes spectral operations along the spatial axes and combines the resulting features. F-FNO updates the combined features through feed-forward transformations with residual connections.

In this study, nine methods based on GANs, DMs, and NOs are applied as image-to-image surrogate models. We compare their accuracy in predicting thickness-distribution images. We then select the method with the highest overall prediction accuracy and the most accurate GAN-based method. Their predictions for the same cases are compared to discuss the relationship between prediction characteristics and model architectures.

III. Diffuser-Plate and Thickness-Distribution Dataset

A. Fluid-Flow Simulation Process

Fig. 1 shows the fluid-flow simulation process for generating a thickness-distribution image from a diffuser-plate image. The input to the fluid-flow simulation is a diffuser-plate image representing the hole diameter and arrangement. Except for the diffuser plate, the geometries of the other components and the simulation conditions were kept unchanged for all diffuser-plate images.

In Process A, a fluid model corresponding to the diffuser-plate image is constructed. In Process B, the computational

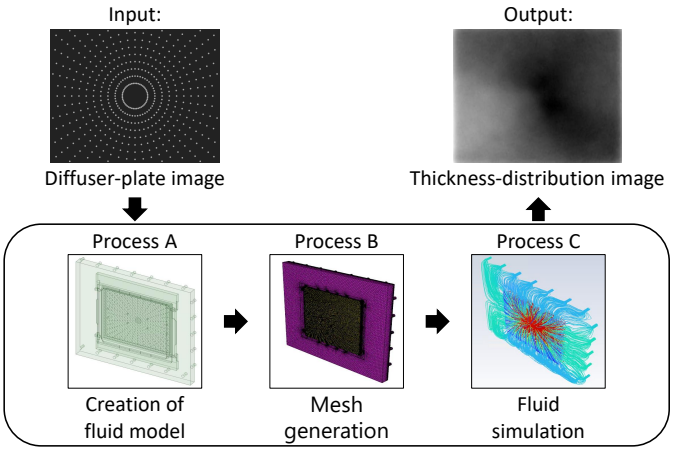

Fig. 1. Fluid-flow simulation process for generating a thickness-distribution image from a diffuser-plate image.

domain of the fluid model is divided into numerous computational cells to generate a mesh for the numerical calculation. In Process C, the material-gas flow is numerically calculated using ANSYS Fluent. Steady-state conditions are assumed for the material-gas flow. The thickness of the organic film formed on the glass substrate is calculated from the simulation results. Finally, the calculated thickness values are sampled at specified coordinates on the glass substrate to generate a thickness-distribution image.

B. Fluid Model

The fluid model represents the space through which the material gas flows inside the deposition system as the computational domain. In Process A shown in Fig. 1, the regions occupied by the system components are excluded to construct the fluid model.

Fig. 2(a) shows the converge plate, diffuser plate, and heating plate. Fig. 2(b) shows a cross-sectional view of the deposition system containing these components. The material gas enters the pipe through the inlet and is supplied to the deposition system through the shower ring. The converge plate concentrates the supplied material gas. The diffuser plate disperses the material gas passing through the converge plate in directions parallel to the glass substrate. The material-gas flow changes according to the hole diameter and arrangement of the diffuser plate. The dispersed material gas receives thermal energy from the heating plate. The material gas forms an organic film as it cools on the surface of the glass substrate. Material gas that is not used to form the organic film is discharged through the outlet.

Fig. 2(c) shows the fluid model obtained by extracting the space through which the material gas flows. The fluid model consists of an inlet region, a fluid region, and an outlet region.

Whenever the hole diameter or arrangement of the diffuser plate is changed, the fluid model must be reconstructed. A computational mesh must also be generated for each reconstructed fluid model. A finer mesh improves the accuracy of

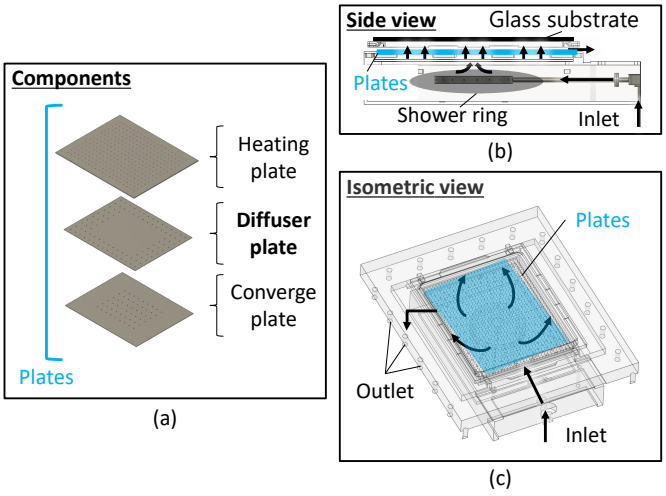

Fig. 2. Overview of the fluid model: (a) components of the deposition apparatus, (b) cross-sectional view of the apparatus, and (c) computational domain for the material-gas flow.

the thickness-distribution calculation. However, increasing the number of computational cells also increases the numerical computation time. Therefore, generating thickness-distribution images for numerous diffuser plates requires substantial computational cost.

C. Construction of the Diffuser-Plate and Thickness-Distribution Dataset

To construct the paired-image dataset, diffuser-plate images with different hole diameters and arrangements were first created. The hole diameters ranged from 2.0 to 5.0 mm, and all holes in each diffuser plate had the same diameter. Three types of hole arrangements were used: circular, grid, and radial arrangements. Fig. 3 shows examples of the three hole arrangements.

For the circular and grid arrangements, Pattern 1 used the center of the diffuser plate as a single reference coordinate, around which a single group of holes was arranged. The spatial extent of the hole group in Pattern 1 varied according to the generation parameters for each arrangement; therefore, the holes were sometimes localized near the center of the diffuser plate. In Patterns 2 and 3, the diffuser plate was divided into a 4×5 array of regions, and identical groups of holes were repeatedly arranged with respect to the center of each region. For the circular arrangement, holes were placed at equal intervals along concentric circles centered at each reference coordinate. For the grid arrangement, holes were placed at equal intervals in the vertical and horizontal directions on a grid centered at each reference coordinate. For the radial arrangement, holes were placed at equal intervals along lines extending outward from the center of the diffuser plate, and the number and angles of the lines were varied among the three patterns.

The generation parameters for each hole arrangement were randomly varied to create 550 diffuser-plate images per arrangement, yielding a total of 1,650 images. In the diffuser-plate images, pixels corresponding to the holes were assigned a value of 255, whereas the remaining pixels were assigned a value of 128.

Next, a fluid model corresponding to each diffuser-plate image was constructed, and the computational domain was divided into numerous cells to generate a mesh. Fluid-flow simulation was performed using each generated mesh to calculate the thickness of the organic film formed on the glass substrate. Within the area of the glass substrate used for the OLED lighting panel, the thickness values were sampled on a uniform grid at 1 mm intervals. The sampled thickness values were arranged on a two-dimensional plane to generate a thickness-distribution image.

For each image, the thickness values were linearly normalized to pixel values ranging from 0 to 254, such that the maximum and minimum thickness values corresponded to pixel values of 0 and 254, respectively. Pixels for which no thickness value was available were assigned a value of 255. Thus, among the pixels representing thickness values, darker pixels indicate relatively thicker regions, whereas lighter pixels indicate relatively thinner regions. The ordering of the pixel values is therefore reversed relative to that of the thickness values output by the fluid-flow simulation.

Finally, the diffuser-plate image and thickness-distribution image were each resized to 256×256 pixels and stored as an image pair. Bilinear interpolation was used for image resizing.

Fig. 4 shows six image pairs, with two selected from each hole arrangement. In this figure, (a) and (b) show examples of the circular arrangement, (c) and (d) show examples of the grid arrangement, and (e) and (f) show examples of the radial arrangement. Each example shows a diffuser-plate image on the left and the corresponding thickness-distribution image obtained by fluid-flow simulation on the right. The constructed dataset included cases in which the thickness-distribution image varied with the hole diameter and arrangement, as well as cases in which different diffuser-plate images produced similar thickness-distribution images.

IV. Evaluation of Image-to-Image Surrogate Models

A. Prediction Processes of the Image-to-Image Surrogate Models

Fig. 5 schematically illustrates the prediction processes of the three categories of image-to-image surrogate models investigated in this study. Each model takes a diffuser-plate image representing the hole diameter and arrangement as input and directly outputs the corresponding thickness-distribution image. Here, because the film-thickness values are normalized separately for each thickness-distribution image, the model predicts relative in-plane variations in film thickness over the glass substrate rather than absolute film thickness. In this study, prediction accuracy is evaluated using thickness-distribution images obtained by fluid-flow simulation as reference images to investigate whether the model can serve as an

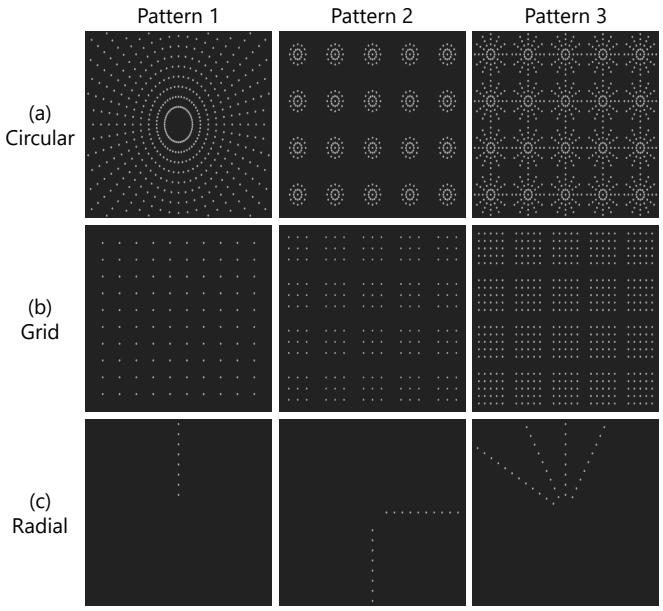

Fig. 3. Examples of hole arrangements in diffuser plates. (a) Circular arrangement. (b) Grid arrangement. (c) Radial arrangement.

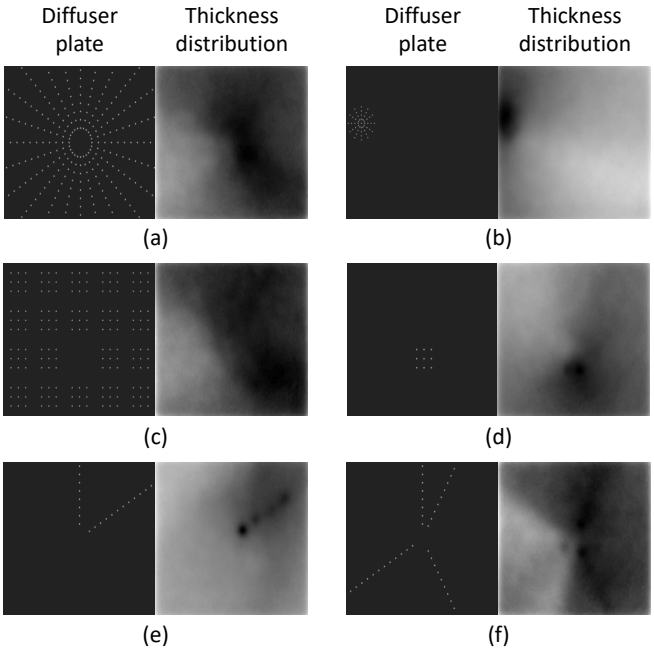

Fig. 4. Examples of diffuser-plate images and thickness-distribution images in the paired-image dataset. (a), (b) Circular arrangement. (c), (d) Grid arrangement. (e), (f) Radial arrangement.

alternative to fluid-flow simulation for comparing thickness-distribution images corresponding to different diffuser-plate geometries.

B. Comparison Methods

To evaluate how the choice of method affects the prediction accuracy of thickness-distribution images, we compared nine methods based on GANs, DMs, and NOs. Unless otherwise

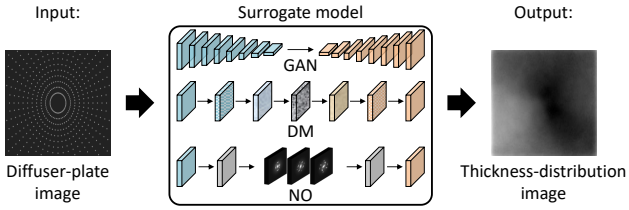

Fig. 5. Schematic prediction processes of the image-to-image surrogate models.

specified, the basic settings for each method were taken from the values reported in the original paper or used in the corresponding publicly available GitHub implementation.

The GAN-based methods were Pix2Pix [5], BicycleGAN [6], and CycleGAN [7]. For Pix2Pix, the weight assigned to the L1 loss was set to 100. For BicycleGAN, the cLR-GAN variant was used. For CycleGAN, the image pool was not used.

The DM-based methods were BBDM [8], I2SB [9], and DiffI2I [10]. For BBDM, the latent dimension, downsampling factor, and codebook size were set to 4, 8, and 16,384, respectively. For I2SB, the number of filters and the time-embedding dimension were set to 48 and 256, respectively. For DiffI2I, the input and output image sizes were set to 128×128 pixels because of implementation constraints.

The NO-based methods were FNO [11], U-FNO [12], and F-FNO [13]. For each method, the numbers of input and output channels were both set to 1. For the eight methods other than DiffI2I, the input and output image sizes were set to 256×256 pixels.

C. Evaluation Method

The constructed paired-image dataset consisting of 1,650 pairs was divided into five subsets of 330 pairs, and five-fold cross-validation was performed. In each trial, 1,320 pairs from four subsets were used as training data, and the remaining 330 pairs from one subset were used as test data. Training and evaluation were repeated five times, using a different subset as the test data in each trial. Owing to GPU memory constraints, the batch size was set to 4 for DiffI2I and 8 for the other eight methods. The eight methods other than DiffI2I were trained for 1,000 epochs. For DiffI2I, pretraining and subsequent diffusion-model training were each performed for 1,000 epochs.

Mean squared error (MSE) was used to evaluate prediction accuracy. The MSE was calculated as the mean of the squared pixel-wise differences between the predicted image and the reference thickness-distribution image obtained by fluid-flow simulation. For DiffI2I, the MSE was calculated using reference and predicted images with a resolution of 128×128 pixels. For the other eight methods, the MSE was calculated using reference and predicted images with a resolution of 256×256 pixels.

For the eight methods other than DiffI2I, the MSE was calculated every 100 epochs during the 1,000-epoch training. For DiffI2I, the MSE was calculated every 100 epochs dur-

ing diffusion-model training following pretraining. At each evaluation epoch, the MSE on the test data was calculated separately for each fold and averaged across the five folds. For each method, the evaluation epoch with the lowest mean MSE between 100 and 1,000 epochs was selected, and the mean and standard deviation (SD) of the five fold-wise MSE values at that epoch were reported.

D. Prediction Accuracy for Thickness-Distribution Images

Table I presents the mean and standard deviation (SD) of the MSE across the five folds for each image-to-image surrogate model. The mean and SD values are rounded to the nearest integers. Among the GAN-, DM-, and NO-based methods, Pix2Pix, BBDM, and F-FNO achieved the lowest mean MSEs of 147 ± 21, 382 ± 33, and 124 ± 48, respectively. F-FNO achieved the lowest mean MSE among all nine methods. However, the mean $\pm$ SD ranges of F-FNO, U-FNO, and Pix2Pix overlapped; therefore, the present evaluation does not establish whether the differences among these methods are statistically significant.

For the image-wise comparison, we selected F-FNO, which achieved the lowest mean MSE among all methods, and Pix2Pix, which achieved the lowest mean MSE among the GAN-based methods. Fig. 6 compares the thickness-distribution images obtained by fluid-flow simulation, F-FNO, and Pix2Pix for the same diffuser-plate images. All images are displayed using the same grayscale range, and the MSE of each prediction, rounded to the nearest integer, is shown below the image. Fig. 6(a) shows a case in which both methods yielded low MSEs. Figs. 6(b) and (c) show cases in which F-FNO yielded lower MSEs than Pix2Pix, whereas Figs. 6(d) and (e) show cases in which Pix2Pix yielded lower MSEs than F-FNO. Fig. 6(f) shows a case in which both methods yielded high MSEs.

In Fig. 6(a), both methods reproduced the relatively thick region near the center and the overall grayscale distribution, each yielding an MSE of 9. In Figs. 6(b) and (c), F-FNO reproduced the locations of the relatively thick regions and the surrounding grayscale distributions more closely than Pix2Pix. In contrast, the Pix2Pix predictions were closer to the fluid-flow simulation results in Figs. 6(d) and (e). In particular, the F-FNO prediction in Fig. 6(e) contained horizontal grayscale variations that were absent from the fluid-flow simulation result. In Fig. 6(f), both methods reproduced the relatively thick region near the center, but differences remained in the surrounding grayscale distributions.

The reproduction of broad grayscale distributions by F-FNO in Figs. 6(b) and (c) is consistent with its architecture, which uses operations in Fourier space to capture relationships between spatially distant regions. However, this ability to capture global relationships is shared by FNO and U-FNO and therefore does not by itself explain why F-FNO achieved the lowest mean MSE among the three NO-based methods. Furthermore, the present comparison did not separately evaluate the contributions of the factorized spectral operations and residual updates in F-FNO.

TABLE I
MSE ACROSS THE FIVE FOLDS (MEAN $\pm$ SD).

Category	Method	MSE $\downarrow$
GAN	Pix2Pix	147 ± 21
GAN	BicycleGAN	193 ± 37
GAN	CycleGAN	$5{,}599 \pm 292$
DM	BBDM	382 ± 33
DM	I2SB	$4{,}882 \pm 451$
DM	DiffI2I	879 ± 61
NO	FNO	437 ± 144
NO	U-FNO	136 ± 23
NO	F-FNO	$\mathbf{124 \pm 48}$

The approximate localization of relatively thick regions by Pix2Pix is also consistent with its U-Net generator, whose skip connections transfer spatial information from the input to the output. Nevertheless, the relative MSE values of F-FNO and Pix2Pix varied among the input diffuser-plate images. Table I reports prediction accuracy over the entire dataset, whereas Fig. 6 presents cases selected to illustrate differences between the predictions of the two methods. Therefore, the number of displayed cases does not represent their frequency in the evaluation dataset.

E. Time Comparison Between Fluid-Flow Simulation and the Image-to-Image Surrogate Model

The mean time required per thickness-distribution image was compared between fluid-flow simulation and the image-to-image surrogate model. F-FNO, which achieved the lowest mean MSE among the nine methods, was used as the image-to-image surrogate model. This comparison evaluated the time required to generate a thickness-distribution image by fluid-flow simulation and the prediction time of F-FNO. The time required to prepare the input diffuser-plate images, generate the training dataset, and train F-FNO was excluded. Therefore, this comparison represents the evaluation time per additional candidate using a trained F-FNO model, rather than the total development time from dataset construction through model training.

The time required for fluid-flow simulation was measured using five diffuser plates randomly selected from the dataset. ANSYS Fluent was used for fluid-flow simulation, and the computations were performed using an Intel Xeon w5-3435X CPU. One CPU core was used for Process A shown in Fig. 1, whereas 12 CPU cores were used for Processes B and C. Prediction using F-FNO was performed on an NVIDIA GeForce RTX 4090 GPU with a batch size of 8. Because different computing hardware was used for fluid-flow simulation and F-FNO, the results apply to the measurement conditions described above.

Table II presents the mean times for fluid-flow simulation and F-FNO. The mean time required to generate one thickness-distribution image by fluid-flow simulation was 11,880 s. The mean prediction time per thickness-distribution image using F-FNO was 0.032 s. This value was calculated as the mean time per image from inference performed with a batch size of 8. Under the measurement conditions used in this study, the

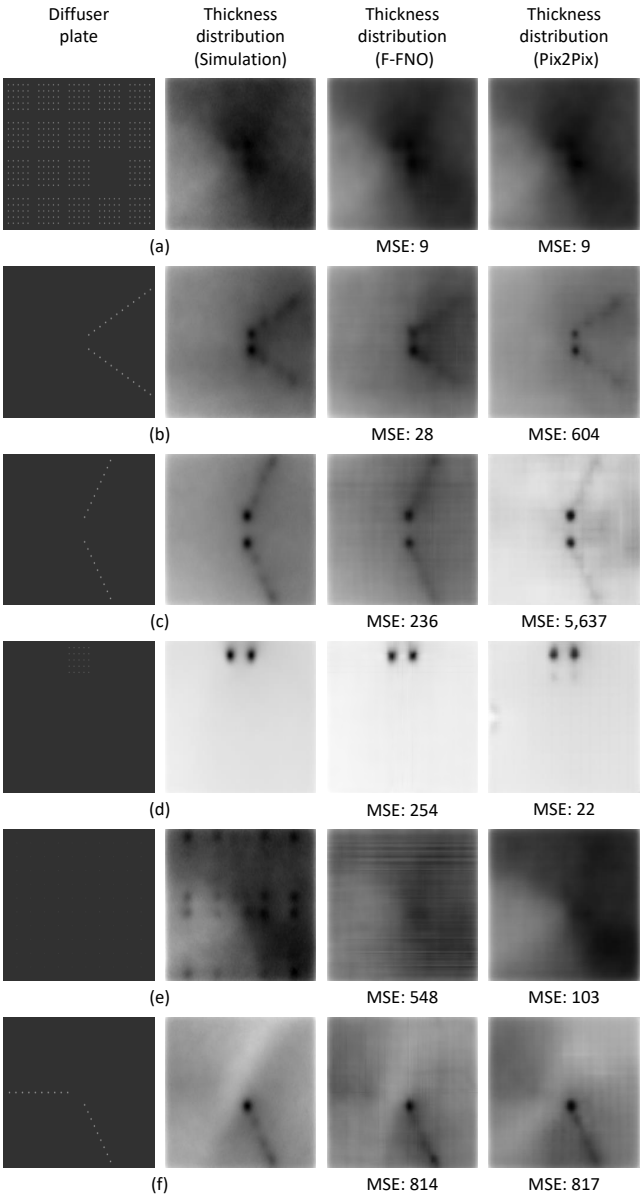

Fig. 6. Thickness-distribution predictions by F-FNO and Pix2Pix for the same inputs. (a) Low MSEs for both methods. (b), (c) Lower MSEs for F-FNO. (d), (e) Lower MSEs for Pix2Pix. (f) High MSEs for both methods. The number of examples does not represent their frequency in the evaluation dataset.

TABLE II
MEAN TIME PER THICKNESS-DISTRIBUTION IMAGE (F-FNO: RTX 4090, BATCH SIZE 8).

Method	Time [s]
Fluid-flow simulation	11,880
F-FNO-based surrogate model	0.032

trained F-FNO model generated thickness-distribution images in less time than fluid-flow simulation.

V. CONCLUSION

In this study, we investigated image-to-image surrogate models for predicting relative thickness-distribution images from diffuser-plate images. We constructed a paired-image dataset consisting of 1,650 image pairs using fluid-flow simulation and compared nine methods based on GANs, DMs, and NOs.

Five-fold cross-validation showed that F-FNO achieved the lowest MSE of 124 ± 48 (mean $\pm$ SD) among the nine methods. A comparison of F-FNO and Pix2Pix using the same diffuser-plate images showed that the prediction accuracy for individual images varied depending on the input. Under the measurement conditions used in this study, the mean prediction time of the trained F-FNO model was 0.032 s per image, which was shorter than the mean fluid-flow simulation time of 11,880 s.

Future work will analyze the relationship between the geometric characteristics of diffuser-plate images and prediction errors and compare the predictions with the thickness distributions of organic films formed in an actual deposition system.

REFERENCES

[1] M. G. Figueiro, J. D. Bullough, A. Thayer, R. Nagare, and M. S. Rea, "Supporting visual and non-visual lighting design without increasing discomfort glare or lighting power density," *Light. Res. Technol.*, vol. 57, no. 1, pp. 5–27, 2025.

[2] M. Riahi, K. Yoshida, H. Hafeez, and I. D. W. Samuel, "Improving the uniformity of top emitting organic light emitting diodes using a hybrid electrode structure," *Adv. Electron. Mater.*, vol. 10, no. 1, 2024, Art. no. 2300675.

[3] S. Bhatnagar, Y. Afshar, S. Pan, K. Duraisamy, and S. Kaushik, "Prediction of aerodynamic flow fields using convolutional neural networks," *Comput. Mech.*, vol. 64, no. 2, pp. 525–545, 2019.

[4] Z. Jin, D. D. Lim, X. Zhao, M. Mamunuru, S. Roham, and G. X. Gu, "Machine learning enabled optimization of showerhead design for semiconductor deposition process," *J. Intell. Manuf.*, vol. 35, no. 2, pp. 925–935, 2024.

[5] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 1125–1134.

[6] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman, "Toward multimodal image-to-image translation," in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.

[7] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 2223–2232.

[8] B. Li, K. Xue, B. Liu, and Y.-K. Lai, "BBDM: Image-to-image translation with Brownian bridge diffusion models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 1952–1961.

[9] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, "I2SB: Image-to-image Schrödinger bridge," in *Proc. 40th Int. Conf. Mach. Learn. (ICML)*, 2023, pp. 22 042–22 062.

[10] B. Xia, Y. Zhang, S. Wang, Y. Wang, X. Wu, Y. Tian, W. Yang, R. Timofte, and L. Van Gool, "DiffI2I: Efficient diffusion model for image-to-image translation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 47, no. 3, pp. 1578–1593, 2025.

[11] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. M. Stuart, and A. Anandkumar, "Fourier neural operator for parametric partial differential equations," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021.

[12] G. Wen, Z. Li, K. Azizzadenesheli, A. Anandkumar, and S. M. Benson, "U-FNO: An enhanced Fourier neural operator-based deep-learning model for multiphase flow," *Adv. Water Resour.*, vol. 163, 2022, Art. no. 104180.

[13] A. Tran, A. Mathews, L. Xie, and C. S. Ong, "Factorized Fourier neural operators," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023.