

A Physics-Stabilized Self-Updating Digital Twin Framework Using Physics-Informed Neural Networks for Thermal Field Prediction

Aswin Karkadakattil*

Independent researcher, Kasaragod, Kerala 671314, India

Abstract – Digital twins increasingly rely on autonomous self-updating mechanisms to remain synchronized with physical systems; however, repeated self-updating can lead to error accumulation, numerical instability, and progressive loss of physical consistency when models iteratively learn from their own predictions. To address this challenge, the study proposes a physics-stabilized self-updating digital twin framework based on Physics-Informed Neural Networks (PINNs) and demonstrates its core principles on a canonical thermal field prediction problem. The framework integrates adaptive physics-loss weighting, a physics-only stabilization stage, and second-derivative smoothness regularization within the self-updating loop, enabling controlled data assimilation while explicitly enforcing governing equation constraints. Numerical results show a monotonic reduction in root mean square error (RMSE) from approximately 1×10^{-3} in the first update cycle to 8×10^{-5} after four update cycles, accompanied by effective suppression of model drift and a substantial reduction in partial differential equation (PDE) residuals compared to a naïve self-updating strategy. Furthermore, the analysis reveals the existence of an update saturation point, beyond which additional autonomous updates yield diminishing accuracy improvements, providing a physically motivated stopping criterion for autonomous updating. By establishing a stable and physically interpretable self-updating architecture, the study provides a foundational framework for the development of reliable digital twins, with clear potential for extension to more complex thermal and Multiphysics systems.

Keywords – Digital twin, intelligent manufacturing, physics-guided learning, physics-informed neural networks, self-updating models, thermal field prediction.

I. INTRODUCTION

A. Digital Twins and the Challenge of Autonomous Updating

Digital twins are virtual representations of physical systems that integrate computational models, data streams, and system knowledge to mirror real-world behaviour across the asset lifecycle. By coupling physical assets with data-driven and physics-based models, digital twins enable real-time monitoring, prediction, optimisation, and decision support in complex engineering systems. Consequently, they have become a central component of modern industrial and cyber-physical infrastructures, particularly in manufacturing, energy systems,

and intelligent control applications [1]–[4]. A defining capability of contemporary digital twins is their ability to update autonomously as new data become available. Continuous updating is essential for capturing evolving operating conditions, material degradation, environmental variability, and long-term system aging [5]. Without such updates, digital twins gradually lose alignment with their physical counterparts, leading to reduced predictive accuracy and diminished practical usefulness. As a result, recent studies have increasingly focused on online learning and self-updating strategies to maintain synchronization between physical and virtual systems [6]. Despite these advances, autonomous self-updating introduces a fundamental challenge. When digital twins repeatedly retrain using their own predictions or limited observational data, error accumulation and numerical drift can occur. Over successive update cycles, small inaccuracies may amplify, leading to instability, loss of physical consistency, and degraded long-term performance. This phenomenon, commonly referred to as *digital twin drift*, poses a significant obstacle to trustworthy deployment, particularly in safety-critical and high-precision engineering applications. Addressing this challenge requires update strategies that preserve stability, accuracy, and physical interpretability over prolonged operational horizons.

B. Physics-Informed Neural Networks in Engineering

Physics-Informed Neural Networks (PINNs) have emerged as a powerful modelling paradigm that integrates data-driven learning with governing physical laws expressed through partial differential equations (PDEs). Unlike conventional black-box machine learning models, which rely solely on data to infer input–output relationships, PINNs embed physical constraints directly into the training process by penalizing violations of governing equations, boundary conditions, and initial conditions [7]. This physics-guided formulation offers several advantages for engineering applications. PINNs improve generalization performance in data-scarce or noisy environments, provide physically interpretable predictions that respect conservation laws, and reduce the risk of non-physical extrapolation commonly observed in purely data-driven models

* Corresponding author's e-mail: ashwinharik20000@gmail.com
Article received 2026-01-11; accepted 2026-03-02

[8]–[13]. Accordingly, PINNs have been successfully applied to a wide range of heat transfer and thermal modelling problems, including porous media, multiphase flows, and non-Fourier heat conduction [9]–[13]. In the context of digital twins, PINNs are particularly attractive because they naturally align with the dual requirements of data assimilation and physics consistency. More broadly, physics-guided and hybrid learning approaches have gained increasing attention as a means to balance data-driven flexibility with physical reliability [14]–[18]. However, while PINNs have been extensively studied for forward and inverse problems, their behaviour under repeated autonomous self-updating remains comparatively underexplored. Without careful design, even physics-informed models may experience instability when the balance between data loss and physics loss evolves across successive update cycles.

C. Research Gap

Despite the rapid growth of digital twin and PINN-based methodologies, several critical gaps remain in the existing literature. First, most self-updating digital twin studies emphasise short-term prediction accuracy, while update saturation behaviour the point beyond which additional updates no longer yield meaningful improvements has not been systematically analysed [19]–[21]. Second, mechanisms for controlling model drift during autonomous updating are often heuristic or implicit, leaving long-term stability insufficiently examined [22]–[25]. Third, many existing approaches do not explicitly enforce physics-dominant updating, allowing data-driven components to progressively overwhelm physical constraints as updates continue [14]–[18]. Importantly, addressing these challenges requires controlled problem settings in which stability, drift, and convergence behaviour can be isolated and rigorously analysed before extension to higher-dimensional or Multiphysics systems. Consequently, there remains a clear need for digital twin methodologies that explicitly balance data assimilation with physics enforcement, quantify the limits of self-updating, and provide actionable criteria for safe and reliable autonomous operation.

D. Contributions

To address the above challenges, this study makes the following key contributions:

- a physics-stabilized self-updating digital twin framework based on Physics-Informed Neural Networks is proposed, explicitly designed to preserve physical consistency during autonomous updates;
- error saturation behaviour is systematically demonstrated and quantified, showing that unlimited self-updating does not necessarily lead to continued accuracy improvement;
- governing equation residual reduction is analysed to verify enhanced physics compliance relative to naïve self-updating strategies;
- a practical stopping criterion for autonomous updating is introduced, enabling stable and computationally efficient digital twin deployment without drift-induced degradation.

Together, these contributions establish a foundational, physics-guided self-updating architecture, validated on a canonical thermal problem, and provide a robust basis for future extension to more complex digital twin applications [26]–[37].

II. GOVERNING EQUATIONS AND PROBLEM FORMULATION

A. One-Dimensional Steady-State Heat Equation

To examine the stability and physical consistency of self-updating digital twins, a one-dimensional steady-state heat conduction problem is adopted as a benchmark case. This problem is widely used in the literature due to its clear physical interpretation and availability of a closed-form analytical solution, making it well suited for validating physics-informed learning frameworks.

The governing equation is given by

$$\frac{d^2 T(x)}{dx^2} + \pi^2 \sin(\pi x) = 0, \quad x \in [0, 1],$$

where $T(x)$ denotes the temperature distribution along the spatial coordinate x . The spatially varying source term $\pi^2 \sin(\pi x)$ is introduced to produce a non-trivial temperature profile, allowing a meaningful assessment of prediction accuracy, smoothness, and physics compliance across the domain. This equation represents a steady-state heat conduction process with an internal heat source and serves as a controlled test case for evaluating the predictive performance and long-term stability of physics-informed digital twins under repeated autonomous updates.

B. Boundary Conditions

The governing equation is supplemented with homogeneous Dirichlet boundary conditions:

$$T(0) = 0, \quad T(1) = 0.$$

These boundary conditions correspond to fixed temperature constraints at both ends of the domain and ensure the existence of a unique and physically meaningful solution. Enforcing boundary conditions is essential in physics-informed modelling, as violations can lead to non-physical predictions and numerical instability, particularly in self-updating scenarios. In the proposed framework, boundary conditions are enforced implicitly through the physics-informed loss formulation, ensuring that the predicted temperature field remains consistent with the prescribed physical constraints throughout all update cycles.

C. Analytical Solution

The analytical solution of the governing equation subject to the specified boundary conditions is

$$T(x) = \sin(\pi x).$$

This closed-form solution provides a reliable reference for evaluating the accuracy and convergence behaviour of the proposed physics-stabilized self-updating PINN digital twin. By comparing the model predictions against the analytical

solution, quantitative error metrics such as the root mean square error (RMSE) can be computed to assess prediction accuracy. Moreover, deviations from the analytical solution enable direct identification of drift and instability effects during successive update cycles. The availability of an exact solution allows rigorous validation of both numerical accuracy and physics consistency, making this formulation particularly suitable for analysing the long-term behaviour of autonomous digital twin updating strategies.

III. PHYSICS-INFORMED DIGITAL TWIN FRAMEWORK

A. PINN Architecture

The physics-informed digital twin is constructed using a fully connected feedforward neural network that serves as a surrogate model for predicting the temperature field. The network takes the spatial coordinate x as input and outputs the corresponding temperature prediction $T(x)$, thereby approximating the continuous solution of the governing heat equation over the spatial domain. The network architecture consists of an input layer, multiple hidden layers, and an output layer. Each hidden layer contains a fixed number of neurons and employs the hyperbolic tangent (tanh) activation function. The tanh activation is selected due to its smooth and bounded nature, which is advantageous for approximating solutions of partial differential equations and for computing higher-order derivatives required in physics-informed learning.

The input-output mapping of the PINN can be expressed as

$$x \xrightarrow{\text{PINN}} T(x),$$

where the network parameters are optimised such that the predicted temperature field simultaneously satisfies observational data and the governing physical laws. Automatic differentiation is used to compute spatial derivatives of the network output, enabling direct evaluation of the PDE residual without the need for numerical discretization. This architecture provides a continuous and differentiable representation of the physical field, which is essential for enforcing physics constraints and maintaining stability during repeated self-updating cycles.

B. Loss Function Formulation

To ensure both data fidelity and physical consistency, training of the physics-informed digital twin is governed by a composite loss function defined as

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}} + \beta \mathcal{L}_{\text{smooth}},$$

where $\mathcal{L}_{\text{data}}$ represents the data loss, $\mathcal{L}_{\text{physics}}$ enforces compliance with the governing equation, and $\mathcal{L}_{\text{smooth}}$ is a regularization term that promotes solution smoothness. The weighting parameters λ and β control the relative influence of physics enforcement and smoothness regularization, respectively.

Data Loss

The data loss quantifies the discrepancy between the predicted temperature field and available reference data at selected spatial locations. It is defined as

$$\mathcal{L}_{\text{data}} = \frac{1}{N} \sum_{i=1}^N \left(T_{\text{pred}}(x_i) - T_{\text{ref}}(x_i) \right)^2,$$

where T_{pred} denotes the PINN prediction, T_{ref} represents the reference or measured temperature values, and N is the number of data points. This term ensures consistency with the observed system behaviour and provides accurate local predictions.

Physics (PDE Residual) Loss

The physics loss enforces compliance with the governing steady-state heat equation by penalizing violations of the PDE across the spatial domain. It is formulated as

$$\mathcal{L}_{\text{physics}} = \frac{1}{N_r} \sum_{j=1}^{N_r} \left(\frac{d^2 T(x_j)}{dx^2} + \pi^2 \sin(\pi x_j) \right)^2,$$

where N_r denotes the number of collocation points used to evaluate the residual. Spatial derivatives are computed using automatic differentiation, ensuring exact gradient evaluation with respect to the network parameters. Minimising this term constrains the neural network to learn solutions that satisfy the governing physics, thereby reducing non-physical oscillations and improving generalization during self-updating.

Second-Derivative Smoothness Regularization

To further enhance numerical stability and suppress high-frequency artefacts that may arise during repeated autonomous updates, a second-derivative smoothness regularization term is introduced:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{N_r} \sum_{j=1}^{N_r} \left(\frac{d^2 T(x_j)}{dx^2} \right)^2.$$

This term penalizes excessive curvature in the predicted temperature field, promoting smooth solutions consistent with the physical characteristics of steady-state heat conduction. The smoothness regularization plays a critical role in preventing numerical instabilities and drift amplification during successive update cycles.

C. Role of Adaptive Physics Weighting in Self-Updating

During autonomous self-updating, the relative importance of data fidelity and physics enforcement naturally evolves as the digital twin assimilates new information. To account for this behaviour, the physics weighting parameter λ is adaptively increased across successive update cycles following a predefined schedule. Specifically, λ is initialized with a lower value during early update cycles to allow data-driven correction and bias reduction, and is progressively increased in later cycles to strengthen physics dominance as the model converges. In this study, the adaptive weighting follows a monotonic update schedule across update cycles, ensuring a smooth transition

from data-assisted learning to physics-dominated stabilization. This design avoids abrupt changes in optimisation behaviour while preventing the data loss term from overwhelming physical constraints during prolonged self-updating. Importantly, the adaptation of λ is performed at the update-cycle level rather than at every training iteration, which

simplifies implementation and improves numerical stability. By dynamically balancing data loss, physics residual enforcement, and smoothness regularization within a unified framework, the proposed physics-informed digital twin maintains accuracy, suppresses drift, and delivers physically interpretable predictions throughout repeated autonomous updates.




Fig. 1. Conceptual architecture of the proposed physics-stabilized self-updating digital twin based on physics-informed neural networks (PINNs). The framework integrates observational data from the physical system with governing physical laws through a PINN core, while a controlled self-updating loop incorporating physics-only stabilization, adaptive physics-loss weighting, and data blending ensures long-term stability, physical consistency, and suppression of model drift during autonomous updates.

IV. SELF-UPDATING STRATEGY

A. Naïve Self-Updating Digital Twin

In a naïve self-updating digital twin, model parameters are iteratively updated by directly assimilating newly generated predictions or limited observational data, without explicit mechanisms to preserve physical dominance. In this setting, the digital twin progressively retrains using its own outputs as surrogate ground truth, implicitly assuming that repeated updates will lead to monotonic improvements in accuracy. While such an approach may produce short-term error reduction, it lacks safeguards against error accumulation and feedback amplification. Small inaccuracies introduced during early updates can be re-ingested in subsequent training cycles, leading to progressively amplified deviations. Moreover, when the data loss term dominates the optimisation process, the model may prioritize local data fitting at the expense of global physical consistency. Over time, this imbalance can cause the predicted solution to drift away from the governing physical laws, resulting in non-physical oscillations, loss of smoothness, and eventual numerical instability. The absence of explicit physics dominance and stabilization mechanisms makes naïve self-updating digital twins particularly vulnerable to model drift, thereby limiting their reliability for long-term autonomous operation.

B. Physics-Stabilized Self-Updating Strategy

To overcome the limitations of naïve self-updating, a physics-stabilized self-updating strategy is proposed that explicitly enforces physical consistency throughout the update process. The strategy integrates three complementary mechanisms: a physics-only refresh stage, adaptive physics-loss weighting, and controlled data blending.

Physics-only refresh stage. Before assimilating new data during each update cycle, the digital twin undergoes a physics-only training phase in which optimisation is driven exclusively by the governing equation residuals and boundary conditions. During this stage, the data loss term is temporarily disabled, allowing the model to re-align with the underlying physical laws and correct deviations accumulated during previous updates. This refresh stage acts as a stabilization step that anchors the digital twin to physically admissible solution manifolds.

Adaptive physics-loss weighting. To ensure long-term stability, the relative weight of the physics loss is progressively increased across successive update cycles. Early updates emphasize data-driven correction to reduce initial bias, while later updates prioritize physics enforcement to suppress drift and maintain global consistency. This adaptive weighting strategy balances data fidelity and physical constraints in a controlled manner, preventing either component from

dominating the optimisation process during prolonged self-updating.

Controlled data blending strategy. Rather than directly reusing predicted values as ground truth, a controlled data blending approach is adopted to assimilate new information gradually. Specifically, the updated reference data at update cycle $k + 1$ are constructed as

$$T_{\text{ref}}^{(k+1)}(x) = \alpha T_{\text{ref}}^{(k)}(x) + (1 - \alpha) T_{\text{pred}}^{(k)}(x),$$

where $T_{\text{ref}}^{(k)}$ denotes the reference data used in update cycle k , $T_{\text{pred}}^{(k)}$ represents the model prediction from the previous cycle, and $\alpha \in [0, 1]$ is a blending coefficient. In this study, α is chosen to favour existing reference data, thereby limiting the influence of purely synthetic predictions during early updates. This gradual assimilation prevents abrupt shifts in the training distribution and reduces the risk of reinforcing prediction errors through feedback loops. By constraining the contribution of predicted data, the digital twin maintains robustness against error amplification while still benefiting from autonomous updating. Collectively, these mechanisms enable stable,

accurate, and physically interpretable self-updating behaviour, even over multiple autonomous update cycles.

C. Update Saturation Concept

Although self-updating is essential for maintaining alignment between the digital twin and the physical system, unlimited autonomous updating is neither necessary nor beneficial. As demonstrated in the results, successive updates eventually lead to diminishing accuracy gains, indicating the presence of an update saturation point. Beyond this saturation point, additional updates yield marginal improvements while increasing computational cost and the risk of numerical instability. In naïve self-updating frameworks, continued updating beyond saturation can exacerbate drift and degrade solution quality. In contrast, the proposed physics-stabilized strategy enables identification of a practical stopping criterion based on convergence behaviour and error reduction trends. Recognising and quantifying update saturation is therefore critical for the efficient and safe deployment of autonomous digital twins. By restricting updates to the region of meaningful improvement, the digital twin achieves optimal predictive performance while maintaining long-term stability and physical consistency.




Fig. 2. Comparison between naïve and physics-stabilized self-updating strategies. The figure illustrates two parallel self-updating workflows. In the naïve strategy, predicted outputs are directly recycled into the training process through an unconstrained feedback loop, leading to progressive drift and instability. In contrast, the physics-stabilized strategy incorporates a physics-only refresh stage, adaptive physics loss weighting, and controlled data blending, resulting in a damped and stable update loop that converges toward a physically consistent solution.

V. EXPERIMENTAL SETUP AND IMPLEMENTATION

A. Training Details

The physics-informed digital twin is implemented using a fully connected neural network trained via gradient-based optimisation. The Adam optimiser is employed due to its robustness and proven effectiveness in handling the non-convex optimisation landscapes commonly encountered in neural network training. Its adaptive moment estimation facilitates stable and consistent convergence, particularly when balancing competing data-driven and physics-based loss components.

Each training phase is performed for a fixed number of epochs to ensure sufficient convergence while mitigating the risk of overfitting. For the naïve self-updating strategy, the model is trained using both data loss and physics loss terms for 1000 epochs per update cycle. In contrast, the proposed physics-stabilized strategy incorporates a two-stage training procedure. First, a physics-only training stage is applied, during which optimisation is driven exclusively by the governing equation residuals and boundary conditions. This is followed by a full training stage that combines data and physics losses. This two-

stage approach ensures that the model is re-aligned with the underlying physical laws before assimilating new data, thereby improving numerical stability during autonomous updating. A constant learning rate of 1×10^{-3} is used across all training phases, providing a balance between convergence speed and numerical robustness throughout repeated update cycles.

B. Update Cycle Definition

An update cycle is defined as a complete sequence of training and data assimilation steps through which the digital twin updates its internal model parameters. Each update cycle consists of the following stages:

- *Model training*: Neural network parameters are optimised using the composite loss function comprising data loss, physics residual loss, and smoothness regularization.
- *Prediction generation*: The trained model produces updated temperature predictions over the spatial domain.
- *Data assimilation*: For the physics-stabilized strategy, newly predicted data are blended with existing reference data using weighted averaging to prevent abrupt distribution shifts and feedback-induced drift.
- *Physics stabilization*: A physics-only refresh stage is applied at the beginning of each update cycle to reinforce compliance with governing equations and boundary conditions.

A total of four update cycles are performed in this study. This choice is motivated by the observed convergence behaviour, wherein the majority of accuracy improvement occurs within the initial update cycles, followed by a clear reduction in incremental gains. As demonstrated in the results, error saturation is reached by the third to fourth update cycle, indicating that additional updates provide marginal benefits while increasing computational cost. Consequently, four update cycles are sufficient to capture convergence trends, update saturation behaviour, and long-term stability characteristics without introducing unnecessary computational overhead.

This update cycle formulation enables a systematic evaluation of both short-term accuracy improvement and stability during autonomous self-updating.

TABLE I
MODEL AND TRAINING PARAMETERS

Parameter	Value
Network depth	3 layers
Neurons per layer	64
Activation function	Tanh
Optimizer	Adam
Learning rate	1×10^{-3}
Update cycles	4

VI. RESULTS AND DISCUSSION

This section presents a comprehensive evaluation of the proposed physics-stabilized self-updating PINN digital twin. The results are analysed in terms of prediction accuracy, convergence behaviour, physics consistency, and long-term stability across successive update cycles.

A. Accuracy and Convergence Analysis

The predictive accuracy of the stabilized digital twin is first assessed by comparing the predicted temperature field with the analytical solution, as shown in Fig. 3. The stabilized PINN demonstrates excellent agreement with the true temperature distribution across the entire spatial domain, including regions near the boundaries. The close overlap between the predicted and analytical temperature profiles indicates that the proposed framework accurately captures the underlying thermal behaviour without introducing spurious oscillations or numerical artefacts.




Fig. 3. Comparison of the analytical temperature field and the stabilized PINN prediction across the spatial domain, demonstrating accurate and physically consistent thermal field reconstruction.

The convergence behaviour across update cycles is quantified using the root mean square error (RMSE), as illustrated in Fig. 4. A clear monotonic reduction in RMSE is observed from the first to the third update cycle, demonstrating the effectiveness of autonomous self-updating when guided by physics-based stabilization. Importantly, no error rebound or oscillatory convergence behaviour is observed, indicating numerically stable learning.




Fig. 4. The evolution of the root mean square error (RMSE) across successive digital twin update cycles. A clear monotonic reduction in RMSE is observed, demonstrating stable convergence of the proposed physics-stabilized self-updating framework and confirming the effectiveness of controlled physics-guided updating.

A direct comparison between the naïve and stabilized self-updating strategies is shown in Fig. 5. Although the naïve approach yields visually reasonable predictions, noticeable deviations emerge in regions of higher curvature. In contrast, the physics-stabilized strategy maintains closer agreement with the analytical solution across the entire domain, highlighting the advantage of enforcing physical constraints even when overall prediction accuracy appears high.




Fig. 5. Comparison of the predictions obtained from the naïve and physics-stabilized self-updating PINN frameworks. The stabilized approach demonstrates improved accuracy and robustness, reflecting the benefits of enforcing physics-guided stabilization during autonomous updates.

B. Physics Consistency and Residual Reduction

Beyond predictive accuracy, physical consistency is a fundamental requirement for trustworthy digital twin deployment. The spatial distribution of the governing equation residual is presented in Fig. 6. The naïve self-updating model exhibits noticeably higher residual magnitudes across the spatial domain, indicating partial violation of the governing heat equation and reduced adherence to the underlying physical constraints.




Fig. 6. The spatial distribution of the governing equation residual for both naïve and physics-stabilized self-updating strategies. The stabilized digital twin achieves an order-of-magnitude reduction in PDE residuals, with residuals remaining low and smoothly distributed across the domain. This behaviour confirms that the adaptive physics-loss weighting and the physics-only refresh stage effectively constrain the learning process to physically admissible solution manifolds.

In contrast, the physics-stabilized framework maintains substantially lower residual values, confirming that explicit enforcement of governing equations effectively constrains the learning process and suppresses non-physical behaviour during successive update cycles.

The evolution of prediction error on a logarithmic scale is shown in Fig. 7, further highlighting the advantages of the stabilized approach. While the naïve self-updating strategy exhibits early stagnation and limited improvement, the physics-stabilized framework demonstrates sustained error reduction across successive update cycles. This contrast underscores the importance of explicit physics dominance in preventing performance degradation and maintaining long-term stability during autonomous updating.




Fig. 7. The log-scale evolution of RMSE for the naïve and physics-stabilized digital twins across successive update cycles. The stabilized framework exhibits sustained error reduction over multiple updates, whereas the naïve approach demonstrates early stagnation. This contrast further highlights the effectiveness of physics-guided stabilization in maintaining long-term predictive performance.

C. Digital Twin Stability and Update Saturation

While autonomous self-updating improves prediction accuracy, unrestricted updating is neither computationally efficient nor physically necessary, even in controlled problem settings. Figure 8 shows the change in RMSE (Δ RMSE) between successive update cycles, revealing a pronounced reduction in error improvement after the third update cycle. This behaviour indicates the presence of an update saturation point, beyond which additional autonomous updates provide diminishing accuracy gains. Importantly, this saturation behaviour reflects an inherent property of the self-updating loop rather than a limitation of the underlying physical model. Identifying such saturation points enables the formulation of a practically meaningful stopping criterion, allowing autonomous digital twins to achieve stable convergence while avoiding unnecessary computation and potential drift amplification during prolonged operation.




Fig. 8. The change in RMSE (Δ RMSE) across successive update cycles, clearly revealing update saturation behaviour and motivating the use of a practical stopping criterion for autonomous digital twin updating.

The stability of the physics-stabilized digital twin is further validated through an analysis of second-derivative smoothness, as shown in Fig. 9. The predicted second derivative closely follows the corresponding analytical trend, confirming the absence of high-frequency numerical artefacts. This smoothness is essential for suppressing drift and ensuring long-term physical interpretability during repeated autonomous updates.




Fig. 9. Second derivative of the predicted temperature field obtained from the stabilized digital twin, confirming smoothness and suppression of non-physical oscillations.

Identifying update saturation provides a practical stopping criterion for autonomous digital twins, enabling efficient deployment while minimising the risk of instability or unnecessary computational cost.

The values in Table II quantitatively confirm the convergence trends observed in Figs. 4 and 8, demonstrating diminishing accuracy gains beyond the third update cycle.

TABLE II
RMSE ACROSS UPDATE CYCLES

Update Cycle	RMSE
1	$\sim 1 \times 10^{-3}$
2	$\sim 4 \times 10^{-4}$
3	$\sim 1 \times 10^{-4}$
4	$\sim 8 \times 10^{-5}$

VII. DISCUSSION: PRACTICAL IMPLICATIONS

The findings of this study provide several important insights for the deployment of autonomous digital twins in engineering systems. One of the most persistent challenges in operational digital twins is model drift, whereby repeated self-updating gradually leads to deviation from physical reality and degradation of predictive reliability. This issue has been widely recognised in recent digital twin and hybrid modelling studies, particularly in systems that rely on online learning and limited observational data [1], [5], [6]. The proposed physics-stabilized self-updating framework directly addresses this challenge by explicitly enforcing governing equations and smoothness constraints throughout the update process. By anchoring the learning dynamics to physical laws, the digital twin maintains long-term reliability and mitigates the accumulation of non-physical errors, thereby enhancing overall trustworthiness. Another key implication of the proposed framework concerns the trade-off between computational cost and long-term reliability. The identification of an update saturation point demonstrates that continuous or indefinite updating is neither necessary nor efficient, even in controlled problem settings. Similar observations regarding diminishing returns and convergence saturation have been reported in autonomous and continual learning systems [32]–[34]. Once convergence is achieved, additional updates yield only marginal accuracy gains while increasing computational cost and the risk of numerical instability. By introducing a principled stopping criterion based on convergence behaviour, the proposed approach enables efficient allocation of computational resources without compromising predictive stability. From an application perspective, the framework provides a reliable architectural foundation for digital twin deployment in manufacturing and cyber-physical systems, where digital twins are increasingly used for process monitoring, thermal management, and decision support under evolving operating conditions [2], [3], [26]–[29]. In such environments, autonomous updating must be accompanied by strong guarantees of physical consistency to maintain operator confidence and system safety. The physics-stabilized strategy ensures that predictions remain interpretable, stable, and aligned with known physical behaviour, positioning the framework as a robust basis for future integration into industrial digital twin platforms and safety-critical applications.

VIII. LIMITATIONS AND FUTURE WORK

While the proposed framework demonstrates stable and accurate autonomous self-updating within the scope of this study, several aspects naturally motivate further extension and validation. First, the present investigation focuses on a steady-

state one-dimensional heat conduction problem, which provides a controlled and transparent setting for isolating and analysing update stability, drift suppression, and saturation behaviour. Such canonical formulations are commonly adopted to establish foundational algorithmic properties. Extending the framework to transient heat transfer problems represents a natural next step and would enable systematic evaluation of time-dependent updating dynamics, temporal stability, and real-time adaptation under evolving boundary and loading conditions. Second, the current implementation assumes noise-free reference data in order to clearly examine the intrinsic effects of physics stabilization and controlled self-updating. In practical deployments, sensor measurements are often affected by noise, bias, and incomplete data availability. Incorporating noisy and sparse observations, together with uncertainty-aware or probabilistic loss formulations, would further enhance robustness under realistic operating conditions while preserving physical consistency. Finally, the framework is demonstrated on a one-dimensional spatial domain to facilitate rigorous analysis and visualisation of stability mechanisms. Many industrial digital twin applications involve higher-dimensional geometries and coupled Multiphysics interactions, such as thermo-mechanical deformation or fluid-thermal coupling. Extending the proposed physics-stabilized self-updating strategy to two- and three-dimensional problems, as well as to coupled physical systems, therefore represents an important area for future research. Collectively, these extensions reflect the generality and modular design of the proposed framework rather than inherent limitations, and outline a clear and practical roadmap toward increasingly realistic and complex digital twin applications.

IX. CONCLUSIONS

The study presented a physics-stabilized self-updating digital twin framework based on Physics-Informed Neural Networks and demonstrated its core principles on a canonical thermal field prediction problem. By integrating adaptive physics-loss weighting, a physics-only stabilization stage, and second-derivative smoothness regularization into the self-updating process, the proposed framework preserves physical consistency and effectively mitigates drift during autonomous model updates. The results show that controlled self-updating consistently outperforms naïve self-updating strategies, achieving monotonic error reduction, substantial minimisation of governing equation residuals, and stable convergence across successive update cycles. Importantly, the analysis identifies the presence of an update saturation point, beyond which additional autonomous updates yield diminishing accuracy gains. This insight provides a physically motivated and practically meaningful stopping criterion, enabling improved computational efficiency while maintaining predictive reliability. Overall, the proposed framework establishes a robust and interpretable self-updating architecture that explicitly balances data assimilation with physics enforcement. By clarifying the mechanisms underlying stability, drift suppression, and update saturation, the study provides a foundational basis for the development of reliable autonomous

digital twins, with clear potential for extension to more complex thermal and Multiphysics applications.

DECLARATIONS

Consent to Publish

Not applicable.

Consent to Participate

Not applicable.

Ethics Approval and Consent to Participate

Not applicable.

Competing Interests

The author declares that there are no competing financial or non-financial interests related to this work.

Availability of Data and Materials

No experimental datasets were generated or analysed in this study. All results were obtained using a physics-informed modelling framework based on governing equations, prescribed boundary conditions, and analytical reference solutions. The model implementation details and representative outputs can be made available by the author upon reasonable request.

Author Contributions

The author conceived and designed the study and developed the physics-stabilized self-updating digital twin framework. All theoretical formulation, model implementation, numerical training, parametric analysis, and post-processing were performed by the author. The author analysed and interpreted the results, prepared all figures and tables, wrote the original manuscript, and reviewed and approved the final version.

REFERENCES

- [1] W. Birk, R. Hostettler, M. Razi, K. Atta, and R. Tammia, "Automatic generation and updating of process industrial digital twins for estimation and control: A review," *Frontiers in Control Engineering*, vol. 3, Aug. 2022, Art. no. 954858. <https://doi.org/10.3389/fcteg.2022.954858>
- [2] A. Kantaros, T. Ganetsos, E. Pallis, and M. Papoutsidakis, "From mathematical modeling and simulation to digital twins: Bridging theory and digital realities in industry and emerging technologies," *Applied Sciences*, vol. 15, no. 16, Aug. 2025, Art. no. 9213. <https://doi.org/10.3390/app15169213>
- [3] H. Peter, Digital Twins and AI: Real-Time Simulation of Complex Systems, Apr. 2025.
- [4] M. Diakité and M. K. Traore, "Formalizing a framework of inference capabilities for digital twin engineering," *SIMULATION*, vol. 100, no. 9, pp. 887–902, Jan. 2024. <https://doi.org/10.1177/00375497241228281>
- [5] Y. Tang, M. R. Dehaghani, and G. G. Wang, "Capturing lifecycle system degradation in digital twin model updating," *arXiv preprint arXiv:2503.08953*, Mar. 2025. <https://doi.org/10.48550/arXiv.2503.08953>
- [6] T. Wu, F. Yang, U. Farooq, X. Li, and J. Jiang, "An online learning method for constructing self-update digital twin model of power transformer temperature prediction," *Applied Thermal Engineering*, vol. 237, Jan. 2024, Art. no. 121728. <https://doi.org/10.1016/j.applthermaleng.2023.121728>
- [7] S. Cai, Z. Wang, S. Wang, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks for heat transfer problems," *Journal of Heat Transfer*, vol. 143, no. 6, June 2021, Art. no. 060801. <https://doi.org/10.1115/1.4050542>

- [8] R. Laubscher, "Simulation of multi-species flow and heat transfer using physics-informed neural networks," *Physics of Fluids*, vol. 33, no. 8, Aug. 2021, Art. no. 087101. <https://doi.org/10.1063/5.0058529>
- [9] J. Xu, H. Wei, and H. Bao, "Physics-informed neural networks for studying heat transfer in porous media," *International Journal of Heat and Mass Transfer*, vol. 217, Dec. 2023, Art. no. 124671. <https://doi.org/10.1016/j.ijheatmasstransfer.2023.124671>
- [10] S. Manavi, T. Becker, and E. Fattahi, "Enhanced surrogate modelling of heat conduction problems using physics-informed neural network framework," *International Communications in Heat and Mass Transfer*, vol. 142, Mar. 2023, Art. no. 106662. <https://doi.org/10.1016/j.icheatmasstransfer.2023.106662>
- [11] D. Jalili, S. Jang, M. Jadidi, G. Giustini, A. Keshmiri, and Y. Mahmoudi, "Physics-informed neural networks for heat transfer prediction in two-phase flows," *International Journal of Heat and Mass Transfer*, vol. 221, Apr. 2024, Art. no. 125089. <https://doi.org/10.1016/j.ijheatmasstransfer.2023.125089>
- [12] J. Zheng, F. Li, and H. Huang, "T-phPINN: Physics-informed neural networks for solving 2D non-Fourier heat conduction equations," *International Journal of Heat and Mass Transfer*, vol. 235, Dec. 2024, Art. no. 126216. <https://doi.org/10.1016/j.ijheatmasstransfer.2024.126216>
- [13] P. Cui, W. Fan, M. Yu, W. Liu, and Z. Liu, "Solution and applications of parameterized convective heat transfer equations using physics-informed neural networks," *International Journal of Heat and Mass Transfer*, vol. 247, Sep. 2025, Art. no. 127040. <https://doi.org/10.1016/j.ijheatmasstransfer.2025.127040>
- [14] S. Pawar, O. San, B. Aksoylu, A. Rasheed, and T. Kvamsdal, "Physics guided machine learning using simplified theories," *Physics of Fluids*, vol. 33, no. 1, Jan. 2021, Art. no. 011701. <https://doi.org/10.1063/5.0038929>
- [15] R. Rai and C. K. Sahu, "Driven by data or derived through physics? A review of hybrid physics guided machine learning techniques with cyber-physical system (CPS) focus," *IEEE Access*, vol. 8, pp. 71050–71073, Apr. 2020. <https://doi.org/10.1109/ACCESS.2020.2987324>
- [16] C. Meng, S. Griesemer, D. Cao, S. Seo, and Y. Liu, "When physics meets machine learning: A survey of physics-informed machine learning," *Machine Learning for Computational Science and Engineering*, vol. 1, no. 1, May 2025, Art. no. 20. <https://doi.org/10.1007/s44379-025-00016-0>
- [17] S. M. Razak, J. Cornelio, Y. Cho, H.-H. Liu, R. Vaidya, and B. Jafarpour, "Physics-guided deep learning for improved production forecasting in unconventional reservoirs," *SPE Journal*, vol. 28, no. 5, pp. 2425–2447, Oct. 2023. <https://doi.org/10.2118/214663-PA>
- [18] E. He, Y. Xie, L. Liu, Z. Jin, D. Zhang, and X. Jia, "Knowledge guided machine learning for extracting, preserving, and adapting physics-aware features," in *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*, 2024, pp. 715–723. <https://doi.org/10.1137/1.9781611978032.82>
- [19] F. Zahra, J. Macedo, and C. Málaga-Chuquitaype, "Hybrid data-driven hazard-consistent drift models for SMRF," *Earthquake Engineering & Structural Dynamics*, vol. 52, no. 4, pp. 1112–1135, Jan. 2023. <https://doi.org/10.1002/eqe.3807>
- [20] Q. Wang, F. Li, Y. Tang, and Y. Xu, "Integrating model-driven and data-driven methods for power system frequency stability assessment and control," *IEEE Transactions on Power Systems*, vol. 34, no. 6, pp. 4557–4568, May 2019. <https://doi.org/10.1109/TPWRS.2019.2919522>
- [21] H. Toubakh and M. Sayed-Mouchaweh, "Hybrid dynamic data-driven approach for drift-like fault detection in wind turbines," *Evolving Systems*, vol. 6, no. 2, pp. 115–129, Oct. 2015. <https://doi.org/10.1007/s12530-014-9119-8>
- [22] F. Li, Q. Wang, Y. Tang, Y. Xu, and J. Dang, "Hybrid analytical and data-driven model based instance-transfer method for power system online transient stability assessment," *CSEE Journal of Power and Energy Systems*, vol. 10, no. 4, pp. 1664–1675, July 2021. <https://doi.org/10.17775/CSEEJPES.2020.03880>
- [23] X.-B. Jin, R. J. R. Jeremiah, T.-L. Su, Y.-T. Bai, and J.-L. Kong, "The new trend of state estimation: From model-driven to hybrid-driven methods," *Sensors*, vol. 21, no. 6, Mar. 2021, Art. no. 2085. <https://doi.org/10.3390/s21062085>
- [24] S. Windmann, "Data-driven fault detection in industrial batch processes based on a stochastic hybrid process model," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 3888–3902, Oct. 2022. <https://doi.org/10.1109/TASE.2021.3138925>
- [25] F. Balali, H. Seifoddini, and A. Nasiri, "Data-driven predictive model of reliability estimation using degradation models: A review," *Life Cycle Reliability and Safety Engineering*, vol. 9, no. 1, pp. 113–125, Feb. 2020. <https://doi.org/10.1007/s41872-020-00111-6>
- [26] K. Kobayashi and S. B. Alam, "Explainable, interpretable, and trustworthy AI for an intelligent digital twin: A case study on remaining useful life," *Engineering Applications of Artificial Intelligence*, vol. 129, Mar. 2024, Art. no. 107620. <https://doi.org/10.1016/j.engappai.2023.107620>
- [27] S. Wang, X. Lai, X. He, Y. Qiu, and X. Song, "Building a trustworthy product-level shape-performance integrated digital twin with multifidelity surrogate model," *Journal of Mechanical Design*, vol. 144, no. 3, Mar. 2022, Art. no. 031703. <https://doi.org/10.1115/1.4052390>
- [28] J. Jeon and G. Theotokatos, "A framework to assure the trustworthiness of physical model-based digital twins for marine engines," *Journal of Marine Science and Engineering*, vol. 12, no. 4, Mar. 2024, Art. no. 595. <https://doi.org/10.3390/jmse12040595>
- [29] N. Bicocchi, M. Fogli, C. Giannelli, M. Picone, and A. Virdis, "Requirements and design architecture for digital twin end-to-end trustworthiness," *IEEE Internet Computing*, vol. 28, no. 4, pp. 31–39, Mar. 2024. <https://doi.org/10.1109/MIC.2024.3376439>
- [30] Z. Pervez, Z. Khan, A. Ghafoor, and K. Soomro, "Signed: Smart city digital twin verifiable data framework," *IEEE Access*, vol. 11, pp. 29430–29446, Mar. 2023. <https://doi.org/10.1109/ACCESS.2023.3260621>
- [31] I. A. Ridhawi, S. Otoum, and M. Aloqaily, "Decentralized zero-trust framework for digital twin-based 6G," *arXiv preprint arXiv:2302.03107*, Feb. 2023. <https://doi.org/10.48550/arXiv.2302.03107>
- [32] K. Shaheen, M. A. Hanif, O. Hasan, and M. Shafique, "Continual learning for real-world autonomous systems: Algorithms, challenges and frameworks," *Journal of Intelligent & Robotic Systems*, vol. 105, no. 1, Apr. 2022, Art. no. 9. <https://doi.org/10.1007/s10846-022-01603-6>
- [33] X. Gu, J. Han, Q. Shen, and P. P. Angelov, "Autonomous learning for fuzzy systems: A review," *Artificial Intelligence Review*, vol. 56, no. 8, pp. 7549–7595, Dec. 2023. <https://doi.org/10.1007/s10462-022-10355-6>
- [34] J. Zeng, S.-B. Lin, Y. Yao, and D.-X. Zhou, "On ADMM in deep learning: Convergence and saturation-avoidance," *Journal of Machine Learning Research*, vol. 22, no. 1, Jan. 2021, Art. no. 199. <https://dl.acm.org/doi/abs/10.5555/3546258.3546457>
- [35] C. Wu, A. R. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, "Flow: A modular learning framework for mixed autonomy traffic," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1270–1286, July 2021. <https://doi.org/10.1109/TRO.2021.3087314>
- [36] K. G. Vamvoudakis, M. F. Miranda, and J. P. Hespanha, "Asymptotically stable adaptive-optimal control algorithm with saturating actuators and relaxed persistence of excitation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 27, no. 11, pp. 2386–2398, Oct. 2015. <https://doi.org/10.1109/TNNLS.2015.2487972>
- [37] W. Senn and S. Fusi, "Convergence of stochastic learning in perceptrons with binary synapses," *Physical Review E*, vol. 71, no. 6, July 2005, Art. no. 061907. <https://doi.org/10.1103/PhysRevE.71.061907>

Aswin Karkadatil received his M.Tech in Materials and Manufacturing Engineering from the Indian Institute of Technology (IIT) and his B.Tech in Mechanical Engineering from the Government College of Engineering, Kannur. He was a recipient of the Prime Minister's Scholarship for Academic Excellence during his undergraduate studies. His research experience spans physics-informed modelling, digital twins, and intelligent manufacturing systems, with a particular focus on integrating physical laws with data-driven learning for reliable and stable engineering predictions. He has authored and co-authored several Scopus- and SCI-indexed journal articles in the areas of manufacturing processes, materials processing, and aerospace-related applications. His current research interests include physics-informed digital twins, laser-based post-processing, AI-driven manufacturing, and hybrid optimisation frameworks for sustainable and intelligent engineering systems.

E-mail: ashwinharik20000@gmail.com

APPENDIX A: PHYSICS-STABILIZED SELF-UPDATING PINN ALGORITHM

This appendix summarises the proposed physics-stabilized self-updating digital twin framework in an algorithmic form. The formulation highlights the integration of physics-only stabilization, adaptive physics weighting, and controlled data blending within the self-updating loop.

Algorithm A1: Physics-Stabilized Self-Updating PINN Digital Twin

I. Inputs

- Initial reference dataset $\mathcal{D}^{(0)} = \{x_i, T_{\text{ref}}^{(0)}(x_i)\}_{i=1}^N$
- Collocation points $\{x_j\}_{j=1}^{N_r}$
- Maximum update cycles K
- Initial physics weight $\lambda^{(0)}$
- Smoothness weight β
- Data blending coefficient $\alpha \in [0, 1]$

II. Output

- Stabilized digital twin model $\mathcal{M}^{(k)}$

III. Initialization

Initialize network parameters $\theta^{(0)}$

Set update index $k = 0$

IV. Self-Updating Procedure

While $k < K$ do

Step 1: Physics-Only Stabilization. Realign the model with governing physical laws by minimising the physics loss:

$$\theta^{(k)} \leftarrow \arg \min_{\theta} \mathcal{L}_{\text{physics}}(\theta) \quad (\text{A1})$$

Step 2: Full Physics-Informed Training. Optimise the composite loss function:

$$\mathcal{L}^{(k)} = \mathcal{L}_{\text{data}}^{(k)} + \lambda^{(k)} \mathcal{L}_{\text{physics}} + \beta \mathcal{L}_{\text{smooth}} \quad (\text{A2})$$

Step 3: Prediction Generation. Compute the updated temperature field:

$$T_{\text{pred}}^{(k)}(x) = \mathcal{M}^{(k)}(x) \quad (\text{A3})$$

Step 4: Controlled Data Blending. Update the reference data for the next cycle:

$$T_{\text{ref}}^{(k+1)}(x) = \alpha T_{\text{ref}}^{(k)}(x) + (1 - \alpha) T_{\text{pred}}^{(k)}(x) \quad (\text{A4})$$

Step 5: Adaptive Physics Weight Update. Increase physics dominance across update cycles:

$$\lambda^{(k+1)} = \lambda^{(k)} + \Delta\lambda \quad (\text{A5})$$

Step 6: Increment Update Counter

$$k \leftarrow k + 1 \quad (\text{A6})$$

End While

V. Output

Return the stabilized digital twin model $\mathcal{M}^{(k)}$

Appendix A.1: Loss Function Definitions

Data Loss

$$\mathcal{L}_{\text{data}} = \frac{1}{N} \sum_{i=1}^N (T_{\text{pred}}(x_i) - T_{\text{ref}}(x_i))^2 \quad (\text{A7})$$

Physics (PDE Residual) Loss

$$\mathcal{L}_{\text{physics}} = \frac{1}{N_r} \sum_{j=1}^{N_r} \left(\frac{d^2 T(x_j)}{dx^2} + \pi^2 \sin(\pi x_j) \right)^2 \quad (\text{A8})$$

Smoothness Regularization

$$\mathcal{L}_{\text{smooth}} = \frac{1}{N_r} \sum_{j=1}^{N_r} \left(\frac{d^2 T(x_j)}{dx^2} \right)^2 \quad (\text{A9})$$