Physics-Informed Diffusion Model for Super-Resolution and Surrogate Modelling of Time-Dependent Partial Differential Equations

M.R. Hasan[1,2], **P. Behnoudfar**[2,3], **F. Nugen**[4,5], **T. Poulet**[2] and **T. Gedeon**[1,6]

[1]Curtin University, [2]Commonwealth Scientific and Industrial Research Organisation (CSIRO), [3]University of Wisconsin-Madison, [4]Seirrowon Labs, [5]University of California, Berkeley, [6]Australian National University

Summary

We present a Physics-Informed Denoising Diffusion Probabilistic Model (PIDDPM) for super-resolution and surrogate modelling of time-dependent physical systems. PIDDPM is conditioned on a coarse-resolution input at the current timestep and high-resolution ground truth from the two preceding timesteps, with the aim of reconstructing fine-scale solutions consistent with the underlying dynamics. PIDDPM acts as a surrogate, approximating the behaviour of nonlinear PDEs such as the Allen-Cahn equation without requiring full numerical simulation. Physics-based penalties are incorporated into the loss function to penalise lack of consistency with the governing equations and boundary conditions, ensuring that the generated outputs remain physically plausible. Our results demonstrate that PIDDPM significantly improves perceptual and physical accuracy compared to baseline DDPM, achieving higher PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index Measure) while reducing MSE (Mean Squared Error) and MSGE (Mean Squared Gradient Error). The model's ability to learn temporal evolution and spatial refinement makes it a scalable and physically grounded alternative to traditional solvers. PIDDPM shows strong potential for resolution enhancement, interpolation and predictive modelling in subsurface workflows, offering a data-driven approach to accelerate simulations and support efficient decision-making in geoscientific domains.

Corresponding author: rakibul.hasan@curtin.edu.au

This is the author's version before final copy-editing. The final published version is available at https://doi.org/10.3997/2214-4609.202576020.

Publisher: European Association of Geoscientists & Engineers
Source: EAGE Workshop on Enhancing Subsurface Practices using AI/ML, Nov 2025, Volume 2025, p. 1 – 5.

Introduction

Recent advances in generative modelling have opened new possibilities for enhancing resolution and accelerating simulations in scientific domains (Hasan et al., 2025a). In subsurface modelling and other time-dependent physical systems, obtaining high-resolution data across space and time is often computationally expensive or practically infeasible. To address this, we propose a physics-informed neural network model that performs both super-resolution and surrogate modelling. The model is guided by physical constraints and temporal context to generate high-resolution predictions from coarse inputs.

We consider a general transient Initial-Boundary-Value Problem (IBVP) defined over a space-time domain, where the solution evolves over time under specified initial and boundary conditions. The governing equation involves a first-order time derivative and a spatial operator that encapsulates physical properties, nonlinearities and source terms. Such a formulation is broad enough to model a wide range of time-dependent Partial Differential Equations (PDEs) relevant to subsurface and fluid dynamics applications. To demonstrate the effectiveness of our method, we focus on a specific class of IBVP, namely the Allen-Cahn reaction-diffusion equation (Allen and Cahn, 1972).

Diffusion models are a class of generative machine learning models used for a variety of tasks, such as denoising (Kulikov et al., 2023), inpainting (Corneanu et al., 2024) and image super-resolution (Yue et al., 2023). They have been used successfully for generating or improving a variety of image types, including still life images (e.g., photographs, paintings, and cartoons) (Ko et al., 2023) and medical images (Yoon et al., 2023). Given the utility of diffusion models for solving such a variety of computer vision tasks, we hypothesised that diffusion models would be valuable for solving the reaction-diffusion equation.

Method

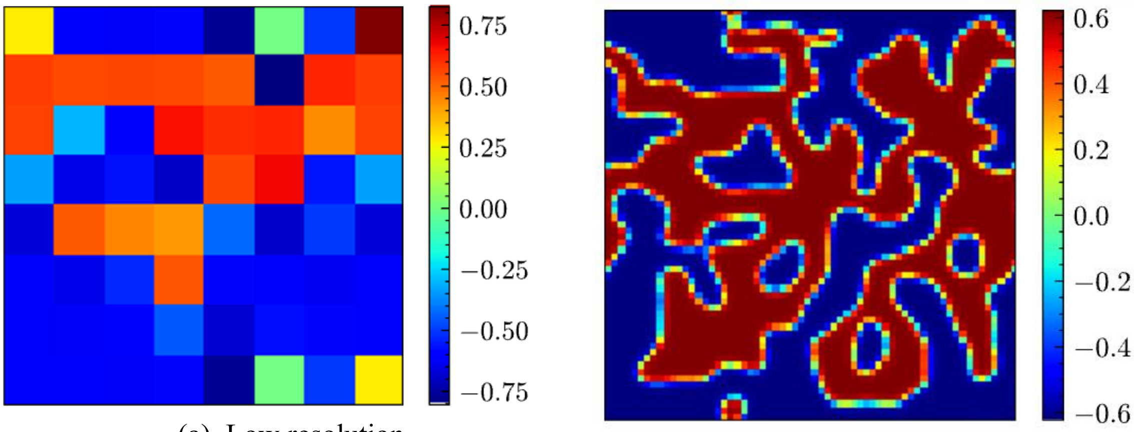

(a) Low resolution

(b) High resolution

Figure 1 *Example (a) low-resolution and (b) corresponding high-resolution image pairs at a particular time step. The low-resolution one is provided as input to the model, while the high-resolution (ground-truth) one is the desired output from the model.*

Dataset. To train the machine learning model, we use a publicly available dataset, which was created using the Allen-Cahn equation with the Periodic boundary condition (Hasan et al., 2025b). This dataset was generated through discretising the spatial domain using the Finite Element Method (FEM) and applying a time-marching scheme to obtain fully discrete solutions. These solutions were used to construct time series of low- and high-resolution images, where each image represents a snapshot of fluid flow at a specific timestep. The resolution was upscaled by a factor of 8 in each spatial dimension (input resolution 8×8 to output resolution 64×64), resulting in a rich dataset that captures the temporal evolution of the system conforming to IBVP. Figure 1 depicts an example low- and high-resolution image pair.

Physics-Informed Denoising Diffusion Probabilistic Model (PIDDPM). We consider the DDPM proposed in (Khosravi et al., 2023) as a baseline model. The training process includes adding Gaussian noise as a forward diffusion process first, and then a U-Net deep learning model learns to denoise the noisy image. While Khosravi et al., (2023) proposed this DDPM for a biomedical image segmentation task, we extend it to super-resolution and surrogate modelling.

To model super-resolution, we condition DDPM using the low-resolution image, interpolated to the high-resolution image size, at the current timestep n. In PIDDPM, we concatenate the interpolated high-resolution image at n with the high-resolution ground truth images at $(n-1)$ and $(n-2)$ timesteps. With these inputs, PIDDPM learns to reconstruct the fine-scale solution at timestep n, effectively approximating the behaviour of the underlying physical system. This setup enables the model to act as a data-driven surrogate, capturing temporal dynamics and spatial structures without solving the full PDE at each step.

PIDDPM includes two additional physics-based loss components in addition to DDPM: one for the inner and another for the boundary regions of the images. These terms are derived from a semi-discrete version of the governing equation and are designed to guide the model toward physically plausible solutions. The physics-based loss components for the inner region of the images, based on the Backward Differentiation Formula (BDF) time integrator, also take the same high-resolution ground truth images at $(n-1)$ and $(n-2)$ timesteps. For details of the physics-based loss components, including their mathematical structure and implementation, we refer interested readers to our recent work, PC-SRGAN (Hasan et al., 2025a). PC-SRGAN introduced a physically consistent super-resolution framework based on Generative Adversarial Network (GAN). In this study, we extend that approach by adopting a DDPM as the backbone.

We utilise open-source software associated with (Khosravi et al., 2023) to streamline the development and training of DDPM and PIDDPM. Training was conducted for 1,000 diffusion timesteps until 100K images (with repetitions) were seen by the model.

Evaluation Metrics. To assess the performance of our super-resolution framework, we employ four metrics: Peak Signal-to-Noise Ratio (PSNR), which quantifies reconstruction fidelity; Structural Similarity Index Measure (SSIM), which evaluates perceptual similarity; Mean Squared Error (MSE), which captures pixel-wise differences; and Mean Squared Gradient Error (MSGE), proposed in (Hasan et al., 2025a), which measures consistency in spatial gradients to reflect physical consistency for scientific application. For PSNR and SSIM, higher is better, and for MSE and MSGE, lower is better.

Results

Figure 2 compares the validation performance of PIDDPM against the baseline DDPM across training steps. PIDDPM consistently outperforms the baseline, showing higher PSNR and SSIM values, indicating better image fidelity and perceptual quality. It also achieves lower MSE and MSGE, demonstrating improved pixel accuracy and stronger adherence to physical consistency.

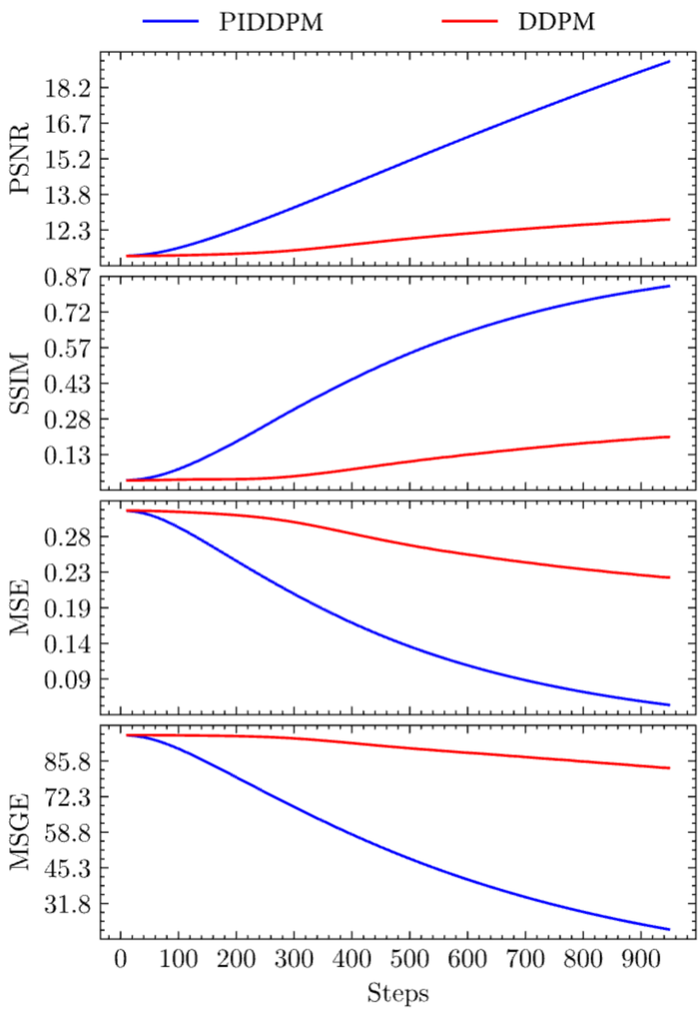

Figure 2 *Comparison between our Physics-Informed Denoising Diffusion Probabilistic Model (PIDDPM) versus baseline DDPM in terms of validation metrics during training for the image super-resolution task. Abbreviations: PSNR: Peak Signal-to-Noise Ratio; SSIM – Structural Similarity Index Measure; MSE – Mean Squared Error; MSGE – Mean Squared Gradient Error.*

Our PIDDPM framework offers a promising direction towards AI-driven subsurface modelling practices. Since the generative process in PIDDPM is conditioned on interpolated low-resolution inputs and temporal context from previous high-resolution states, it not only performs super-resolution but also acts as a data-driven surrogate model capable of predicting future states. This dual capability is particularly valuable in subsurface applications where high-fidelity simulations are computationally expensive and temporally rich data is scarce. The integration of physical constraints ensures that the generated outputs remain consistent with governing equations, making the approach suitable for tasks such as flow prediction and geophysical interpolation. As such, PIDDPM contributes toward scalable, physics-aware ML solutions that can accelerate decision-making and reduce reliance on costly numerical solvers in subsurface workflows.

Conclusions

We present PIDDPM, a physics-informed generative framework that performs super-resolution and surrogate modelling for time-dependent physical systems. By conditioning on coarse, low-resolution input and high-resolution temporal contexts, our model accurately reconstructs the fine-scale solution at the current timestep. PIDDPM approximates temporal derivatives and spatial structures without solving the full PDE, offering a data-driven surrogate for accelerating simulations. The integration of physics-based loss components drives consistency with governing equations, making the model particularly suitable for subsurface applications where resolution enhancement and predictive modelling are essential.

References

Allen, S. M., and Cahn, J. W. [1972]. Ground state structures in ordered binary alloys with second neighbor interactions. *Acta Metallurgica*, *20*(3), 423–433. https://doi.org/10.1016/0001-6160(72)90037-5

Corneanu, C., Gadde, R., and Martinez, A. M. [2024]. LatentPaint: Image Inpainting in Latent Space With Diffusion Models. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 4334–4343.

Hasan, M. R., Behnoudfar, P., MacKinlay, D., and Poulet, T. [2025a]. PC-SRGAN: Physically Consistent Super-Resolution Generative Adversarial Network for General Transient Simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–8. https://doi.org/10.1109/TPAMI.2025.3596647

Hasan, M. R., Behnoudfar, P., Mackinlay, D., and Poulet, T. [2025b]. *Transient Simulation Datasets Using Allen-Cahn and Eriksson-Johnson Equations* [Data set]. https://doi.org/10.5281/ZENODO.16250888

Khosravi, B., Rouzrokh, P., Mickley, J. P., Faghani, S., Mulford, K., Yang, L., … Wyles, C. C. [2023]. Few-shot biomedical image segmentation using diffusion models: Beyond image generation. *Computer Methods and Programs in Biomedicine*, *242*, 107832. https://doi.org/10.1016/j.cmpb.2023.107832

Ko, H.-K., Park, G., Jeon, H., Jo, J., Kim, J., and Seo, J. [2023]. Large-scale Text-to-Image Generation Models for Visual Artists' Creative Works. *Proceedings of the 28th International Conference on Intelligent User Interfaces*, 919–933. https://doi.org/10.1145/3581641.3584078

Kulikov, V., Yadin, S., Kleiner, M., and Michaeli, T. [2023]. SinDDM: A Single Image Denoising Diffusion Model. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett (Eds), *Proceedings of the 40th International Conference on Machine Learning* (pp. 17920–17930). PMLR.

Yoon, J. S., Zhang, C., Suk, H.-I., Guo, J., and Li, X. [2023]. SADM: Sequence-Aware Diffusion Model for Longitudinal Medical Image Generation. In A. Frangi, M. de Bruijne, D. Wassermann, and N. Navab (Eds), *Information Processing in Medical Imaging* (pp. 388–400). Cham: Springer Nature Switzerland.

Yue, Z., Wang, J., and Loy, C. C. [2023]. ResShift: Efficient Diffusion Model for Image Super-resolution by Residual Shifting. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds), *Advances in Neural Information Processing Systems* (Vol. 36, pp. 13294–13307). Curran Associates, Inc.