
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Haitsiukevich, Katsiaryna; Poyraz, Onur; Marttinen, Pekka; Ilin, Alexander

Diffusion models as probabilistic neural operators for recovering unobserved states of dynamical systems

Published in:

34th IEEE International Workshop on Machine Learning for Signal Processing, MLSP 2024 - Proceedings

DOI:

[10.1109/MLSP58920.2024.10734762](https://doi.org/10.1109/MLSP58920.2024.10734762)

Published: 01/09/2024

Document Version

Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:

Haitsiukevich, K., Poyraz, O., Marttinen, P., & Ilin, A. (2024). Diffusion models as probabilistic neural operators for recovering unobserved states of dynamical systems. In *34th IEEE International Workshop on Machine Learning for Signal Processing, MLSP 2024 - Proceedings* (IEEE International Workshop on Machine Learning for Signal Processing). IEEE. <https://doi.org/10.1109/MLSP58920.2024.10734762>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

DIFFUSION MODELS AS PROBABILISTIC NEURAL OPERATORS FOR RECOVERING UNOBSERVED STATES OF DYNAMICAL SYSTEMS

Katsiaryna Haitsiukevich, Onur Poyraz, Pekka Marttinen, Alexander Ilin

Department of Computer Science, Aalto University, Espoo, Finland

ABSTRACT

This paper explores the efficacy of diffusion-based generative models as neural operators for partial differential equations (PDEs). Neural operators are neural networks that learn a mapping from the parameter space to the solution space of PDEs from data, and they can also solve the inverse problem of estimating the parameter from the solution. Diffusion models excel in many domains, but their potential as neural operators has not been thoroughly explored. In this work, we show that diffusion-based generative models exhibit many properties favourable for neural operators, and they can effectively generate the solution of a PDE conditionally on the parameter or recover the unobserved parts of the system. We propose to train a single model adaptable to multiple tasks, by alternating between the tasks during training. In our experiments with multiple realistic dynamical systems, diffusion models outperform other neural operators. Furthermore, we demonstrate how the probabilistic diffusion model can elegantly deal with systems which are only partially identifiable, by producing samples corresponding to the different possible solutions.

Index Terms— Diffusion Models, Neural Operator, Physical Systems Modelling, Partial Differential Equations

1. INTRODUCTION

Physical dynamical systems can be expressed with partial differential equations (PDEs) and traditionally solved using methods like the finite difference method. This conventional approach cannot easily incorporate measurements of the process, which motivates the development of neural network methods that can combine the prior knowledge with data-driven learning. Differential equations can be combined with neural networks by integrating the PDEs directly into the architecture [1, 2] or by including the PDE error in the loss to represent prior knowledge [3]. The first option complies with the underlying physics by design, but is system-specific and non-trivial to extend. In the second approach the architecture

does not depend on the equations but the training requires the functional form of the equations and knowledge of the parameters, which may be unavailable. Alternatively, it is possible to incorporate domain knowledge through a physics simulator, which can be integrated into the network architecture [4] or utilized to augment the training data [5]. The simulator allows to collect a rich data set from the underlying process by randomizing its parameters, which alleviates the need for system identification at inference time, enables a wider class of models, and increases robustness against noise and distortions in the data. If the PDE is known at test time, it can be used to fine-tune the model’s prediction [6].

Neural operators are data-driven and learn a mapping from the input parameters or observations to the output. Most of them are deterministic, which suffices for approximating the mapping from the PDE parameters to the solution. However, for the inverse problems of recovering parameters or reconstructing an unobserved state, deterministic predictions may not capture all possible outcomes. To resolve this, we opt for probabilistic generative models, the diffusion models [7, 8], which have excelled in many domains. We view the system state evolution, discretized across spacial and temporal dimensions, as an image with state variables stacked as channels. Consequently, the reconstruction of an unobserved state is analogous to a color restoration task. The diffusion models can be pre-trained on complete image data, and conditioning applied only at inference time [9]. This ‘unconditional’ approach, called RePaint in [9], is adaptable to many use cases, but with increased inference complexity. Alternatively, the model’s input can be augmented with the conditioning information already during training [10, 11], yielding a specialized model suitable for a single task. A similar approach has been used for multivariate time series prediction and imputation [12].

Our novel contributions are: **1)** We define and compare multiple diffusion models for unobserved state reconstruction and prediction. Unlike existing works in this domain [13, 14], we train our model with a ‘mixed conditional’ objective, allowing it to solve multiple tasks with a single training (see Fig. 1). **2)** We empirically demonstrate the strong performance of diffusion models in general, and the model with mixed conditional training in particular, against other neural

E-mail: {firstname.lastname}@aalto.fi

This work was supported by the Research Council of Finland (Flagship programme: Finnish Center for Artificial Intelligence FCAI, and grants 352986, 358246) and EU (H2020 grant 101016775 and NextGenerationEU). We thank Ti John and Nicola Dainese for the help with graphics design.

operator baselines. 3) We study systems which are not fully identifiable, and demonstrate that diffusion models can represent the variability of possible solutions.

2. RELATED WORK

Neural operators [15] have strong performance in modelling of physical systems. State-of-the-art results on rectangular domains are produced by FNO [13]. Recently proposed neural operators use the transformer architecture, e.g., OFormer [14]. Existing methods solve either the forward (parameters to solution) or the inverse (solution to parameters) problem. Unlike the neural operator approaches, a diffusion model can solve both of these problems at once by conditioning with known information. Concurrent to our work, [16] combines forward and inverse operators but in a non-diffusion model.

Diffusion models for physical systems have been successfully applied for predicting the dynamics of the system in the future, either by autoregressive unrolls [17] or by taking the desired future time as a model input [18]. These models focus on prediction only while we consider training jointly for prediction and reconstruction. A work [19] utilized the diffusion model for data generation for several physical systems. It focused on unconditional generation that adheres to physical constraints of the system while we recover the unobserved states conditioned on partial information.

Incorporation of PDE prior information into diffusion models. A physics-informed diffusion model [20] includes a scaled PDE residual into the training objective, encouraging solutions that comply with the PDE. Alternatively, the gradient of the PDE residual guide the sampling [21] or be directly fed into the model as an extra input channel during training [22]. All of the approaches can be applied together with the methods in this paper.

3. METHOD

3.1. Background on diffusion models

Denoising diffusion probabilistic models (DDPMs) [8] are a class of generative models which learn underlying data distribution by progressively removing noise from the distorted input. In particular, the training procedure requires defining a schedule for the magnitude of noise added to model input at every training iteration. The highest noise magnitude in the schedule makes the model input close to pure Gaussian noise. A neural network, predicting a denoised version of the input, is trained by optimizing a variational lower bound or a simplified objective [8] between the added noise and the network prediction, or some other refined combinations (see, e.g., [23, 24]). Denoising diffusion models use the U-net architecture [25] with conditioning on the step t of the noise schedule [8]. The original architecture has been made more expressive by improvements in residual blocks [26, 27] while

inefficiencies in training dynamics have been remedied by input and output preconditioning [24].

One downside of diffusion models is the slow inference, which, starting from random noise, iteratively denoises the input into the final prediction. The original DDPM [8] sampling uses the same number of steps as the training noise schedule (typically around 1000 steps). However, the number of inference steps can be significantly decreased by defining a separate inference schedule [23, 28, 24]. The approaches [23, 28] are uniformly skipping some steps in the training noise schedule. However, a better strategy for allocating the inference resources is skipping more steps in the beginning of the sampling process and preserving more steps closer to the end where the noise magnitude is smaller. This idea combined with a second order Heun method has been proposed by [24] and together with input and output preconditioning is referred to as EDM.

3.2. Mixed conditional training

We aim at training a model capable of solving several tasks, such as the unobserved state reconstruction, the system dynamics prediction from the parameter value, the underlying parameter recovery from the dynamics, and the prediction of future system states. The parameter values can be concatenated with the system states as an additional input channel making parameter recovery and dynamics prediction analogous to the state reconstruction task. Similar problem formulation arises in image-to-image translation considered in Palette [10] and in time series imputation and forecasting addressed in SSSD^{S4} [12]. We propose to model the dynamics by a conditional diffusion model trained on a mixture of the tasks of interest as shown in Fig. 1.

To simultaneously train for several tasks, we generate a mask corresponding to the unknown part of the system (e.g. upper or lower channel in Fig. 1) for each training batch. For the observed channel(s), we sample the number of time steps observed (in the direction of the x-axis in Fig. 1). For example, Task 3 in Fig. 1 corresponds to observing both channels partially for the first half of the time span, while in Tasks 4 and 5 only one channel is partially observed.

Since a part of the full state is known, the diffusion process is applied only to the unobserved part. As shown in Fig. 2, the input for one denoising step consists of system states, such that noise is added to the unobserved part, and the observed part is kept noise free and given as conditioning information. The input is the same both at training and inference times. The conditioning information can be optionally accompanied by the corresponding binary masks as in SSSD^{S4} which we also found beneficial. To take into account that the denoising should be applied only to the unknown part, the observed part is masked out from the loss.

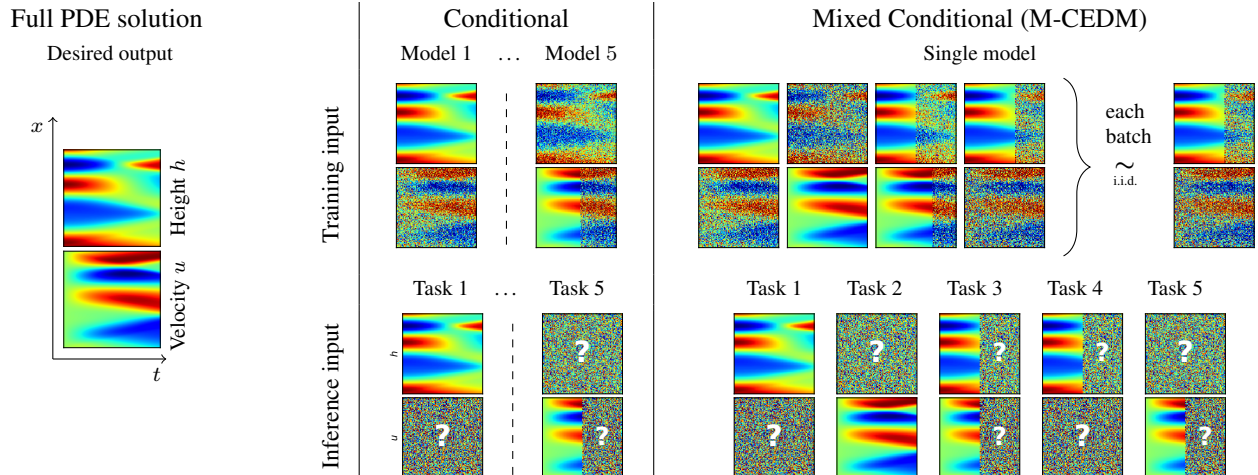


Fig. 1. Conditional and mixed conditional training of diffusion models demonstrated with the Shallow-Water Equation (SWE-orig). The system has two variables (channels) represented by the colored rectangles with time (t) and spatial coordinate (x) on the x - and y -axes. The ‘desired output’ is the full state which we train the model to reconstruct from partial information. The clean parts of the input represent the conditioning information (at training and inference time) and the noisy parts are reconstructed by denoising. Each **conditional** model is trained with conditioning information for a single task. The **mixed conditional** training yields a single model for all defined tasks by sampling one task for each mini-batch during training.




Fig. 2. Model input for mixed conditional training.




Fig. 3. Example simulations from two of the studied systems.

4. EXPERIMENTS

4.1. Dynamical systems considered

We demonstrate the performance of the proposed diffusion models on four systems. The first three: Darcy flow, shallow water equation with periodic initial conditions, and the fixed-bed tubular reactor are fully identifiable given the conditioning information. The fourth system is the shallow water equation with an initial perturbation of the water level, and it is partially non-identifiable, meaning that multiple solutions are compatible with the given conditioning information. Below we specify these systems and their data generation.

Darcy Flow equation is defined over a unit square with a 2D steady-state solution governed by

$$\nabla(a(x)\nabla u(x)) = f(x), x \in (0, 1)^2, \quad (1)$$

$$u(x) = 0, x \in \partial(0, 1)^2, \quad (2)$$

with viscosity $a(x)$ being piece-wise constant and the force term $f(x) = 1$. We use a subset of 1000 simulations from PDEBench [29], of which one instance is shown in Fig. 3(a).

Shallow water equation (SWE-orig) is derived from the compressible Navier-Stokes equation and describes free-surface fluid flow. The system is characterized by height h and velocity u of the fluid evolving as:

$$\frac{\partial h}{\partial t} = -\frac{\partial hu}{\partial x}, \quad \frac{\partial hu}{\partial t} = -\frac{\partial(u^2 h + 0.5gh^2)}{\partial x} \quad (3)$$

where g is the gravitational constant. Here, we set $g = 1$ and randomize the initial conditions for h . For $t \in [0, 0.128]$, $x \in [-0.5, 0.5]$, and $u_0 = 0$, h_0 is calculated as:

$$\tilde{h}(0, x) = \sum_{k=-N}^N \lambda_k \cos(2\pi kx) + \gamma_k \sin(2\pi kx), \quad (4)$$

$$h(0, x) = 1 + \frac{\tilde{h}(0, x) - \min(\tilde{h}(0, x))}{\max(\tilde{h}(0, x)) - \min(\tilde{h}(0, x))}, \quad (5)$$

where $N = 3$ and $\lambda_k, \gamma_k \sim \mathcal{N}(0, 1)$. An example of this system is used in the illustrations in Fig. 1.

Shallow water equation (SWE-init) is generated for $x \in [-2.5, 2.5]$ and $t \in [0, 1.28]$ with initial conditions:

$$h(0, x) = h_{\text{in}} + \epsilon \exp\left(\frac{(x - x_0)^2}{2\sigma^2}\right), \quad u(0, x) = u_0, \quad (6)$$

Table 1. Model comparison on state reconstruction for the forward (Task 1, columns 2 – 4) and inverse (Task 2, columns 5 – 7) problems in case of identifiable systems. For each problem we present the mean absolute error (MAE) and the PDE residual error. All MAE values should be multiplied by 10^{-3} .

Model	Darcy, $a \rightarrow u$		SWE, $h \rightarrow u$		Reactor, $x_p, T \rightarrow \theta, x_a$		Darcy, $u \rightarrow a$		SWE, $u \rightarrow h$		Reactor, $\theta, x_a \rightarrow x_p, T$	
	MAE	PDE	MAE	PDE	MAE	PDE	MAE	PDE	MAE	PDE	MAE	PDE
FNO [13]	7.55	38.08	0.44	0.9963	0.07	0.0146	3.38	3.236	2.21	1.1370	5.64	0.1132
OFormer [14]	6.82	21.00	0.31	0.9966	0.13	0.0083	4.06	6.683	0.81	0.9929	7.76	0.0282
RePaint [9]	9.80	6.909	0.69	1.0462	0.73	0.1521	191.8	84.11	9.68	1.0429	188.3	0.2017
CDDPM [8]	6.23	1.478	0.25	0.9925	0.09	0.0039	0.25	1.180	0.84	0.9906	4.62	0.0039
CEDM [24]	5.52	1.344	0.10	0.9923	0.05	0.0015	0.18	1.192	0.44	0.9924	3.42	0.0017
M-CEDM	5.46	1.272	0.10	0.9915	0.07	0.0020	0.17	1.180	0.43	0.9906	2.69	0.0017

where $h_{\text{in}} \sim \mathcal{U}(1.2, 5.2)$, $\epsilon \sim \mathcal{U}(0.05, 1)$, $x_0 \sim \mathcal{U}(-1, 1)$, $\sigma \sim \mathcal{U}(0.2, 2)$, $(hu)_0 \sim \mathcal{U}(-2.2, 2.2)$. This equation models an abrupt change in the water level, e.g. caused by a dam break. Depending on the initial water velocity u_0 , several possible outcomes may correspond to the same observation of the height h (Fig. 4).

For both SWE datasets the data are simulated using Py-Claw [30] Python package implementing a finite volume method for this equation.

Fixed-bed tubular reactor is a system with hydrogenation of aromatics with concentration x_a on the catalyst surface with activity θ . The reactor’s input contains a poisoning agent with concentration x_p which deteriorates activity θ [31]. The system dynamics in space z and time t is modelled as

$$\frac{\partial x_{a,p}}{\partial t} = -U \frac{\partial x_{a,p}}{\partial z} - \alpha(T)r_{a,p}, \quad (7)$$

$$\frac{\partial T}{\partial t} = -\beta(x_a, T)U \frac{\partial T}{\partial z} + \gamma r_a, \quad \frac{d\theta}{dt} = -r_d, \quad (8)$$

where U is the fluid velocity, $\alpha(T)$ and $\beta(x_a, T)$ are functions of temperature T and the concentration x_a , γ is a constant and the functions r_a , r_p and r_d are reaction, poisoning adsorption, and catalyst deactivation rates (see [2] for more details). A simulation of the system is shown in Fig. 3(b).

4.2. Baselines

We compared the proposed approach against the (non-diffusion) supervised neural operators FNO [13] and OFormer [14], which have been trained using the implementations provided by the authors. As diffusion model baselines, we compared the mixed conditional training (M-CEDM) with the ‘unconditional’ RePaint approach and the conditional diffusion models based on the DDPM (CDDPM) and EDM (CEDM). RePaint and CDDPM are trained on the simplified objective [8] by predicting added noise, while CEDM and M-CEDM use the EDM objective [24] and predict a mixture of noise and denoised image. All diffusion models rely on

the 2nd order Heun sampler [24]. In all experiments we take 100 samples from each diffusion model and calculate the mean absolute error (MAE) and the PDE residual error by averaging across those samples. The number of parameters in the different models is approximately equal. The code and the datasets will be released upon acceptance.

4.3. Reconstruction of the unobserved variables

In this experiment (results in Table 1), we consider two tasks: forward and inverse. In the forward task we predict the solution u conditioned on the parameter a (Darcy), the velocity u from the height h (SWE-orig) and both the concentration x_a and the catalyst θ from the temperature T and poisoning concentration x_p (Reactor). For the inverse task the inputs and the targets are flipped (Tasks 1, 2 in Fig. 1). For these tasks the supervised neural operators, FNO and OFormer, provide strong results. However, the PDE residuals are consistently smaller for the conditional and mixed conditional diffusion models (CDDPM, CEDM and M-CEDM), which means that their solutions are more compatible with the prior knowledge of the system dynamics. Since CEDM results were better than CDDPM, we selected it as the basis for training the mixed conditional approach. Comparing RePaint trained ‘unconditionally’ to generate both states jointly against the specialized conditional models, we conclude that the conditioning during training is beneficial. However RePaint uses a single model for both forward and inverse problems, while the conditional models (including FNO and OFormer) use a dedicated model for each task. Similar to RePaint, M-CEDM uses a single model for both tasks but with performance similar to or better than the other conditionally trained models.

4.4. Future state prediction

Next, we compared the methods on Task 3 (future prediction) and Tasks 4-5 (reconstruction and prediction) in Fig. 1, results shown in Table 2. For RePaint and M-CEDM, we used

Table 2. Results for SWE-orig dataset in reconstruction and prediction (Tasks 3-5). MAE values should be multiplied by 10^{-3} . Models with * are the same as in Table 1 and trained for full state reconstruction.

Model	Task 3		Task 4		Task 5	
	MAE	PDE	MAE	PDE	MAE	PDE
FNO* [13]	0.70	0.98	51.2	47.5	90.7	150
OFormer* [14]	0.40	1.00	30.3	105	92.8	47.5
FNO [13]	0.70	0.98	0.56	0.98	1.34	1.04
OFormer [14]	0.40	1.00	0.35	1.00	1.04	1.01
RePaint [9]	0.65	1.08	0.79	1.04	54.2	4.72
M-CEDM	0.08	0.99	0.14	0.99	0.48	0.99

Table 3. Results of state reconstruction (Task 1) in the non-identifiable case (SWE-init). All MAE values should be multiplied by 10^{-3} . We see that samples from the diffusion model include the correct prediction (‘closest’), and additional prior information (extra observed points or PDE error) can be useful to identify the most plausible samples.

Model	MAE	PDE
Kalman Filter [32]	52.2	0.2858
FNO [13]	2.43	0.1300
OFormer [14]	2.74	0.0400
Diffusion model, mean prediction	55.1	0.0193
- prediction selected by PDE error	13.7	0.0159
- prediction selected by 2 corner points	1.92	0.0161
- prediction closest to the target	1.53	0.0162

the same trained model as in Table 1. We selected M-CEDM as a representative of the conditional diffusion models as it was the best in the previous experiment and did not require retraining. For FNO and OFormer we had two models, one to reconstruct the state, the other for prediction. Both FNO and OFormer can handle inputs with different amounts of observed time steps at test time; however, the performance drops significantly when the models trained on the full state are applied to partial observations (marked by * in Table 2). By retraining FNO and OFormer for the correct input size, the performance improved but stayed below the diffusion model.

4.5. Recovering several possible outcomes

Here we compare the models for velocity reconstruction (Task 1) in the SWE-init system, and the results are summarized in Table 3. The system is partially non-identifiability, which means that without additional information multiple reconstructions can be compatible with the conditioning information. The purpose of this experiment is to demonstrate the benefits of the probabilistic diffusion models in this scenario,




Fig. 4. State reconstruction results for non-identifiable system SWE-init for a test sample with multiple possible outcomes.

namely their ability to generate multiple different samples reflecting the variability of possible solutions, and including the correct target value. Here, we include also the Kalman Filter [32], obtained by linearizing the PDEs [33], as another probabilistic baseline.

As shown in Table 3, the diffusion model mean prediction, obtained by averaging over the samples, has a higher MAE than the other methods, as expected. To illustrate this further, Fig. 4 presents results for one test example with a large mean MAE. We see that the diffusion model produced samples which all have the same overall shape, but which differ in the initial velocity, which reflects true non-determinacy of the system. Further, the correct result is among the samples (‘closest’), and just by including two additional points in the conditioning set it can already be identified. The deterministic FNO and OFormer models have a relatively small MAE but they also introduce artifacts, making the prediction not compatible with the physical dynamics, as seen also in the large PDE-errors. The Kalman filter yields a plausible prediction but is less precise due to the linearization.

Finally, Fig. 5 highlights the potential of the PDE-error as a prior to identify the most plausible samples. For all test examples the PDE-error is clearly correlated with the MAE (left and middle panels in Fig. 5). However, even among the samples with the smallest PDE-errors, there can still be variability, as demonstrated on the middle row in Fig. 4. The rightmost panel in Fig. 5 confirms the earlier finding that the diffusion model produces samples with a smaller PDE-error, meaning that they better comply with the underlying physics.

5. DISCUSSION AND FUTURE WORK

We proposed a diffusion model that can be used as a neural operator to solve several tasks without retraining. Additionally, when the problem is not fully defined, the probabilistic




Fig. 5. Detailed results for the non-identifiable system SWE-init. **Left:** PDE error vs. MAE for 100 samples for a single test case. **Middle:** Histogram of correlations between MAE and PDE error across the test set (red line shows the case on the left). **Right:** PDE errors of different models in the test set.

model, unlike the deterministic alternatives, can produce several possible results compatible with the conditioning information, and we studied using prior knowledge, either additional measured points or the PDE error, as a way of selecting the most plausible samples. As a limitation, the inference with the diffusion model requires sampling, which, despite recent improvements, is still slower than with the deterministic baselines. As future work, the model architecture could be tailored for the specifics of the dynamical systems.

6. REFERENCES

- [1] S. Greydanus et al., “Hamiltonian neural networks,” in *NeurIPS*, 2019, vol. 32.
- [2] K. Haitsiukevich et al., “A grid-structured model of tubular reactors,” in *2021 IEEE 19th Int. Conf. on Ind. Informatics (INDIN)*, 2021, pp. 1–6.
- [3] M. Raissi et al., “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *J. Comput. Physics*, vol. 378, pp. 686–707, 2019.
- [4] T. Yang et al., “Learning physics constrained dynamics using autoencoders,” *NeurIPS*, vol. 35, pp. 17157–17172, 2022.
- [5] R. Lam et al., “Learning skillful medium-range global weather forecasting,” *Science*, p. eadi2336, 2023.
- [6] Z. Li et al., “Physics-informed neural operator for learning partial differential equations,” *ACM/JMS J. Data Science*, 2021.
- [7] J. Sohl-Dickstein et al., “Deep unsupervised learning using nonequilibrium thermodynamics,” in *ICML*, 2015, pp. 2256–2265.
- [8] J. Ho et al., “Denoising diffusion probabilistic models,” *NeurIPS*, vol. 33, pp. 6840–6851, 2020.
- [9] A. Lugmayr et al., “Repaint: Inpainting using denoising diffusion probabilistic models,” in *Proc. IEEE/CVF Conf. on Comp. Vision and Pattern Recognition*, 2022, pp. 11461–11471.
- [10] C. Saharia et al., “Palette: Image-to-image diffusion models,” in *ACM SIGGRAPH 2022 Conf. Proc.*, 2022, pp. 1–10.
- [11] S. Yang et al., “Uni-paint: A unified framework for multimodal image inpainting with pretrained diffusion model,” in *Proc. 31st ACM Int. Conf. on Multimedia*, 2023, pp. 3190–3199.
- [12] J. Lopez Alcaraz et al., “Diffusion-based time series imputation and forecasting with structured state space models,” *Trans. Mach. Learn. Research*, 2023.
- [13] Z. Li et al., “Fourier neural operator for parametric partial differential equations,” in *ICLR*, 2021.
- [14] Z. Li et al., “Transformer for partial differential equations’ operator learning,” *Trans. Mach. Learn. Research*, 2023.
- [15] N. Kovachki et al., “Neural operator: Learning maps between function spaces with applications to pdes,” *J. Mach. Learn. Research*, vol. 24, no. 89, pp. 1–97, 2023.
- [16] D. Long et al., “Invertible fourier neural operators for tackling both forward and inverse problems,” *arXiv preprint arXiv:2402.11722*, 2024.
- [17] G. Kohl et al., “Turbulent flow simulation using autoregressive conditional diffusion models,” *arXiv preprint arXiv:2309.01745*, 2023.
- [18] G. Yang et al., “A denoising diffusion model for fluid field prediction,” *arXiv preprint arXiv:2301.11661*, 2023.
- [19] R. Apte et al., “Diffusion model based data generation for partial differential equations,” *arXiv preprint arXiv:2306.11075*, 2023.
- [20] J. Bastek et al., “Physics-informed diffusion models,” *arXiv preprint arXiv:2403.14404*, 2024.
- [21] C. Jacobsen et al., “Cocogen: Physically-consistent and conditioned score-based generative models for forward and inverse problems,” *arXiv preprint arXiv:2312.10527*, 2023.
- [22] D. Shu et al., “A physics-informed diffusion model for high-fidelity flow field reconstruction,” *J. Comput. Physics*, vol. 478, pp. 111972, 2023.
- [23] A.Q. Nichol et al., “Improved denoising diffusion probabilistic models,” in *ICML*, 2021, pp. 8162–8171.
- [24] T. Karras et al., “Elucidating the design space of diffusion-based generative models,” *NeurIPS*, vol. 35, pp. 26565–26577, 2022.
- [25] O. Ronneberger et al., “U-net: Convolutional networks for biomedical image segmentation,” in *Med. Image Comput. and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241.
- [26] Y. Song et al., “Score-based generative modeling through stochastic differential equations,” *ICLR*, 2021.
- [27] P. Dhariwal et al., “Diffusion models beat gans on image synthesis,” *NeurIPS*, vol. 34, pp. 8780–8794, 2021.
- [28] J. Song et al., “Denoising diffusion implicit models,” *ICLR*, 2021.
- [29] M. Takamoto et al., “Pdebench: An extensive benchmark for scientific machine learning,” *NeurIPS*, vol. 35, pp. 1596–1611, 2022.
- [30] D.I. Ketcheson et al., “Pyclaw: Accessible, extensible, scalable tools for wave propagation problems,” *SIAM J. Scientific Comput.*, vol. 34, no. 4, pp. C210–C231, 2012.
- [31] T.H. Price et al., “Catalyst poisoning and fixed bed reactor dynamics—II: Adiabatic reactors,” *Chem. Eng. Sci.*, vol. 32, no. 4, pp. 393–412, 1977.

- [32] R.E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [33] M. Rafiee et al., "Kalman filter based estimation of flow states in open channels using lagrangian sensing," in *Proc. 48th IEEE Conf. on Decision and Control (CDC)*, 2009, pp. 8266–8271.