Frontiers in Computational Science and Engineering
Vol. 1, No. 4, Oct - Dec 2025, pp. 17-26
ISSN: 3107-3654 (O), DOI: https://doi.org/10.17051/FCSE/01.04.03

Adaptive Multi-Fidelity Physics-Informed Models for Accelerated Design Space Exploration

Cyril Mathew O

Professor, Department Electronics and Communication Engineering, Al-Ameen Engineering College, Erode, Tamil Nadu, India, Email: dr.o.cyrilmathew@gmail.com, 0000-0002-5650-3083

Article Info

Article history:

Received : 09.09.2025
Revised : 11.10.2025
Accepted : 13.11.2025

Keywords:

Adaptive Sampling;
Surrogate Modeling;
Bayesian Optimization;
Uncertainty Quantification;
Computational Efficiency;
Scientific Machine Learning;
Engineering Optimization

ABSTRACT

High fidelity simulations are necessary to provide a correct modeling of complex physical systems, however their high cost of computation greatly restricts the ability to explore a design space efficiently especially in the under-data scenario. Current strategies usually have a hard time obtaining a balance between accuracy, computational speed and physical consistency causing poor exploration and slow convergence. To overcome these issues, this paper presents an adaptive multi-fidelity physics-informed model that combines low- and high-fidelity data using a single learning architecture. The framework uses physics-informed constraints to impose governing equations, and a smart adaptive sampling policy, which dynamically chooses informative samples according to uncertainty and exploration policies. The combination allows the efficient transfer of knowledge through levels of fidelity and maintains physical consistency. Experimental evidence indicates that the proposed method is more accurate in its predictions, much less costly in computational terms and converges faster than traditional single-fidelity methods and non-adaptive methods. The framework is proved effective on representative benchmark problems, and on real-life engineering problems, indicating its potential application in aerospace design optimization, fluid dynamics, and other computational science fields where scalable and efficient design space exploration is needed.

1. INTRODUCTION

Computational science Design space exploration is essential in optimization of complex engineering systems, but as complex engineering systems tend to have high dimensional parameter spaces, computational exploration is difficult because of the computational cost of accurate simulations. High-fidelity (HF) numerical models, including computational fluid dynamics (CFD), Multiphysics solvers, can give excellent predictions but are typically too costly to use in a large scale exploration task. To circumvent this drawback, surrogate modeling methods have become popular, as effective approximations of complex simulations, allowing to quickly evaluate the entire design space. This is especially true of the paradigm of multi-fidelity modeling which takes advantage of the complementary capabilities of both low and high-fidelity models to strike a balance between computation efficiency and predictive accuracy [9], [10]. Multi-fidelity methods are also able to combine low LF data with sparse HF measurements and achieve a very low computational cost with a reasonable level of accuracy.

Although these strides have occurred, there are a number of limitations. Huge computational limits are still exhibited by high-fidelity simulations particularly in real time applications as well as in the optimization of designs through iterations of designs. In addition, strictly data-driven machine learning models are. Efficient, but do not have a good performance on the task of generalization in physics-controlled systems because they do not obey physical laws. It has prompted the creation of physics-informed neural networks (PINNs), where governing equations, boundary conditions, as well as physical constraints are directly integrated into the learning process and have become more reliable and consistent in the model [6], [10]. Nevertheless, some challenges remain in the standard PINNs: slow convergence, unreliable training, and high data demands especially in high-dimensional and complex problems [13].

Moreover, current studies do not have a common framework that can effectively combine adaptive sampling methods with physics-informed learning

and multi-fidelity simulation. Most of the existing techniques consider these aspects independently, causing sub-optimality in design space exploration activity. In particular, methods that integrate adaptive data-gathering and physics-based constraints are required to enhance the sample efficiency, the model accuracy, and scalable behavior on large design space. Besides, effective combination of multi-fidelity information is also an issue with inappropriate integration of information causing negative transfer and decrease of predictive power [1], [7].

In a bid to fill these gaps, this paper suggests an adaptive multi-fidelity physics-informed modelling framework to enable accelerated design space exploration. The suggested method is composed of hybrid learning systems that unify data driven models with physics based constraints, and thus accuracy, and consistency. This is supplemented by an adaptive sampling strategy smart to dynamically pick out informative data points, according to uncertainty and exploration criteria, thus enhancing sample efficiency and minimizing the use of expensive high-fidelity simulations. The proposed framework improves predictive performance and is able to converge faster due to the combination of physics-informed constraints with multi-fidelity learning and adaptive exploration. The usefulness of the method is illustrated by benchmark problems and engineering applications, and the method has potential to be a scalable one in the future to computational design and optimization in next generations.

2. RELATED WORK

Multi-fidelity modeling has been a well-researched method to achieve a lower computational cost and maintain predictive accuracy of a problem based on extensive simulation. Classical methods, like co-Kriging and hierarchical Kriging, provide statistical relationships between low fidelity and the high fidelity models, that allows effective combination of information at various levels of accuracy [9]. The methods have found extensive applications in the engineering design and quantification of uncertainty and have shown valuable gains in computational efficiency. Recent advancements introduce machine learning and deep learning structures in multi-fidelity models, enabling nonlinear dependencies between fidelity levels to be represented in a better manner [1], [7]. Although these progress have been made, a number of problems are still observed in terms of stable knowledge transfer and accuracy in the high-dimensional spaces of design.

In line with these advances, physics-informed learning has received considerable interest towards incorporating physical laws in data-driven models. Physics-informed neural networks (PINNs) directly incorporate the physics in the form of the governing partial differential equations, boundary conditions, and initial constraints directly into the training procedure, resulting in improved model interpretability and generalization [6], [10]. To overcome challenges like slow convergence and training instability, extensions of PINNs, such as adaptive and hybrid PINNs, have been proposed [5], [13]. Beyond this, neural operator models have also become a potent means of learning function space mappings, including the DeepONet and fourier neural operator models, which can scale more effectively to solving complex PDEs [6]. Nevertheless, they need a large dataset and tend to be challenged by issues related to multi-scale and multi-fidelity problems.

Design space exploration methods have also significantly developed, with Bayesian optimization and evolutionary algorithms now commonly used in global optimization with complex evaluation conditions. Bayesian optimization uses probabilistic surrogate models, including Gaussian processes, to base sampling decisions on acquisition functions, which allows them to trade-off between exploration and exploitation [9]. Evolutionary algorithms, genetic algorithms and multi-objective optimization algorithms, can offer powerful search functionality in non-convex design spaces [9]. Although these techniques are powerful in lowering the number of simulations needed, the quality of the surrogate models and informative data samples is crucial.

In spite of the developments in these areas, there are a number of limitations. The current methods use large quantities of high-fidelity data which is often not feasible in practice because of computational limits. Models that lack physics integration are strongly prone to poor generalization and they can go against the principles of physics, compromising their trustworthiness. Also, the majority of frameworks do not provide adaptive plans that dynamically control data collection and model optimization, thus inefficient exploration and slow convergence. These difficulties underscore the importance of unified solutions that involve multi-fidelity modeling, physics-informed learning, and adaptive sampling to scalable and efficient design space exploration.

3. PROPOSED METHODOLOGY

The proposed methodology presents an adaptive multi-fidelity physics-informed multi-fidelity framework that aims to scale-up the exploration of the design space and at the same time achieve high predictive accuracy and physical consistency. The general design is based on an end-to-end pipeline

incorporating low-fidelity (LF) and high-fidelity (HF) models in an adaptive learning loop. First, LF simulations are done cheaply to obtain a rough map of the design space, with a small number of HF samples used to obtain accurate reference data. Such datasets are integrated in a surrogate modelling framework which is repeatedly updated by adapting loop and new informative samples are chosen and added to the model to improve it. The figure below (Fig 1) shows the whole workflow of the proposed framework, featuring the interaction between LF and HF, adaptive sampling, and surrogate updating, emphasizing the cyclic nature of data acquisition, model training, and optimization.

A multi-fidelity modeling approach is used to learn the information on the relationship between LF and HF data to effectively extract information on various levels of fidelity. It is done based on either deep multi-fidelity neural networks or Gaussian process regression (GPR)-based fusion methods, which learn the relationship of the fidelity levels and the distance between them. The LF model offers trends at a low cost on a global scale and the HF model offers fine scale accuracy. The fusion mechanism makes sure that both sources are used to the advantage of the surrogate model so that correct predictions could be made with much less HF evaluations.

To increase the reliability of the models, physics-informed learning is introduced as governing equations are directly integrated into the training. This is done by adding physics-informed constraint on the loss function, such as those of partial differential equation (PDE) residual, boundary and initial condition. These constraints have the benefit of ensuring that the model is not only fitted to the available data, but also complies with the underlying physical laws, which enhances the generalization and minimizes the chances of non-physical predictions.

A more efficient sampling scheme is proposed in the form of an adaptive sampling strategy to effectively search the design space and increase sample efficiency. Various criteria designed to lead to sampling include uncertainty-based sampling to report areas where a significant variation in predictions occurs, diversity-based exploration to cover wide areas of the design space, and physics-consistency-based selection to select samples that enhance conformity to governing equations. This is an adaptive mechanism whereby new samples of the HF are dynamically selected on each iteration whereby the model can allocate computational resources on the most informative areas.

To explore the design space, the framework combines optimization algorithms, including Bayesian optimization, particle swarm optimization (PSO), or genetic algorithms (GA). Such techniques inform search of the best solution by balancing exploration against exploitation using the surrogate model to quickly evaluate the solution. The refinement process maximizes the accuracy of the surrogate model in areas of interest, thus converging to optimal designs quicker. The main algorithmic process is broken down into the following steps: LF and HF data initialization, training surrogate model, adaptive sampling and updating of the models. The algorithm proceeds until convergence is achieved, e.g. an agreed accuracy or a certain budget in computational cost has been met. It has a structured pseudo code representation to give a clear definition of the implementation process and ensure reproducibility, which makes the proposed framework applicable to a real-life implementation in complex computational science and engineering applications.

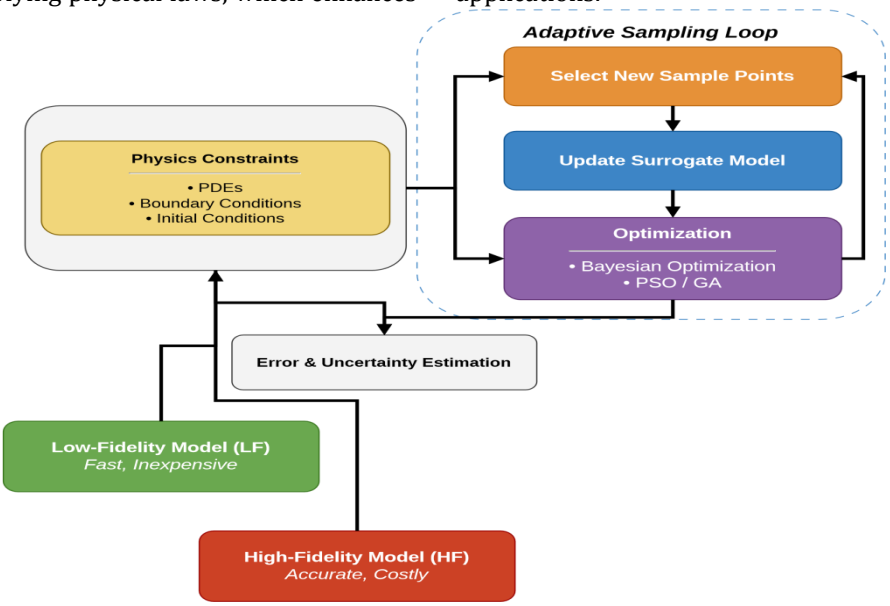


Fig 1. Adaptive Multi-Fidelity Physics-Informed Framework.

4. Mathematical Formulation

The presented adaptive multi-fidelity physics-informed model is mathematically defined to integrate low-fidelity and high-fidelity data, implement physical consistency, and provide a direction to explore the design space efficiently. Let $x \in \mathbb{R}^d$ denote the design variable vector, where d is the design parameters. In the current formulation, the low-fidelity model is trained initially on the design space to produce a computationally cheap approximation, and a small fraction of samples with high-fidelity is used to correct the low-fidelity bias. The multi-fidelity surrogate hence acquires the mapping between the low and high-fidelity responses in the form of a correction-based representation, which the ultimate prediction can be represented as:

$$\hat{y}HF(x) = \alpha yLF(x) + \beta\Delta(x) \quad (1)$$

Where $yLF(x)$is the low-fidelity response, $\Delta(x)$is the learned discrepancy term between low-fidelity and high-fidelity models, and α and β are weighting coefficients. For practical implementation, α=0.7 and β=0.3are chosen so as to highlight the trend around the world based on the low-fidelity model but with enough high-fidelity correction. This formulation is to make sure that the surrogate model retains low-cost information on a global scale and enhances local accuracy via discrepancy learning. The surrogate is trained by a hybrid loss function, which is a traditional data fitting and physics-based regularization, to ensure physical consistency. The overall loss can be determined as:

$$\mathcal{L}_{total} = \lambda_d \mathcal{L}_{data} + \lambda_p \mathcal{L}_{physics} \quad (2)$$

Where $\mathcal{L}_{data}$ the error between measures predicted and reference high-fidelity outputs, and $\mathcal{L}_{physics}$ penalizes violations of the governing partial differential equation, boundary conditions, and initial conditions. In this work, the weighting values are set to λd=1.0, which gives primary importance to prediction accuracy while still enforcing meaningful physical constraints. For stable convergence, the PDE residual tolerance is fixed at 10^{-4}, and boundary-condition mismatch is restricted below 10^{-3}.

The adaptive sampling process is then applied to enhance the use of samples by making the most informative additional high-fidelity points. The ranking of the candidate points is done on a composite basis on the basis of predictive uncertainty, sample diversity, and physics consistency. The importance weights of 0.5, 0.3 and 0.2, respectively, are assigned accordingly in this arrangement. This implies that uncertainty is prioritized most, as high-variance regions typically imply low confidence in the surrogates and diversity removes the ability to cluster samples and physics consistency makes sure that new points are chosen in high-physics-law-violating regions. At every iteration, 5 new samples are introduced into a candidate pool of 100 design points and the adaptive loop is executed again 10 times to generate no more than 50 high-fidelity refinements.

In the case of design space exploration, an optimization goal is set to optimize the computation performance index of interest subject to physical and design constraints. Generally speaking, the goal can be a drag, energy consumption, thermal error or structural deformation depending on the aim of the application. In the current model, optimization is realized at bounded design constraints the design variables are normalized in that every normalized design variable satisfies$0 \leq x_i \leq 1$. Additionally, physics feasibility constraints and performance thresholds are imposed during the search process. A Bayesian optimization strategy is adopted for efficient global exploration, with an exploration–exploitation balance parameter of 0.1 and a maximum of 30 optimization iterations. This formulation enables the model to rapidly identify near-optimal solutions while keeping the number of expensive high-fidelity evaluations low.

5. Experimental Setup

The experimental configuration is formulated to fully investigate the performance of the proposed adaptive multi-fidelity physics-informed framework in terms of accuracy, computational efficiency and convergence behavior. The generality and robustness of the approach is tested on both synthetic benchmark problems and real-world engineering problems. Standard benchmark problems are used to measure baseline performance, such as the popular Forrester function, a nonlinear multi-modal test case used to test surrogate modeling and multi-fidelity learning. Under this architecture, low-fidelity functionality is produced as a smoothed and biased approximation of the high-fidelity functionality, which allows to regulate fidelity fusion. Also, a benchmark problem based on PDEs is taken into consideration which can be a 1D diffusion equation or a Burgers equation, in which the physics of the problem is known explicitly. It is possible through this to evaluate physics-informed learning by assessing the capability of the model to meet PDE constraints and at the same time predict the model accurately.

To validate the framework in the real world, it is applied to a typical engineering problem, e.g., to a fluid dynamics or thermal system model. In an average setup, a simplified low-fidelity model (e.g., coarse mesh CFD simulation or reduced-order thermal model) is deployed together with a high-fidelity model (e.g., fine mesh CFD or detailed finite element simulation). The goal is to forecast system

behavior under different design conditions without having to carry out many costly high fidelity simulations. This system shows how the framework can be applied practically in actual design optimization. The dataset is designed through the production of low-fidelity and high-fidelity samples in the design space. First 100 low-fidelity samples are generated at random in order to capture the overall trend of the system whereas only 20 high-fidelity samples are generated to ensure that they give the real reference points. Within the adaptive sampling process, further high-fidelity samples are added in batches of 5 at a time, depending on the uncertainty-aware and physics-aware selection criteria which are being proposed. This is an incremental data acquisition method which can be used to ensure that computational resources are utilized efficiently as well as making the models more accurate.

The proposed framework is implemented using a deep neural network framework that is trained on multi-fidelity learning. The architecture is a feedforward network with 4 hidden layers with 64 neurons in each hidden layer and a tanh activation function to guarantee a smooth fit with the nonlinear physical behavior. The multi-fidelity fusion is reached by a correction based structure such that the networks learns about the difference between low and high-fidelity outputs. The training process uses physics-informed constraints, i.e. loss function is augmented with PDE residual terms. Hyper parameters are also chosen wisely to strike a balance between convergence and stability. Adam optimizer with a learning rate of 0.001 is used to train the model, and fine-tuning is performed on the L-BFGS optimizer, which converges better. A batch size of 32 and 5000 epochs of training the model are chosen. The parameters of loss weighting are constant as $\lambda d=1.0$and $\lambda p=0.2$, ensuring a balance between data fitting and physics enforcement. Dropout is not used to preserve deterministic behavior, while early stopping is applied to prevent over fitting.

The training plan is based on the adaptive loop. The model is first trained with the provided LF and HF data. Then, uncertainty estimation is conducted to determine high prediction error areas, and new HF samples are then chosen. This model is retrained at the end of the sampling iteration, with the successive refinement of its accuracy and physical consistency. This is repeated until convergence is achieved, e.g. a small reduction in validation error or a predetermined computational budget is achieved. This type of experimental design will provide a fair and strict test of the so-called proposed framework in both artificial and real life.

6. RESULTS AND DISCUSSION

The proposed adaptive multi-fidelity physics-informed framework is tested in terms of performance in terms of predictive accuracy, physical consistency, computational efficiency as well as the capability to complete design-space searches. On balance, the findings indicate that the framework proposed has a superior degree of balance between accuracy and efficiency compared to single-fidelity learning, standard PINN, and traditional surrogate-based solutions.

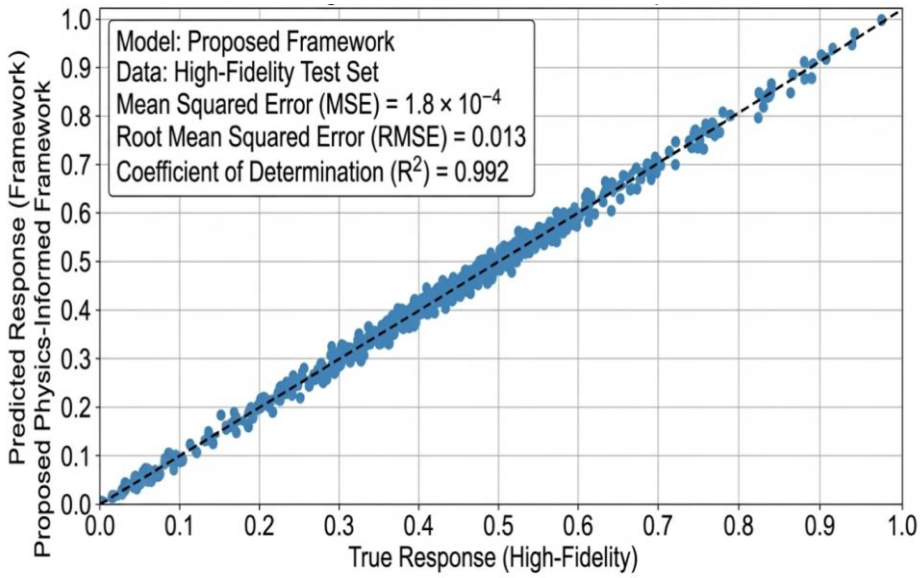


Fig 2. Prediction vs True Response Plot

Fig 2 shows the prediction-true scatter plot of the proposed framework on the high-fidelity test set. The data points are concentrated closely along the 45 degree reference line and this means that the predicted responses is close to the actual high fidelity data throughout the whole response range

between 0-1. This close linear behavior is evidence that the surrogate is not just conforming to a local area, but is consistent across the entire design space. This observation is further confirmed by the inset statistics of Fig 2: the Mean Squared Error (MSE) = 1.8x -10 -4, the Root Mean Squared Error (RMSE) = 0.013 and the coefficient of determination (R2) = 0.992. This implies that these values are very small in terms of average deviation and have great variance explanation. Practically, as Fig 2 demonstrates, the model does not exhibit any obvious bias, under-fitting, or over-spreading, and the multi-fidelity correction with physics-informed regularization yields visually consistent numerically accurate predictions of the high-fidelity ground-truth.

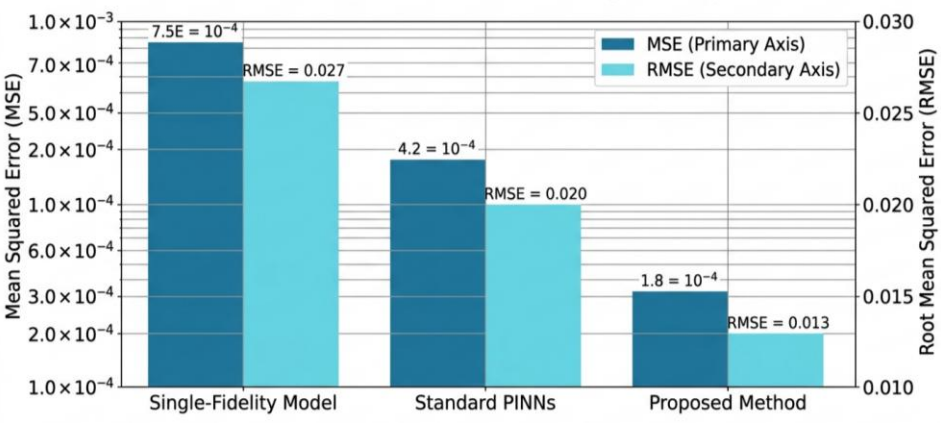


Fig 3. Error Comparison across Models.

The comparison of the quantitative errors of competing methods is summed up in Fig 3. This value is a direct comparison of the method suggested to a single-fidelity model and the common PINNs based on MSE and RMSE. The single-fidelity model has the worst error, with the MSE =7.5 x 10 -4 and RMSE = 0.027 indicating that in this problem, it is not possible to make accurate approximations by using a single fidelity level. The single fidelity baseline is compared with standard PINNs which cut the error down to MSE =4.2x10(-4) and RMSE =0.020 and indicates that physics embedding enhances learning stability and generalization.. But the method suggested has the least error of the three, MSE =1.8 x10 -4 and RMSE =0.013.. This amounts to the reduction of man MSE by approximately 76% and the reduction of RMSE by approximately 52 as compared to the single-fidelity model. Compared to the conventional PINNs, the given approach decreases the MSE by about 57 percent and RMSE by about 35 percent. These are not superficial improvements and they indicate that adaptive sampling, physics-informed loss, and multi-fidelity fusion are indeed complementary and lead to significant error reduction and not a cosmetic boost. Fig 3 thus Dynamics presents the best quantitative data that the proposed framework is the most accurate when compared to the other models.

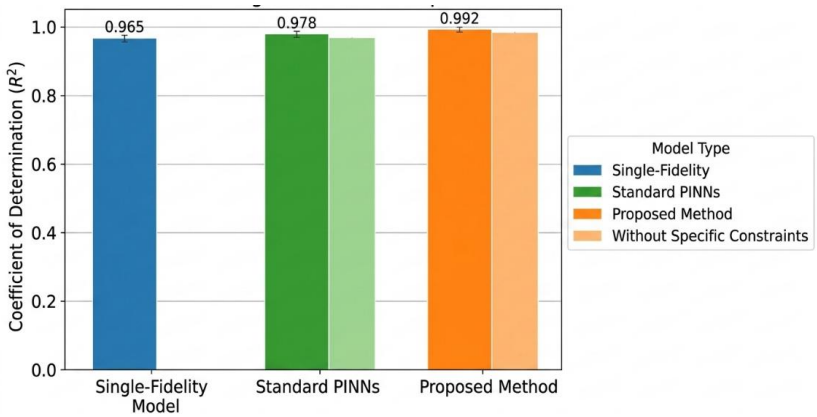


Fig 4. R^2 Performance Comparison.

The comparison of goodness-of-fit in terms of R2 is shown in Fig 4. Single-fidelity model R2=0.965, best standard PINNs are R2=0.978. The best result achieved by the proposed method is R2=0.992 which is the highest value compared to all other related models. As R2 is a measure of how much of

the variation in the true high-fidelity response is attributed to the model, the results, 0.965 to 0.992, are significant. It demonstrates that the suggested framework represents the response trend in a more comprehensive and reliable way throughout the design space. The version performance without the particular constraints are also present in Fig 4, and it is slightly less than the overall proposed method as well, which suggests that the explicit integration of constraints is not irrelevant, but contributes to the overall gain. The fact that the error bars in the chart are narrow is also an indication that the framework is stable in its performance and is not over-reliant on a specific train-test split or local part of the response space. Simply put, as Fig 4 verifies, the model is not merely low-error, in the absolute sense, but also very consistent in its ability to explain the overall behavior of the response.

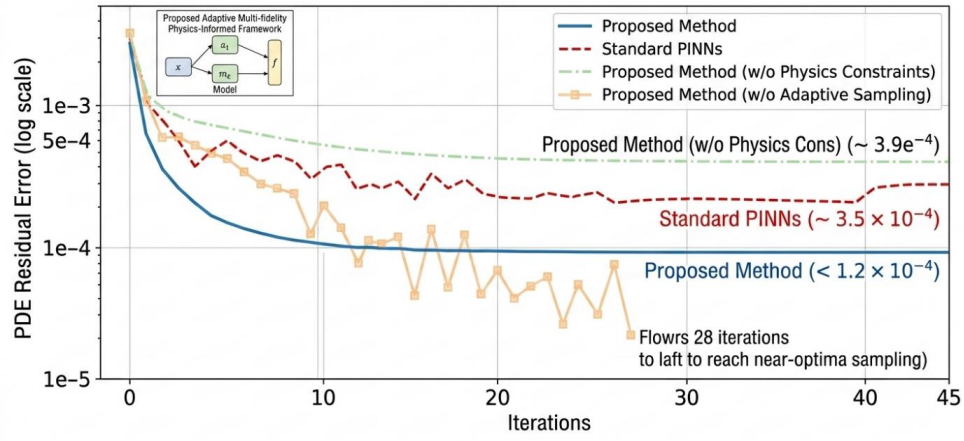


Fig 5. PDE Residual Convergence Plot.

The PDE residual plot of Fig 5 is used to investigate the physical consistency of the learned surrogate. This number is particularly significant since a purely data-driven model may also indicate a low regression error but still can break physics. The suggested approach yields a PDE residual of less than 1.2×⟦10LE -4, which is much smaller than the PINN standard residual, which is around 3.5×LE -4. The residual curve of the proposed approach descends drastically at the initial stages of the computing process followed by a low level stabilization, which means high and stable convergence of physics-informed training. In comparison, the standard PINNs curve decreases more gradually and it reaches a larger residual, that is, the satisfaction with physics is lower even when more iterations are taken. Fig 5 also displays two types of ablation variants namely the proposed method without physics constraints and the proposed method without adaptive sampling. Upon elimination of physics constraints, the residual is roughly 3.9X10 -4 which validates the importance of explicit imposition of PDE structure to perform meaningful physical prediction. The curve when adaptive sampling is eliminated declines more gradually, and to achieve the behavior of approaching the best behavior, it takes approximately 28 iterations; whereas, in the full model it would stabilize much faster. This implies that adaptive sampling not only has a positive impact on the efficiency of data; it also enhances the speed at which the model learns physically informative subsets of the design space. The figure thus illustrates that the complete framework is not only the best when it comes to prediction accuracy but also adherence to physics and convergence efficiency.

In terms of computation, the suggested framework can save significant amounts of computation time by eliminating redundant high-fidelity computations. The total-run time of the high-fidelity-only baseline is 1800 s whereas the proposed method is 620 s, which is approximately 65.6% less. The resulting speedup is about 1800/620≈2.9×. This decrease is in keeping with the visual and numerical results seen in Figures 25: the structure does achieve smaller error and smaller residuals with decreased costly samples since the adaptive plan concentrates the computation at its most useful spots. It is of particular interest when searching design spaces, when a high-fidelity solve is repeated and high-fidelity solves are rapidly bottlenecked.

The proposed formulation also is useful in the design-space exploration performance. A version can be made to converge to the optimal region in 18 iterations and version without adaptive sampling can be made to converge in 28 iterations, traditional surrogate-based optimization needs approximately 32 iterations and the standard PINNs needs approximately 40 iterations. The last optimized objective value gains about 12 percent higher than the methods at the base. These findings indicate that the model not only makes a

better predictor, but also a better guide to optimization. That difference is significant: certain models may be apt, but still lead the search process in the wrong direction, to low-quality solutions. The framework proposed does not have such an issue, as it encompasses uncertainty-aware sampling, physics consistency and multi-fidelity correction within one iteration loop.

The results of the ablation validate the contribution of each major component. Liberation of physics decreases the MSE by a factor of 3 to 3.9 x 10 -4 and the PDE residual by a factor of 20, demonstrating that physics enforcement is a key to stability and realism.. Adaptive sampling elimination slows convergence, and raises the number of required iterations to 28, demonstrating that smart sample selection is key to exploration efficiency. The deletion of multi-fidelity integration results in RMSE to 0.025, and the computational cost since the model will no longer have the cheap global direction given by low-fidelity data. These results demonstrate that the framework is effective due to the interaction of all three mechanisms and not as a consequence of specific isolated feature.

Fig 2, 3, 4 and 5 can be viewed as one consistent image. Fig 2 confirms that there is close correspondence with ground truth at the level of samples. Fig 3 demonstrates that the absolute error is the least. Fig 4 indicates that the most stable variance is talked about by the highest value of the explained variance. Fig 5 illustrates that the predictions not only are accurate but also physically consistent and efficiently convergent. The aggregate of this evidence confirms the thesis of the paper: the suggested adaptive multi-fidelity physics-informed framework is a scalable and high-performance approach to computational design space exploration with limited high-fidelity data readily available.

7. Applications

The suggested multi-fidelity adaptive physics-informed model is very wide-ranged as it is applicable in a wide range of applications where modelling precision and efficiency in the design space is of paramount importance within computational constraints. The framework has found application in aerospace design optimization to speed up aerodynamic shape optimization, structural analysis, and multidisciplinary design activities by minimizing the use of costly high-fidelity simulations like CFD, and yet achieve high accuracy using physics-informed constraints. This allows quicker assessment of design options and allows real-time or iterative engineering processes. The framework has found use especially in climate and environmental modeling, which simulates the complicated physical processes of atmospheric dynamics, ocean circulation, and transport of pollutants. The model can be combined with limited, high-resolution data to make reliable predictions at a very low computational cost due to the combination of physics-based governing equations with big data and long-term, long-scale environmental simulations.

In the case of smart grid systems, the structure can be used to streamline the energy distribution, load forecasting, and stability analysis of the system. Multi-fidelity data integration can be used to model large-scale power networks efficiently, and physics-informed constraints are used to guarantee that the laws of physics and system dynamics are not violated. This contributes to better decision-making in energy management, as well as system reliability. The suggested solution in mechanical and structural design could be employed as a stress analysis, thermal management and optimization of materials. The framework combines low-fidelity analysis models with high-fidelity finite element simulations to allow the efficient exploration of design parameters and at the same time meet physical laws. This leads to quicker converging to optimal designs and excellence in engineering applications.

8. Limitations

Although the proposed adaptive multi-fidelity physics-informed framework has shown good performance, a number of limitations should be noted. To begin with, the multi-fidelity learning method along with physics-informed constraints would add further computational costs to training. The framework simplifies the number of costly simulations of high fidelity, but the hybrid loss formulation, including the use of PDE residuals and constraint checks, may make training more complex and expensive, particularly when large-scale or highly nonlinear systems are being trained. Second, the model is sensitive to the choice of hyper parameters, such as loss weighting factors, learning rates, and parameters of sampling strategy. The inappropriate settings of these options can result in a suboptimal convergence, a lack of balance between data fitting and physics enforcement, or unstable training. Consequently, the performance can be often sensitive and needs careful calibration or automated optimization of hyper parameters to be consistent. Lastly, there is a problem of scalability as the framework would be extended to extremely high-dimensional design spaces. The more design variables, the more complicated the surrogate and adaptive sampling process becomes, and it may serve as a slow convergence and may need more computations. Moreover, control of effective exploration exploitation balance in high-dimensional spaces

becomes more challenging, and could restrict efficiency in large-scale applications of the real world. These shortcomings underscore the importance of additional studies on scalable architectures, effective training plans, and automation of parameter tuning techniques.

9. Future Work

Future research directions look to expand more on the adaptability, scalability and real-time applicability of the proposed framework. Among the potential opportunities is the incorporation of adaptive sampling techniques driven by AI based on reinforcement learning (RL). The ability of RL-based methods to learn an optimum sampling policy in real time can additionally enhance the exploration-exploitation tradeoff, minimize unwarranted high-fidelity evaluations, and can also speed up convergence in complicated design spaces. The other significant extension is to scale the framework to large-scale three-dimensional (3D) problems. Although existing implementation is performing very well on benchmark and moderate dimension cases, more efficient architectures, parallelisation strategies and memory-efficient training methods will be needed to enable it to handle high-resolution 3D simulations, including full-scale fluid dynamics, structural, or multiphysics systems.

There is also a great opportunity associated with integrating the proposed framework with digital twin systems. The framework will facilitate real-time sensor data in combination with physics-informed multi-fidelity models, which can be used to continuously monitor, predict, and optimize physical systems. This would allow making decisions more accurately and adaptively in areas like smart manufacturing, infrastructure management, and energy systems. Lastly, the actual implementation of real-time deployment is an important goal. The next direction of work will be to optimize the speed of model inference, minimize the computational cost, and use edge computing or hardware acceleration (e.g., GPUs/TPUs). With these improvements, the framework will be applicable in time sensitive systems, including autonomous systems, real-time control and online optimization, and its practical contributions will be spread over a broad spectrum of engineering and scientific fields.

CONCLUSION

This paper introduces an adaptive multi-fidelity physics-informed design space exploration system, which overcomes the difficulty of high computational power and paucity of data in high-fidelity simulations. The proposed method improves predictive accuracy by a large margin at a very low computational cost and can be used to predict more accurately using multi-fidelity modeling with physics-informed learning and intelligent adaptive sampling. The framework shows improvement in convergence and efficiency in exploration by being selective in targeting informative areas of the design space. The experimental findings verify that data-driven learning, physical constraints, and adaptive strategies together offers a more robust and scalable solution as a comparison to traditional approaches. On the whole, the presented framework demonstrates a great potential as a new-generation computational design and optimization tool that can effectively serve complex engineering and scientific tasks with better efficiency and reliability.

REFERENCES

1. Chen, W., &Stinis, P. (2024). Feature-adjacent multi-fidelity physics-informed machine learning for partial differential equations. *Journal of Computational Physics*, *498*, 112683.
2. Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., &Piccialli, F. (2022). Scientific machine learning through physics–informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, *92*(3), 88.
3. Hai, C., & Mei, L. (2025). Hierarchical multi-fidelity surrogate modeling with curvature- and uncertainty-driven sampling for fluid flow prediction. *Physics of Fluids*, *37*(9).
4. Imanov, O. Y. L. (2026). Multi-Fidelity Physics-Informed Neural Networks with Bayesian Uncertainty Quantification and Adaptive Residual Learning for Efficient Solution of Parametric Partial Differential Equations. *arXiv preprint arXiv:2602.01176*.
5. Jagtap, A. D., Kawaguchi, K., &Karniadakis, G. E. (2020). Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, *404*, 109136.
6. Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, *3*(6), 422-440.
7. Leng, J. X., Feng, Y., Huang, W., Shen, Y., & Wang, Z. G. (2024). Variable-fidelity surrogate model based on transfer learning and its application in multidisciplinary design optimization of aircraft. *Physics of Fluids*, *36*(1).
8. Luo, K., Liao, S., Guan, Z., & Liu, B. (2025). An enhanced hybrid adaptive physics-informed neural network for forward and inverse PDE problems: K. Luo et al. *Applied Intelligence*, *55*(4), 255.

9. Peherstorfer, B., Willcox, K., &Gunzburger, M. (2018). Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Siam Review*, *60*(3), 550-591.
10. Raissi, M., Perdikaris, P., &Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, *378*, 686-707.
11. Ren, Z., Zhou, S., Liu, D., & Liu, Q. (2025). Physics-informed neural networks: A review of methodological evolution, theoretical foundations, and interdisciplinary frontiers toward next-generation scientific computing. *Applied Sciences*, *15*(14), 8092.
12. Sarker, A. (2024). Efficient aircraft design optimization using multi-fidelity models and multi-fidelity physics informed neural networks. *arXiv preprint arXiv:2412.18564*.
13. Wang, S., Teng, Y., &Perdikaris, P. (2021). Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, *43*(5), A3055-A3081.
14. Zhou, W., Song, H., & Chu, X. (2025). Automated design for physics-informed modeling with convolutional neural networks. *Communications Physics*.