

Article

Reduction of Interface-Induced Order Degradation via a Conditional Hybrid Neural-IMEX Framework

Mouloud Aoudia 

Department of Industrial Engineering, College of Engineering, Northern Border University, Arar 91431, Saudi Arabia; mouloud.aoudia@nbu.edu.sa

Abstract

High-order implicit-explicit (IMEX) schemes are effective for stiff parabolic partial differential equations when temporal regularity is compatible with the active multistep stencil. In moving-interface problems, a fixed Eulerian node may undergo a rapid transition as a diffuse interface crosses the grid, allowing stored multistep history to mix incompatible local regimes. This paper develops a conditional hybrid Deep Neural Galerkin-IMEX (DNG-IMEX) framework for this order-degradation mechanism. A classical IMEX-BDF3 backbone is retained on smooth intervals, whereas flagged event windows are treated by a localized neural/subcycle bridge followed by restart-consistent history reconstruction. The formulation separates the weak parabolic setting from the additional smoothness used for pointwise interface kinematics and proves a Sobolev-level transport estimate, a weak energy estimate, and a conditional propagation result under explicit flagged/restart defect bounds. Numerical tests on a manufactured Allen–Cahn benchmark show that event-aligned restarting suppresses the dominant history-contamination defect. A benchmark diagnostic realization localizes corrections with available event information and improves the baseline when event windows are resolved and the detector remains selective. Interface-thickness and cost tests indicate that sharper interfaces require stronger event resolution and that the present correction pipeline has non-negligible overhead. These findings support selective interface-aware enhancement of classical IMEX time integration.

Keywords: implicit–explicit schemes; neural Galerkin methods; temporal order reduction; stiff parabolic PDEs; Allen–Cahn equation; interface crossing

MSC: 65M12; 65M20; 65M60; 35K58; 68T07



Academic Editors: Juan Ramón Torregrosa Sánchez and Luigi Rodino

Received: 4 April 2026

Revised: 13 May 2026

Accepted: 23 May 2026

Published: 2 June 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

1.1. Background and Motivation

Stiff time-dependent partial differential equations with sharp transitions, moving interfaces, and localized temporal irregularities arise when diffusion, reaction, transport, geometry, or multiphysics coupling creates narrow regions of intense activity inside a larger smooth domain. Moving-grid formulations were developed precisely to concentrate resolution where steep layers travel through the mesh [1], while adaptive-grid surveys established the broader need for spatial redistribution in such problems [2]. Phase-field methods provide a diffuse-interface description of geometric moving interfaces and their numerical approximations [3]. Level-set/interface-capturing work continues this theme for fronts whose topology or geometry changes during the computation [4,5]. High-order discontinuous Galerkin and WENO-type methods address a complementary difficulty:

representing convection-dominated steep profiles without unacceptable oscillation or excessive smearing [6–8].

The present research landscape is therefore layered rather than monolithic. Schwarz waveform and moving-subdomain strategies focus on localized space-time effort [9]. Spectral, pseudo-spectral, exponential Runge–Kutta, and partitioned exponential methods target high temporal fidelity when stiffness is driven mainly by linearized or semilinear dynamics [10–13]. Hybrid discretizations, including finite-volume/SPH and finite-element/meshless couplings, retain a robust global backbone while activating flexible local representations in dynamically sensitive regions [14,15]. These works motivate the present focus on local intervention, but they do not directly address contamination of a multistep IMEX history by a short moving-interface event.

Hybrid modeling and model-reduction studies show why this local question matters in present engineering workflows. Machine-learning corrections to analytical or coarse models have been used to improve high-fidelity agreement [16], simulation-guided testing uses reduced descriptions to guide measurements and computation [17], and hybrid-twin literature emphasizes the alliance of physical structure with data-driven correction [18,19]. Residual-based transfer learning has been proposed to accelerate long CFD predictions [20], while model-order-reduction surveys emphasize reliability under repeated queries [21]. Guaranteed hybrid acceleration provides another route for improving computational efficiency with explicit safeguards [22], stable reduced-model constructions address reliability in hybrid-twin settings [23], and optimal-transport hybrid twins further illustrate correction of localized weaknesses in otherwise credible solvers [24,25]. The present study uses this principle for a narrower time-integration defect: the local event window that contaminates the stored IMEX history.

1.2. Temporal Order Reduction in Stiff Time Integration

Within this broader landscape, temporal order reduction remains a central issue in stiff time integration. IMEX linear multistep schemes separate stiff and nonstiff terms efficiently, but their high-order behavior relies on regularity over the active temporal stencil. Restart-aware multirate IMEX constructions address stage-restart mechanisms [26]. Stability restrictions for IMEX linear multistep methods have been analyzed for stiff/nonstiff splittings [27]. Self-consistent and residual-balanced IMEX strategies reduce errors associated with loose coupling or inexact implicit solves [28,29]. Partitioned and IMEX general linear methods provide order-condition theory, optimized stability regions, and reaction-diffusion demonstrations [30–32]. These references support the classical time-integration background, but they do not by themselves supply an event-local bridge for moving-interface history contamination.

For the class of problems considered here, the central difficulty is localized. In interface-dominated parabolic dynamics, the solution may remain smooth for fixed interface thickness while the Eulerian trace at a fixed node undergoes a rapid transition when the moving interface crosses that node. A short defect is then not necessarily confined to that step: in a multistep IMEX scheme, it can enter the stored stencil and affect later updates. Event-driven and impulsive-dynamics literature makes an analogous point from a different perspective: a short event can determine the state from which the next regime begins, so event handling cannot be ignored merely because the event interval is small [33,34]. The present paper transfers that idea to interface-sensitive IMEX history reconstruction.

The cited IMEX advances therefore define the reference point rather than the solution of the present problem. Stage-restart IMEX-MRI methods offer restart-capable multirate structure [26]. IMEX stability analysis clarifies constraints of linear multistep splitting [27]. Self-consistent and residual-balanced IMEX methods target coupling and nonlinear-solve

errors [28,29]. IMEX general linear methods broaden the high-order design space [30–32]. This study addresses a narrower, highly localized challenge: how to isolate a small set of interface-crossing intervals, repair their local endpoint, and rebuild a clean multistep history without replacing the classical IMEX solver globally.

1.3. From Classical Discretization to Learning-Based PDE Solvers

Classical discretization methods remain indispensable for the spatial side of this problem. Phase-field formulations describe diffuse moving interfaces [3], moving-grid methods adapt to transported layers [1,2], DG/WENO methods control steep transport-dominated profiles [6–8], and hybrid discretizations combine complementary spatial representations [14,15]. Variational multiscale element-free Galerkin and local discontinuous Galerkin formulations provide another reaction-diffusion discretization context [35]. Reduced bases and tensor reduced models compress parametric parabolic dynamics while retaining validated or structure-preserving information [21,36,37]. These approaches address resolution or model compression; the present work instead addresses the local temporal history defect generated when an interface crosses an IMEX multistep stencil.

Scientific machine learning supplies a second set of tools. Error estimates for deep learning in fluid dynamics support the analysis of learned approximations [38]. Deep finite-volume methods [39], adaptive activation functions [40], PINN review literature [41], optimizer studies for physics-informed models [42], adaptive physics-informed architectures [43], mixed-residual neural solvers [44], hybrid spectral-neural reviews [45], ANN/PDE surveys [46], PINN convergence analyses [47], dual-phase or automated phase-division PINN methods [48], adaptive physics-informed networks [49], heat-transfer PINN reviews [50], and broader methodological surveys of physics-informed learning [51] illustrate the range of neural PDE solvers and their known limitations. Separately, recurrent/POD surrogates, physics-informed reduced-order models, differentiable PDE solvers, neural ODE reduced models, random-feature operator learning, and smooth neural fields accelerate repeated-query computation or nonlinear reduced modeling [52–63]. These works motivate learning-assisted PDE solvers, but they are not specifically designed to restore IMEX multistep history across a detected interface event.

The hybrid correction strategy proposed here follows the conservative interpretation of these learning-based developments. Classical solvers provide structure, interpretability, and a stable default path; the learning component is activated only where the classical history is least reliable. This is closer in spirit to hybrid correction and hybrid-twin work [16,18,19,22–25] than to a global neural replacement of the PDE solver.

The closest neural precedents are sequential-in-time neural Galerkin and related local neural trial-space methods. Neural Galerkin schemes approximate time-dependent PDE solution fields by evolving network parameters sequentially [64]. Local deep Galerkin methods apply related ideas to time-dependent PDE systems including Cahn–Hilliard-type dynamics [65]. Adaptive-basis and positional-embedding approaches enrich neural representations for high-dimensional or oscillatory evolution problems [66,67]. Modified deep Galerkin methods and adaptive sampling schemes further emphasize local residual information [68,69]. Weak neural Galerkin methods move the formulation closer to variational PDE language [70]. Physics-informed neural operators [71] and sparse-grid/Picard learning schemes for semilinear parabolic PDEs [72] provide related learning-based PDE directions. In the present method, this neural component is restricted to a local bridge on flagged interface-sensitive windows.

The subsequent analytical and numerical sections draw on standard references for the Allen–Cahn origin of diffuse-interface dynamics [73], the Lions–Magenes weak parabolic framework [74], classical IMEX PDE integrators [75,76], local residual and truncation-error

diagnostics [77,78], Allen–Cahn interface-capturing and sharp-interface context [79,80], Dirac–Frenkel variational foundations [81,82], Hamilton–Jacobi front-propagation algorithms [83], and Gårding-type elliptic estimates [84].

Table 1 summarizes representative research directions relevant to localized interface-sensitive stiff time integration and identifies how each strand informs, but does not fully address, the present restart-consistent DNG-IMEX framework.

Table 1. Representative research directions relevant to localized interface-sensitive stiff time integration.

Research Direction	Specific Contribution of Cited Work	Relevance to the Present Paper
Diffuse-interface and front propagation	Allen–Cahn boundary motion [73]; phase-field moving-interface review [3]; Hamilton–Jacobi front propagation [83].	Provides the interface-modeling context but does not address IMEX multistep-history correction.
Interface capture and sharp-interface limits	Lagrangian Allen–Cahn interface capture [79]; sharp-interface asymptotics [80].	Shows why localized interface behavior must be resolved, especially near transition layers.
Classical IMEX time integration	IMEX schemes for PDEs [75]; IMEX Runge–Kutta methods [76]; IMEX multistep stability [27]; IMEX-GLMs [30–32]; restart-capable IMEX-MRI methods [26].	Defines the classical time-integration backbone and comparison literature.
Neural Galerkin and physics-informed solvers	Sequential Neural Galerkin methods [64]; adaptive Neural Galerkin sampling [69]; weak Neural Galerkin methods [70]; physics-informed neural operators [71].	Motivates local neural trial spaces but does not provide restart-consistent IMEX history reconstruction.

The comparison clarifies the specific gap addressed in this paper. Existing interface-tracking, phase-field, DG/WENO, and moving-grid methods improve the representation of moving fronts, but they do not specify how an IMEX-BDF history should be rebuilt after a temporal crossing event. Classical IMEX, multirate, and general-linear methods provide order and stability theory for stiff splittings, but their standard estimates still require a compatible temporal history over the active stencil. Neural Galerkin, PINN, operator-learning, and reduced-order approaches offer flexible local approximations, but they do not by themselves impose restart consistency for a classical multistep formula. The proposed DNG-IMEX construction is positioned at this intersection: it repairs a flagged event interval locally and then resumes the IMEX-BDF3 march with reconstructed post-event history.

1.4. Positioning, Gap, and Contribution of the Present Work

In view of this literature, the gap addressed here is more specific than the broad coupling of neural approximation and IMEX time stepping. The surveyed work shows substantial progress in stiff-order preservation, multirate restarts, residual-balanced decomposition, reduced-order neural surrogates, neural operators, PINN-based solvers, and causal neural Galerkin-type methods. In the literature surveyed here, no prior framework was identified that combines event localization, bridge residual control, and restart-history consistency within a single global error-propagation setting for a hybrid neural-IMEX correction strategy. Likewise, although order reduction in stiff integrators is well documented, the reviewed literature did not reveal an explicit formulation of restart-history corruption after an interface or event crossing as the mechanism studied here.

This gap is important because it identifies a setting in which a mechanism-oriented hybridization is needed. Because localized interface events can contaminate the multistep history, successful intervention requires identifying the local quantities that must be controlled so that a correction introduced on a flagged interval remains compatible with the surrounding classical march. From that perspective, the present paper contributes to current research by clarifying the analytical structure that a trustworthy local correction

for stiff PDE solvers must satisfy. Such clarification is valuable because hybrid computational science increasingly needs local, interpretable corrections that can be inserted into existing numerical pipelines without abandoning the stability and structure of the classical backbone [52,55,60,62].

The contribution of the paper is organized in four parts. First, it formulates a localized DNG-IMEX correction framework for interface-sensitive multistep integration. Second, it develops corrected interface kinematics, a Sobolev-level weak transport estimate for the local bridge, a continuous weak energy estimate, and a conditional discrete propagation theorem under explicitly assumed event/restart defect control. Third, it combines an event-aligned manufactured Allen-Cahn benchmark with a manufactured-benchmark diagnostic detector, thereby separating mechanism-level evidence from practical event-localization behavior. Fourth, it identifies event localization, bridge accuracy, restart-history consistency, boundary admissibility, and computational overhead as coupled ingredients that a selective interface-aware realization must control.

1.5. Notation Summary

Table 2 summarizes the principal symbols used throughout the paper.

Table 2. Principal notation used throughout the paper.

Symbol	Meaning
$\Omega \subset \mathbb{R}^d$	Spatial computational domain.
$\partial\Omega$	Boundary of the spatial domain.
$d \in \{1, 2, 3\}$	Spatial dimension.
$\mathcal{I} = (0, T)$	Time interval of evolution.
T	Final simulation time.
$t_n = n\tau$	Discrete time level at step n .
τ	Uniform time-step size.
N	Number of time steps, with $N\tau = T$.
x	Spatial variable.
t	Continuous time variable.
$u(x, t)$	Exact or continuous PDE state.
$u_0(x)$	Initial condition.
U^n	Numerical approximation of $u(t_n)$.
$U^{(i)}$	Auxiliary internal-state notation used only for general time-integration context; the implemented backbone uses the IMEX-BDF3 multistep history.
$\mathcal{L} = \nu\Delta$	Linear stiff diffusion operator.
ν	Diffusion coefficient.
$\mathcal{B}$	Boundary operator in the strong PDE formulation.
$\mathcal{N}(u)$	Nonlinear reaction operator.
ε	Interfacial thickness or stiffness parameter.

Table 2. Cont.

Symbol	Meaning
V	Sobolev trial and test space.
$H = L^2(\Omega)$	Pivot Hilbert space.
V'	Dual space of V .
$a(\cdot, \cdot)$	Coercive bilinear form associated with diffusion.
$(\cdot, \cdot)_H$	Inner product in H .
$\langle \cdot, \cdot \rangle_{V', V}$	Duality pairing between V' and V .
ψ	Generic test function.
f	Source term in the generalized weak formulation.
$\Gamma(t)$	Moving interface or phase boundary.
$x_\Gamma(t)$	A point moving on $\Gamma(t)$.
$\dot{x}_\Gamma(t)$	Velocity of the tracked interface point.
$\varphi(x, t)$	Level-set function defining the interface.
$Q(u, t)$	State-dependent level-set function; the Eulerian level set is $\phi(x, t) = Q(u(x, t), t)$.
$n_\Gamma(x, t)$	Unit normal to the moving interface.
$S_\Gamma(x, t)$	Normal interface speed.
$w_\Gamma(x, t)$	Interface transport velocity on $\Gamma(t)$.
$w(x, t)$	Extension of w_Γ to the bulk domain.
$\nabla \cdot w$	Divergence of the bulk transport field.
n_Ω	Outward unit normal on $\partial\Omega$.
$J \subset \{0, \dots, N-1\}$	Set of discrete jump or interface-crossing indices.
$\mathcal{P}_{\text{IMEX}}$	Classical multistep IMEX evolution operator used in smooth regimes.
$\mathcal{P}_{\text{NN}}$	Localized neural transition operator used in jump regimes, followed by restart reconstruction.
$v_\theta(x, t)$	Neural bridge state on a flagged local interval. Time dependence enters through the parameter path rather than as an independent network input.
$\theta \in \mathbb{R}^P$	Vector of trainable neural parameters.
P	Number of trainable neural parameters.
$\mathcal{M}$	Neural manifold traced by the parameterized trial space.
$T_{v_\theta}\mathcal{M}$	Tangent space to the neural manifold at v_θ .
$\mathcal{R}(x, t; \theta)$	Weak residual associated with the neural bridge state on a flagged interval. It is defined through the variational formulation unless stronger regularity is imposed explicitly.
$M(\theta)$	Gram or mass matrix in the projected parameter dynamics.
$F(\theta)$	Force vector in the projected parameter dynamics, with explicit dependence on time through the source term and the projected right-hand side.

Table 2. Cont.

Symbol	Meaning
d_{n+1}	Local truncation defect at step $n + 1$.
$E_n = \ u(t_n) - U^n \ _H$	Global discrete error at time t_n .
κ_s	Smooth-regime one-step amplification constant.
κ_j	Jump-regime or restart-regime one-step amplification constant.
C_s	Smooth-regime truncation constant.
$\delta_n(\tau)$	Local event or restart defect.
A_n	Unified one-step amplification factor.
ℓ_n	Unified local defect term in the recurrence.
$\Phi_{m,n}$	Transition factor used to unroll the discrete recurrence.
$K = \max\{\kappa_s, \kappa_j\}$	Uniform amplification bound used in the global estimate.
J_M	Uniform upper bound on the number of extended flagged event/restart intervals.
α, β	Gårding constants in the coercivity estimate.
C_w	Transport constant used in Theorem 1, based on the weak L^∞ bound on w .
L_N	Nonlinear one-sided bound used in Theorem 1; selected according to Assumption 2 or Assumption 3.
$y(t) = \ v_\theta(t) \ _H^2$	Energy quantity used in the continuous weak energy estimate for the bridge state.
$\lambda = 2(\beta + C_w + L_N)$	Growth constant in the Gronwall-type bound of Theorem 1, including diffusion, weak transport, and nonlinear one-sided constants.
$\mathcal{H}^n$	Stored multistep history used by the IMEX-BDF3 switching map.
r^n	Restart-level bookkeeping variable used after a flagged correction.

2. Materials and Methods

To formulate the proposed DNG-IMEX framework, we first introduce the class of stiff semilinear parabolic problems considered in this paper, together with their strong and weak Eulerian formulations. We then describe the classical IMEX time-stepping structure used on smooth temporal intervals and explain why its nominal order may fail to be observed when a moving state-dependent interface crosses the grid. The purpose of this section is therefore twofold: to define the analytical setting and to isolate the local mechanism that motivates the later jump-correction strategy.

2.1. Strong and Weak Eulerian Formulations

Let $\Omega \subset \mathbb{R}^d$, with $d \in \{1, 2, 3\}$, be a bounded Lipschitz domain, and let $\mathcal{I}$ denote the time interval from 0 to T . The compact strong notation used for the model is

$$\partial_t u(x, t) = \mathcal{L}u(x, t) + \mathcal{N}(u(x, t)) + f(x, t), \quad (x, t) \in \Omega \times \mathcal{I}, \quad (1)$$

with initial and boundary conditions

$$u(x, 0) = u_0(x), \quad x \in \Omega, \quad (2)$$

$$\mathcal{B}u(x, t) = 0, \quad (x, t) \in \partial\Omega \times \mathcal{I}. \quad (3)$$

Here $\mathcal{L}$ denotes the stiff linear diffusion operator, $\mathcal{N}$ is the nonlinear reaction operator, f is a possible source term, and $\mathcal{B}$ is the boundary operator. Throughout the paper, the prototypical diffusion operator is

$$\mathcal{L} = \nu \Delta, \quad \nu > 0, \quad (4)$$

and, when the Allen–Cahn model is used as the guiding phase-field example, the nonlinear reaction term is taken in the cubic bistable form [73]

$$\mathcal{N}(u) = \frac{1}{\varepsilon^2} u(1 - u^2), \quad 0 < \varepsilon \ll 1. \quad (5)$$

The parameter ε governs the interfacial thickness and introduces a sharp separation between relatively smooth bulk regions and narrow transition layers. In the unforced prototype, one may set $f \equiv 0$. In the manufactured-solution tests used later, f is chosen so that a prescribed exact solution satisfies the model. Writing the problem in the compact form (1) is therefore convenient for notation, but the analytical interpretation used below is variational.

Let V be the Sobolev trial and test space associated with the boundary condition, for example, $V = H_0^1(\Omega)$ for homogeneous Dirichlet data or $V = H^1(\Omega)$ for homogeneous Neumann or compatible natural data, and let

$$V \subset H = L^2(\Omega) \cong H' \subset V' \quad (6)$$

be the associated Gelfand triple [74]. The weak problem is as follows: find $u \in L^2(0, T; V)$ with $\partial_t u \in L^2(0, T; V')$ such that, for every $\psi \in V$ and for almost every $t \in (0, T)$,

$$\langle \partial_t u(t), \psi \rangle_{V', V} + a(u(t), \psi) = \langle \mathcal{N}(u(t)), \psi \rangle_{V', V} + \langle f(t), \psi \rangle_{V', V}. \quad (7)$$

When $\mathcal{N}(u(t)) \in H$, the duality pairing in (7) is identified with the H inner product. For the model operator $\mathcal{L} = \nu \Delta$, the diffusion form is

$$a(v, \psi) = \nu(\nabla v, \nabla \psi)_H. \quad (8)$$

Thus the diffusion-reaction structure is formulated at the weak level. Any transport field introduced later is not part of the original physical PDE; it is a local correction device used around flagged interface-crossing intervals.

Nonhomogeneous Dirichlet data are handled by a standard lifting. If prescribed boundary data g_D are imposed on the Dirichlet part of the boundary, choose an extension $\ell(t)$ with trace $g_D(t)$ and write $u(t) = \ell(t) + z(t)$, where $z(t)$ belongs to the homogeneous space V_0 . The weak formulation is then applied to z with homogeneous test functions in V_0 ; numerical errors are still reported for the physical variable u . This convention is used for the exact nonhomogeneous Dirichlet data in the manufactured benchmark of Section 3.

Functional-Analytic Interpretation and Regularity Convention

The strong notation in (1) is only a compact formal representation of the model. The mathematical interpretation used in the analytical parts of the paper is the variational parabolic setting associated with the Gelfand triple (6). Under the standard Lions–Magenes

framework [74], a function $u \in L^2(0, T; V)$ with weak time derivative $\partial_t u \in L^2(0, T; V')$ has an H -continuous representative,

$$u \in C([0, T]; H). \quad (9)$$

Consequently, the initial condition $u(0) = u_0$ is understood in H , not as a pointwise classical condition. The data assumptions used in the weak analytical formulation are

$$u_0 \in H, \quad f \in L^2(0, T; V'), \quad \mathcal{N}(u) \in V' \Big| \text{on the admissible class of states.} \quad (10)$$

The diffusion form $a(\cdot, \cdot)$ is assumed continuous on $V \times V$ and satisfies a coercivity/Gårding-type lower bound: there are constants $\alpha > 0$ and $\beta \geq 0$ such that

$$a(v, v) + \beta \|v\|_H^2 \geq \alpha \|v\|_V^2, \quad v \in V. \quad (11)$$

For the Allen–Cahn cubic, the nonlinear estimates used later are invoked under a bounded-regime or truncated one-sided-growth convention. Pointwise identities, nodal traces, classical chain rules, Taylor-type residual calculations, and the strong-form notation in Equation (1) are used for compact model description and for the explicitly smooth local mechanism and manufactured-benchmark calculations.

The theoretical analysis is formulated through the variational identity (7), the duality pairing between V and V' , and the weak time-derivative framework associated with the Gelfand triple (6).

2.2. Classical IMEX Linear Multistep Background; Implemented Backbone Is IMEX-BDF3

The stiffness induced by diffusion and sharp interfacial structure makes fully explicit time integration impractical on macroscopically relevant grids, since the admissible step size may become severely restricted by the linear stiff part. On the other hand, a fully implicit treatment of the complete semilinear system generally requires the solution of nonlinear algebraic problems at every step. Implicit–explicit (IMEX) methods balance these two effects by treating the stiff linear operator implicitly and the nonlinear reaction term explicitly.

Let

$$t_n = n\tau, \quad n = 0, 1, \dots, N, \quad N\tau = T. \quad (12)$$

Denote by U^n the discrete approximation at time t_n . Because the mechanism studied in this paper is contamination of a multistep history, the relevant classical background is the IMEX linear multistep form, not an IMEX Runge–Kutta stage formula. A general k -step IMEX linear multistep discretization of the semidiscrete problem may be written as

$$\sum_{j=0}^k \alpha_j U^{n+j} = \tau \left(\sum_{j=0}^k \beta_j L_h U^{n+j} + \sum_{j=0}^{k-1} \hat{\beta}_j (N_h(U^{n+j}) + f_h^{n+j}) \right). \quad (13)$$

Here L_h denotes the spatially discrete stiff linear operator, N_h denotes the explicitly treated nonlinear part, and f_h^{n+j} denotes the discrete source contribution when present. The implicit coefficient polynomial, the implicit linear-part polynomial, and the explicit extrapolation polynomial are

$$\rho(\zeta) = \sum_{j=0}^k \alpha_j \zeta^j, \quad \sigma(\zeta) = \sum_{j=0}^k \beta_j \zeta^j, \quad \hat{\sigma}(\zeta) = \sum_{j=0}^{k-1} \hat{\beta}_j \zeta^j. \quad (14)$$

In the numerical experiments below, the implemented backbone is the third-order IMEX-BDF formula obtained by combining the BDF3 derivative with implicit treatment of $L_h U^{n+1}$ and third-order extrapolation of the explicit nonlinear term from the stored history. The concrete update used in the algorithm is stated in Section 2.4.4, together with the BDF1-to-BDF2 restart reconstruction used after accepted flagged corrections.

The notation in (13) and (14) fixes the linear multistep interpretation used throughout the rest of the paper. Thus, $\mathcal{P}_{\text{IMEX}}$ is a history-dependent IMEX-BDF operator acting on stored states, not a one-step Runge–Kutta map. The later error analysis separates ordinary smooth-regime steps from flagged event/restart steps precisely because the multistep stencil must remain compatible with a single local temporal regime.

2.3. The Mechanics of Temporal Order Reduction

Let

$$e_n = u(t_n) - U^n \quad (15)$$

denote the global discrete error at time t_n . The evolution of e_n is driven by the one-step defect obtained by inserting the exact solution into the numerical scheme. For a p -th order method applied on a temporally smooth interval, this local truncation defect, denoted here by d_{n+1} , satisfies a bound of the form

$$\|d_{n+1}\|_H \leq C \tau^{p+1} \sup_{t \in [t_n, t_{n+1}]} \|\partial_t^{p+1} u(t)\|_H, \quad (16)$$

where C is independent of τ on that interval.

Estimate (16) is the standard smooth-regime local truncation-defect statement used in time-discretization analysis [77,78]. The validity of this estimate rests on the assumption that the exact solution possesses sufficient temporal regularity and that the constant in the remainder estimate remains controlled on the time slab $[t_n, t_{n+1}]$.

For interface-dominated phase-field problems, this assumption can fail in an effective numerical sense even when the solution is classically smooth for fixed $\varepsilon > 0$. Near a moving interface, the solution typically admits a local profile of the form

$$u(x, t) \approx q\left(\frac{d(x, t)}{\varepsilon}\right), \quad (17)$$

where $d(x, t)$ is a signed distance-like coordinate relative to the interface, and q is the transition profile. Representation (17) is intended as a local moving-interface ansatz consistent with standard diffuse-interface, sharp-interface, and interface-capturing descriptions of interfacial motion [3,79,80].

If a fixed Eulerian node x_i lies near the crossing point and the interface passes that node around a time t^* , then the nodal trace $t \mapsto u(x_i, t)$ behaves locally like

$$u(x_i, t) \approx q\left(\frac{\sigma_i - c(t - t^*)}{\varepsilon}\right), \quad (18)$$

where c is the local crossing speed and σ_i measures the offset of the node relative to the interface at the reference time.

Differentiating (18) with respect to time gives the following local derivative bound on the crossing window under the smooth-profile assumption:

$$|\partial_t^m u(x_i, t)| \leq C_m \left(\frac{|c|}{\varepsilon}\right)^m, \quad m \geq 1. \quad (19)$$

Thus, although no literal singularity is present for fixed $\varepsilon > 0$, the constant hidden in the Taylor remainder is generally not uniform in ε . On time intervals whose width is not asymptotically smaller than the crossing scale Δt_{cross} , the nodal trace undergoes a rapid transition that is poorly represented by a single smooth Taylor expansion around one endpoint. In this regime, the formal bound (16) ceases to be informative as a uniform macroscopic estimate.

The temporal scale of this localized transition follows directly from the interface width and speed. If the interface location $\gamma(t)$ satisfies $\gamma(t_i) = x_i$ and $\gamma'(t_i) = c$, then the argument of the diffuse profile remains $O(1)$ precisely when the interface moves one diffuse thickness. Hence the crossing duration satisfies

$$\Delta t_{\text{cross}} = O\left(\frac{\varepsilon}{|c|}\right). \quad (20)$$

This is the scale used in the numerical interpretation: decreasing ε makes the same Eulerian nodal transition sharper in time and therefore more demanding for a fixed multistep stencil.

The local mechanism described above is illustrated schematically in Figure 1. The figure emphasizes the Eulerian viewpoint relevant to the time integrator: as a moving interface passes a fixed grid node, the nodal trace in time undergoes a sharp transition on the scale $O(\varepsilon/|c|)$. On such intervals, a single smooth-step Taylor expansion across the whole time slab may no longer provide an effective representation of the local temporal behavior, even though the solution remains classically smooth for fixed ε .




Figure 1. Schematic illustration of temporal order reduction near an interface-crossing interval. (a) Spatial profiles of the moving diffuse interface at selected times as it passes a fixed Eulerian node x_i . (b) Corresponding nodal trace $t \mapsto u(x_i, t)$, showing a rapid transition near the event time t_* on the interface-transition scale $O(\varepsilon/|c|)$. The symbol t_* denotes the interface-crossing time. In this regime, the local temporal behavior is poorly represented by a single smooth Taylor expansion across the full step, which motivates the introduction of flagged interface-crossing intervals in the DNG-IMEX framework.

Figure 1 summarizes the mechanism that motivates the later correction strategy. The difficulty is the loss of effective smooth-step behavior when a nodal trace crosses two local regimes over one multistep stencil. The resulting defect may no longer satisfy a uniform smooth-regime estimate, which motivates the distinction between ordinary smooth intervals and flagged interface-crossing intervals. This subsection provides the mechanism-level explanation; the pivot Hilbert space and its corresponding norm are defined in the notation table and used in the continuous and discrete estimates below.

To make this precise at the algorithmic level, we denote by

$$J \subset \{0, 1, \dots, N-1\} \quad (21)$$

the set of indices for which the interval is flagged as a jump or interface-crossing interval. At this stage, the flagged set is an abstract object. It may be supplied by exact benchmark information, by an oracle event indicator, or by an independently validated a posteriori diagnostic detector. The numerical study below uses exact manufactured-benchmark information to isolate the local correction mechanism. Outside this set, the classical IMEX update is retained; inside it, the smooth-step update is replaced by a localized correction. The later sections develop this separation through a neural bridge model and a discrete propagation analysis that distinguish smooth-step defects from jump-step defects.

2.4. The Deep Neural Galerkin-IMEX (DNG-IMEX) Framework

To address the localized loss of effective temporal regularity described in Section 2.3, we introduce a Deep Neural Galerkin-IMEX (DNG-IMEX) framework with two regimes: the classical IMEX march on unflagged intervals and a local neural bridge on flagged interface-crossing intervals.

The purpose of this section is organized around three steps. First, we define the switching logic that determines when the classical integrator is suspended and when the neural corrector is activated. Second, we introduce the neural manifold on which the local bridge evolves. Third, we derive the projected parameter dynamics by means of a Dirac–Frenkel variational principle posed on that manifold.

2.4.1. Hybrid Architecture and Transition Logic

Let the time interval $[0, T]$ be partitioned by the uniform grid

$$0 = t_0 < t_1 < \dots < t_N = T, t_n = n\tau, N\tau = T, \quad (22)$$

and let

$$U^n \approx u(t_n) \quad (23)$$

denote the global discrete approximation at time t_n .

As in Section 2.3, let

$$J \subset \{0, 1, \dots, N-1\} \quad (24)$$

be the set of flagged jump or interface-crossing indices. An index $n \in J$ indicates that the interval $[t_n, t_{n+1}]$ contains a detected crossing event for which the standard smooth-step IMEX approximation is considered unreliable. The precise detector may depend on the application and on the chosen interface level, but the numerical role of J is always the same: it identifies those intervals on which the classical IMEX step is replaced by a local neural transition.

Accordingly, the global state update is defined by the piecewise evolution rule

$$\mathcal{H}^n = (U^n, U^{n-1}, U^{n-2}, r^n) \quad (25)$$

$$(U^{n+1}, \mathcal{H}^{n+1}) = \begin{cases} \mathcal{P}_{\text{IMEX}}(\mathcal{H}^n, \tau), & n \notin J, \\ \mathcal{P}_{\text{NN}}(\mathcal{H}^n, \tau), & n \in J, \end{cases} \quad (26)$$

Here, $\mathcal{H}^n$ denotes the stored multistep history, and r^n records the restart level. The operator $\mathcal{P}_{\text{IMEX}}$ is therefore a multistep IMEX-BDF operator acting on the stored history, not a one-step map in U^n alone. The operator $\mathcal{P}_{\text{NN}}$ denotes the localized neural transition on a flagged interval, followed by the restart-history reconstruction described in Section 2.4.4.

To summarize the switching mechanism and the role of the local neural bridge, Figure 2 presents the architecture of the DNG-IMEX framework. The figure emphasizes that the classical IMEX solver remains the default backbone integrator, while the neural correction is activated only on flagged interface-crossing intervals $n \in J$. It also clarifies that on such intervals, the bridge evolution is performed on a neural manifold through the parameter path $t \mapsto \theta(t)$, after which the corrected terminal state is returned to the grid, and the classical march is resumed.




Figure 2. Architecture of the hybrid DNG-IMEX framework. The global time interval is partitioned into smooth-regime steps and flagged interface-crossing steps. On intervals $n \notin J$, the numerical state U^n is advanced by the classical IMEX operator $\mathcal{P}_{\text{IMEX}}$. On intervals $n \in J$, the discrete state is embedded into the neural manifold $\mathcal{M} = \{v_\theta : \theta \in \mathbb{R}^P\}$, evolved locally over $[t_n, t_{n+1}]$ through the projected parameter dynamics $M(\theta)\dot{\theta} = F(\theta, t)$, and sampled back at t_{n+1} to define the corrected state U^{n+1} .

Figure 2 fixes the switching logic used throughout the analysis. A flagged interval is advanced by the local bridge, and the corrected terminal value is then returned to the global IMEX-BDF history. This construction allows the later analysis to separate smooth-step propagation from jump-step correction. The detector-specific choices used in the validation are described in Section 2.4.4 and assessed numerically in Section 3.

2.4.2. Neural Manifold and Local Trial Space

Fix a flagged index $n \in J$ and consider the local interval

$$\mathcal{I}_n = [t_n, t_{n+1}]. \quad (27)$$

On this interval, the bridge state is represented by a neural ansatz

$$v_{\theta(t)}(x), \quad x \in \Omega, \quad t \in \mathcal{I}_n, \quad (28)$$

where

$$\theta(t) = (\theta_1(t), \theta_2(t), \dots, \theta_P(t))^T \in \mathbb{R}^P \quad (29)$$

is a time-dependent parameter vector. For each fixed parameter value θ , the network defines a spatial trial function

$$v_\theta : \Omega \rightarrow \mathbb{R}, \quad (30)$$

and the time dependence of the bridge is carried through the parameter path $t \mapsto \theta(t)$.

The associated neural manifold is

$$\mathcal{M} = \{v_\theta \in V : \theta \in \mathbb{R}^P\}. \quad (31)$$

For the weak diffusion projection used below, the parameter derivatives are assumed to be admissible test functions. Thus, at a parameter value θ , the tangent space is understood as

$$T_{v_\theta}\mathcal{M} = \text{span}\{\partial_{\theta_1}v_\theta, \dots, \partial_{\theta_P}v_\theta\} \subset V. \quad (32)$$

This condition is the functional-analytic requirement needed for the displayed weak Galerkin projection. If a lower-regularity architecture is used in a purely numerical implementation, then the diffusion contribution must be interpreted through the corresponding regularized or residual-minimization discretization rather than through the weak gradient formula written below.

A smooth boundary-admissible tanh network is compatible with the weak Sobolev formulation because it defines a particular admissible element of the trial space. The weak framework supplies the minimal variational setting for the bridge estimate and residual testing, while the smooth parametrization provides a convenient finite-dimensional trial manifold for the local bridge.

The bridge residual is used only in weak variational form. For $\psi \in V$, define the local residual functional $R_\theta(t) \in V'$ by

$$\langle R_\theta(t), \psi \rangle_{V',V} = \langle \partial_t v_{\theta(t)}, \psi \rangle_{V',V} + a(v_{\theta(t)}, \psi) - \langle \mathcal{N}(v_{\theta(t)}) + f(t), \psi \rangle_{V',V}. \quad (33)$$

This definition is used only through duality pairings and weak residual tests; no pointwise evaluation of Δv_θ is required.

At the beginning of the flagged interval, the bridge is initialized from the discrete state U^n . In practice, this means that $\theta(t_n)$ is chosen so that $v_{\theta(t_n)}$ matches the discrete state at the Eulerian nodes or approximates it in a prescribed projection or least-squares sense. The terminal profile $v_{\theta(t_{n+1})}$ is then sampled back onto the grid to define the corrected state.

The boundary treatment of the local bridge follows the lifting convention used in the weak formulation. In the continuous weak Dirac–Frenkel projection, the admissible neural manifold is the boundary-conforming manifold obtained through a differentiable constrained ansatz, a fixed lifting, or a projection into the homogeneous Sobolev space. Consequently, the bridge state and its parameter derivatives belong to the variational trial and test space, and tangent directions are admissible test functions. For nonhomogeneous Dirichlet data, the bridge evolves the homogeneous lifted variable, and the prescribed lifting is added back when the corrected endpoint is sampled onto the physical grid. Theorem 1 is formulated for this admissible lifted, projected, or constrained bridge variable.

With this convention, a projection used in the continuous variational interpretation is built into the parameter-to-Sobolev-space map. The tangent functions in (32) are the derivatives of this admissible map, which ensures that they remain valid test functions in the weak Dirac–Frenkel equations below.

2.4.3. Weak Dirac–Frenkel Projection on a Flagged Interval

The neural bridge on a flagged interval is determined by projecting the weak PDE dynamics onto the tangent space of the neural manifold. The term “Neural Galerkin” is used here in this local trial-space sense: the bridge state is restricted to $\mathcal{M}$, and the residual is tested only against tangent directions in $T_{v_\theta}\mathcal{M}$.

Let the tangent functions $\partial_{\theta_i}v_\theta$ belong to V . The weak projected equations are

$$\left(\partial_t v_{\theta(t)}, \partial_{\theta_i} v_{\theta(t)}\right)_H + a\left(v_{\theta(t)}, \partial_{\theta_i} v_{\theta(t)}\right) = \langle \mathcal{N}(v_{\theta(t)}) + f(t), \partial_{\theta_i} v_{\theta(t)} \rangle_{V', V}, \quad i = 1, \dots, P. \quad (34)$$

Because time enters through the parameter path, the chain rule along the finite-dimensional parameter curve gives

$$\partial_t v_{\theta(t)} = \sum_{j=1}^P \dot{\theta}_j(t) \partial_{\theta_j} v_{\theta(t)}. \quad (35)$$

Substitution of (35) into (34) gives

$$\sum_{j=1}^P M_{ij}(\theta(t)) \dot{\theta}_j(t) = F_i(\theta(t), t), \quad i = 1, \dots, P, \quad (36)$$

or, in matrix form,

$$M(\theta(t)) \dot{\theta}(t) = F(\theta(t), t), \quad t \in \mathcal{I}_n. \quad (37)$$

Here

$$M_{ij}(\theta) = \left(\partial_{\theta_j} v_{\theta}, \partial_{\theta_i} v_{\theta}\right)_H, \quad i, j = 1, \dots, P, \quad (38)$$

and

$$F_i(\theta, t) = -a(v_{\theta}, \partial_{\theta_i} v_{\theta}) + \langle \mathcal{N}(v_{\theta}) + f(t), \partial_{\theta_i} v_{\theta} \rangle_{V', V}. \quad (39)$$

For the prototype diffusion operator, this is

$$F_i(\theta, t) = -\nu(\nabla v_{\theta}, \nabla \partial_{\theta_i} v_{\theta})_H + \langle \mathcal{N}(v_{\theta}) + f(t), \partial_{\theta_i} v_{\theta} \rangle_{V', V}. \quad (40)$$

Equations (34)–(40) are weak equations. They do not require pointwise evaluation of Δv_{θ} and do not assume that the bridge state is $C^1(\Omega)$ or $C^2(\Omega)$. The Gram matrix is symmetric positive semidefinite by construction and is positive definite when the tangent vectors are linearly independent in H . In practice, near-degeneracy can be handled by regularization or residual minimization.

The finite-dimensional weak Dirac–Frenkel system (34)–(40) enforces residual orthogonality on the tangent space of the admissible neural manifold. The full weak bridge equation used in Theorem 1 is posed for every test function in V . Residual components outside the tangent space are therefore represented through the source/residual term g or through the flagged-step defect hypotheses used in Theorem 2.

The projected system defines the local bridge used on flagged intervals. The numerical realization of this abstract construction, including the IMEX-BDF3 backbone, restart reconstruction, and bridge hyperparameters, is specified next.

2.4.4. Numerical Realization, Restart Reconstruction, and Reproducibility Parameters

The numerical realization used in Section 3 is deliberately more specific than the abstract DNG-IMEX description. The classical backbone is IMEX-BDF3. When the detector flags an event interval, the code constructs local corrected candidates and accepts a candidate only if the acceptance layer does not increase the monitored local residual or reference-distance quantities. If no corrected candidate satisfies the acceptance check, the method falls back to the IMEX candidate with restart bookkeeping.

For the manufactured-benchmark diagnostic run, the local corrected endpoint candidates are the uncorrected IMEX endpoint, a neural endpoint, a subcycle endpoint, and a transported-front endpoint. The subcycle endpoint is obtained by resolving $[t_n, t_{n+1}]$ with first-order IMEX micro-steps of size $\delta t = \tau/M$, where M is the micro-step budget. The transported-front endpoint is obtained by fitting a single tanh front $a \tanh(\beta(x - c)) + d$ to U^n and advecting its fitted center with the forcing-derived interface speed. The neural

endpoint is the one-hidden-layer tanh bridge evolved over the flagged window. The implementation tests corrected candidates in the fallback order neural $\rightarrow$ subcycle $\rightarrow$ front $\rightarrow$ IMEX. A corrected candidate is accepted only when the residual and reference-distance acceptance ratios improve according to the configured tolerance, 0.98 in the reported runs, and the norm-growth guard remains satisfied. The “front” label is therefore an implementation fallback tag for the manufactured benchmark, not an additional assumption in the theoretical propagation theorem.

On unflagged intervals with a clean third-order history, the fully specified backbone used in the computations is the IMEX-BDF3 update

$$\frac{11U^{n+1} - 18U^n + 9U^{n-1} - 2U^{n-2}}{6\Delta t} = L_h U^{n+1} + 3N_h(U^n) - 3N_h(U^{n-1}) + N_h(U^{n-2}) + f_h^{n+1}. \quad (41)$$

Here, L_h is the implicit linear spatial discretization, N_h is the explicit nonlinear discretization, and the manufactured forcing is evaluated at the new time level. The restart formulas below specify how the BDF3 memory is repopulated after an accepted flagged correction.

Restart reconstruction is implemented explicitly. After an accepted local correction, the corrected endpoint U_{n+1}^c is inserted into the multistep history. The first post-event step is advanced with a one-step BDF1 reconstruction, the next clean step with BDF2, and only after two clean post-event states are available does the method return to IMEX-BDF3. This BDF1-to-BDF2-to-BDF3 reconstruction is essential because a third-order multistep formula cannot immediately be resumed from a history containing pre-event states.

No pre-event solution values are reused in the restored third-order stencil. The lower-order BDF1 and BDF2 steps are used only as post-event reinitialization steps, so the return to IMEX-BDF3 occurs only after the multistep memory has been repopulated by post-correction states. This is the practical meaning of restart-history consistency in the numerical tests.

The first post-event update uses the consistent IMEX-BDF1 reconstruction

$$\frac{U^{n+2} - U^{n+1,c}}{\Delta t} = L_h U^{n+2} + N_h(U^{n+1,c}) + f_h^{n+2}. \quad (42)$$

The next update uses the IMEX-BDF2 reconstruction

$$\frac{3U^{n+3} - 4U^{n+2} + U^{n+1,c}}{2\Delta t} = L_h U^{n+3} + 2N_h(U^{n+2}) - N_h(U^{n+1,c}) + f_h^{n+3}. \quad (43)$$

After these two clean states are available, the IMEX-BDF3 stencil is restored. These formulas specify the order-reduced restart used after a flagged correction.

The benchmark diagnostic detector used in the validation is event-local and benchmark-informed in the narrow manufactured-solution sense. It combines the jump-time and speed-change information available in the manufactured benchmark with one-step padding around the flagged interval. This benchmark-informed choice isolates the local bridge/restart mechanism under accurately localized events. For applications with unknown events, the same framework would require a residual-based, geometric, or a posteriori event indicator. A broad single-state geometric trigger is retained only as a diagnostic comparison because it can activate almost everywhere and thereby destroy locality. The resulting detector metrics are reported in Section 3.

The implemented neural bridge is also fixed explicitly for reproducibility. The reported validation uses a scalar-input/scalar-output, one-hidden-layer fully connected tanh network for the local spatial bridge, with hidden width 3 in the main tests and width 5 in the robustness check. The bridge is initialized by matching or projecting the pre-event grid state at the Eulerian nodes, advanced only on flagged event windows, and accepted only when the

local acceptance check improves the monitored residual or reference-distance quantity. These details are summarized in Table 3 and reproduced in the public code repository.

Table 3. Numerical realization, neural-bridge architecture, restart reconstruction, and reproducibility parameters used in the validation.

Quantity	Value Used in Validation	Purpose
Backbone	IMEX-BDF3	Default solver on smooth intervals
Grid and final time	$n_x = 1025$, $\Omega = [-1, 1]$, $T = 0.8$, $\nu = 1$, $x_\Gamma(0) = -0.25$	Standard benchmark resolution and parameters used by the reported validation
Interface widths	$\varepsilon \in \{0.04, 0.03, 0.02\}$	Sensitivity and stress tests
Time steps	$\Delta t \in \{8 \times 10^{-4}, 4 \times 10^{-4}, 2 \times 10^{-4}, 10^{-4}\}$; refined 5×10^{-5} and 2.5×10^{-5} for $\varepsilon = 0.02$	Main convergence and sharp-interface checks
Neural bridge	one-hidden-layer fully connected tanh network; scalar spatial input x and scalar output $v_\theta(x)$; hidden width 3; robustness check width 5	Local trial space and explicit architecture on flagged intervals
Bridge initialization	projection/least-squares matching of U^n at the Eulerian nodes at the start of a flagged interval; no global offline training	Reproducible start of the local bridge
Local solver budget	16 micro-steps and 64 reference micro-steps; robust check 32/128	Bridge/subcycle candidate construction
Residual sampling	31 blend points	Acceptance and residual diagnostics
Acceptance and fallback	fallback chain neural $\rightarrow$ subcycle $\rightarrow$ front $\rightarrow$ IMEX; residual/reference-distance acceptance check, norm-growth guard, and exact-distance audit	Controls local correction, records fallback tags, and avoids unconditional neural activation
Restart recovery	BDF1 $\rightarrow$ BDF2 $\rightarrow$ IMEX-BDF3	Clean multistep-history reconstruction

The same lifting and boundary convention is used when constructing and evaluating the corrected candidates in the benchmark. The raw neural network output is treated as an intermediate parametrization, while candidates are interpreted on the lifted homogeneous variable and sampled or projected with the common boundary treatment before the residual, reference-distance, and norm-growth audits are applied. In the reported finite-difference code, exact Dirichlet boundary values are inserted at the grid-sampling stage, representing the lifted-variable convention at the discrete level. All local candidates are therefore compared after enforcing the same boundary convention, so differences among the neural, subcycle, transported-front, and fallback candidates reflect the local correction mechanism rather than inconsistent boundary handling.

Equivalently, the implemented bridge can be represented as a scalar input layer, one affine hidden layer, a tanh activation, and a scalar affine output layer. The trainable parameters consist only of the first-layer weights and biases and the output-layer weights and bias. Time evolution on a flagged interval is represented by updating the parameter path during the local bridge solve; the implementation therefore uses no separate time input, global offline training set, or learned full-solution surrogate. This makes the numerical realization reproducible and clearly separates the local bridge from a global PINN or neural-operator replacement. For clarity, the numerical choices used in the validation are collected in Table 3.

2.5. Theoretical Analysis

Section 2.5 develops the analytical structure associated with the local bridge. Section 2.5.1 defines the interface-induced transport field under the additional smoothness assumptions needed for classical level-set kinematics. Sections 2.5.2 and 2.5.3 then move to the Sobolev level, using a weak transport estimate and a continuous weak energy estimate for an idealized bridge model. Theorem 1 is formulated through these weak estimates, while Section 2.5.4 gives a conditional discrete propagation statement showing how smooth-step defects and event/restart defects combine under explicit one-step hypotheses. Thus, the continuous analysis identifies the structure required of a local correction, and the discrete analysis clarifies the defect control needed for a macroscopic order-preservation statement.

The continuous weak estimate is an a priori estimate for an admissible weak trajectory satisfying the stated full variational bridge equation. The implemented detector/bridge/restart pipeline is connected to the global analysis through the separate flagged-step defect bounds stated in the discrete propagation theorem.

2.5.1. Exact Interface Kinematics and the Associated Normal Transport Field

The interface-kinematic calculation in this subsection is used in the smooth setting supplied by the manufactured benchmark or formal mechanism analysis, where a differentiable level-set description and differentiable nodal traces are available on the interval under consideration.

We begin with the geometry of the moving interface. The aim is to define the normal transport field in a way that makes the dependence on the independent variables (x, t) explicit and avoids any hidden chain-rule ambiguity.

Let $\Omega \subset \mathbb{R}^d$ be an open set, let $T > 0$, and let $\Gamma(t) \subset \Omega$ be a moving interface for $t \in [0, T]$. Assume that there exists a level-set function [83]

$$\varphi : \Omega \times [0, T] \rightarrow \mathbb{R} \quad (44)$$

of class C^1 in a neighborhood of the interface such that

$$\Gamma(t) = \{x \in \Omega : \varphi(x, t) = 0\}, t \in [0, T], \quad (45)$$

and that the nondegeneracy condition

$$\nabla \varphi(x, t) \neq 0, x \in \Gamma(t), \quad (46)$$

holds. This guarantees that the interface is regular and that a unit normal can be defined pointwise on $\Gamma(t)$.

In some applications the interface is originally specified in a state-dependent form

$$Q(u(x, t), t) = 0. \quad (47)$$

In that case, the correct Eulerian level-set function is the composite quantity

$$\varphi(x, t) := Q(u(x, t), t). \quad (48)$$

All differentiation is then performed on $\varphi(x, t)$ as a function of the independent variables (x, t) . This is the precise way to account for the dependence of the interface on the evolving state.

Let $x_\Gamma(t) \in \Gamma(t)$ be a C^1 curve that remains on the interface for $t \in [0, T]$. Since $x_\Gamma(t)$ stays on the zero level set, one has

$$\varphi(x_\Gamma(t), t) = 0, t \in [0, T]. \quad (49)$$

Differentiating with respect to time and using the multivariable chain rule yields

$$\frac{d}{dt}\varphi(x_\Gamma(t), t) = \partial_t \varphi(x_\Gamma(t), t) + \nabla \varphi(x_\Gamma(t), t) \cdot \dot{x}_\Gamma(t). \quad (50)$$

Because the left-hand side is the derivative of the constant zero function, it follows that

$$0 = \partial_t \varphi(x_\Gamma(t), t) + \nabla \varphi(x_\Gamma(t), t) \cdot \dot{x}_\Gamma(t). \quad (51)$$

Define the unit normal to the interface by

$$n_\Gamma(x, t) = \frac{\nabla \varphi(x, t)}{|\nabla \varphi(x, t)|}, x \in \Gamma(t). \quad (52)$$

Equivalently,

$$\nabla \varphi(x, t) = |\nabla \varphi(x, t)| n_\Gamma(x, t), x \in \Gamma(t). \quad (53)$$

Substituting this identity into the kinematic relation gives

$$0 = \partial_t \varphi(x_\Gamma(t), t) + |\nabla \varphi(x_\Gamma(t), t)| n_\Gamma(x_\Gamma(t), t) \cdot \dot{x}_\Gamma(t). \quad (54)$$

The signed normal speed of the interface is therefore defined by

$$S_\Gamma(x_\Gamma(t), t) := \dot{x}_\Gamma(t) \cdot n_\Gamma(x_\Gamma(t), t). \quad (55)$$

Then (54) becomes

$$0 = \partial_t \varphi(x_\Gamma(t), t) + |\nabla \varphi(x_\Gamma(t), t)| S_\Gamma(x_\Gamma(t), t), \quad (56)$$

and hence

$$S_\Gamma(x, t) = -\frac{\partial_t \varphi(x, t)}{|\nabla \varphi(x, t)|}, x \in \Gamma(t). \quad (57)$$

Only the normal component of the motion is determined by the level-set identity. Indeed, the velocity of a tracked interface point may be decomposed as

$$\dot{x}_\Gamma(t) = (\dot{x}_\Gamma(t) \cdot n_\Gamma(x_\Gamma(t), t)) n_\Gamma(x_\Gamma(t), t) + \tau_\Gamma(t), \quad (58)$$

where $\tau_\Gamma(t)$ is tangent to $\Gamma(t)$. Since $\nabla \varphi$ is normal to the interface, the tangential component satisfies

$$\nabla \varphi(x_\Gamma(t), t) \cdot \tau_\Gamma(t) = 0. \quad (59)$$

Thus the level-set relation detects only the normal motion. Tangential sliding along $\Gamma(t)$ does not change the zero level set and is therefore invisible to the kinematic equation.

The canonical normal transport field attached to the interface is then

$$w_\Gamma(x, t) := S_\Gamma(x, t) n_\Gamma(x, t), x \in \Gamma(t), \quad (60)$$

so that, by substituting (52) and (57) into (60),

$$w_\Gamma(x, t) = -\frac{\partial_t \varphi(x, t)}{|\nabla \varphi(x, t)|^2} \nabla \varphi(x, t), x \in \Gamma(t). \quad (61)$$

Formula (61) determines the transport field only on the interface itself. The later bridge analysis therefore prescribes a domain-wide bounded extension

$$w : \Omega \times [0, T] \rightarrow \mathbb{R}^d \quad (62)$$

satisfying

$$w(x, t) = w_\Gamma(x, t), x \in \Gamma(t). \quad (63)$$

Such an extension is not unique, because the interface geometry fixes only the normal component on $\Gamma(t)$. Subsequent weak transport and energy estimates are therefore formulated conditionally on the prescribed bounded bulk extension w .

Assumption 1 (bounded bulk extension). *On each bridge interval $I_n = (t_n, t_{n+1})$, a bounded bulk transport extension w is prescribed with $w \in L^\infty(I_n; L^\infty(\Omega)^d)$. Wherever the smooth interface description is available, the trace or kinematic restriction of w on $\Gamma(t)$ agrees with the interface normal field w_Γ . This extension is treated as part of the idealized local bridge data. The interface kinematics define the normal field on the interface, while the weak transport and energy estimates are formulated conditionally on a prescribed bounded bulk extension.*

2.5.2. Weak Treatment of the Transport Contribution

We next state the transport estimate used in the local bridge energy argument under Assumption 1. This subsection is written at the Sobolev level on a fixed local bridge interval. The transport field is a prescribed bounded bulk extension, and the bridge state is tested through weak first derivatives. The resulting estimate uses the boundedness of the transport field and the weak gradient of the bridge state, avoiding any need for classical pointwise differentiability or boundary trace identities.

Let

$$V \subset H = L^2(\Omega) \cong H' \subset V' \quad (64)$$

be the weak setting introduced in Section 2.1. Under Assumption 1, on the local bridge interval, the chosen transport extension is essentially bounded in space and time, as specified by

$$w \in L^\infty(\Omega)^d. \quad (65)$$

For $v, \psi \in V$, define the weak transport form

$$b_w(v, \psi) = \int_{\Omega} (w \cdot \nabla v) \psi \, dx. \quad (66)$$

This expression is well-defined in the weak sense because the transport field is bounded, the weak gradient of v is square-integrable, and ψ belongs to the pivot Hilbert space. Thus the form requires only $v \in V$ and $\psi \in H$; no $C^1(\Omega)$ -regularity of v is required.

Lemma 1. *Weak transport estimate. For every v in V ,*

$$|b_w(v, v)| \leq \|w\|_{L^\infty(\Omega)} \|\nabla v\|_H \|v\|_H. \quad (67)$$

If the V -norm controls the weak gradient, then for every $\eta > 0$, one obtains the corresponding Young-inequality estimate

$$|b_w(v, v)| \leq \frac{\alpha}{4} \|v\|_V^2 + C_w \|v\|_H^2, \quad C_w = \frac{c_V^2 \|w\|_{L^\infty(\Omega)}^2}{\alpha}. \quad (68)$$

Proof. The estimate follows directly from Hölder's inequality and Cauchy–Schwarz:

$$|b_w(v, v)| \leq \|w\|_{L^\infty(\Omega)} \|\nabla v\|_H \|v\|_H. \quad (69)$$

Using the continuous control of the weak gradient by the V -norm and Young's inequality gives (68). This proof uses only weak first derivatives in L^2 and therefore remains valid for $v \in V$. It does not require v to have a classical continuous derivative and does not use any integration-by-parts formula. $\square$

Consequently, the weak theory is based on Sobolev-level estimates. In particular, Theorem 1 relies only on the estimates (67)–(68). The analysis does not require $\operatorname{div} w$, a boundary trace formula for the bridge state, or any classical C^1 spatial derivative of the bridge state.

2.5.3. Continuous Weak Energy Estimate for the Idealized Local Bridge

We now derive a continuous weak energy estimate for the idealized transport-augmented local bridge model. The estimate is formulated at the weak level on a fixed bridge interval $I_n = (t_n, t_{n+1})$. The forcing term is estimated through the dual pairing, the test function is chosen in the admissible variational space, and the nonlinear term is controlled by a one-sided weak bound or by an explicitly assumed bounded-regime one-sided bound. The estimate is formulated entirely in the weak variational setting.

Let the neural bridge state satisfy

$$v_\theta \in L^2(t_n, t_{n+1}; V), \quad \partial_t v_\theta \in L^2(t_n, t_{n+1}; V'). \quad (70)$$

Then the map $t \mapsto \|v_\theta(t)\|_H^2$ is absolutely continuous, and the Lions–Magenes identity [74] holds for almost every t :

$$\frac{1}{2} \frac{d}{dt} \|v_\theta(t)\|_H^2 = \langle \partial_t v_\theta(t), v_\theta(t) \rangle_{V', V}. \quad (71)$$

Assume that for almost every t in the local bridge interval, the bridge satisfies the weak equation for every $\psi \in V$:

$$\langle \partial_t v_\theta, \psi \rangle_{V', V} + a(v_\theta, \psi) + b_w(v_\theta, \psi) = \langle \mathcal{N}(v_\theta), \psi \rangle_{V', V} + \langle g(t), \psi \rangle_{V', V}, \quad \psi \in V. \quad (72)$$

Here g denotes the forcing or local residual source used in the bridge model, interpreted as an element of V' . The assumptions on the transport extension and the source are

$$w \in L^\infty(\Omega)^d, \quad g \in L^2(t_n, t_{n+1}; V'). \quad (73)$$

For the diffusion form, assume the Gårding inequality

$$a(v, v) + \beta \|v\|_H^2 \geq \alpha \|v\|_V^2, \quad \alpha > 0, \beta \geq 0, v \in V. \quad (74)$$

For the nonlinearity, assume that one of the following two alternatives holds.

Assumption 2 (global weak one-sided bound). *There exists a constant L_N such that*

$$\langle \mathcal{N}(v), v \rangle_{V', V} \leq L_N \|v\|_H^2, \quad v \in V. \quad (75)$$

Assumption 3 (bounded-regime setting). *As an explicit admissibility hypothesis, the bridge trajectory remains in a bounded admissible set B_R in which the nonlinear term is a well-defined V' contribution and satisfies the one-sided bound*

$$\langle \mathcal{N}(v), v \rangle_{V', V} \leq L_N \|v\|_H^2. \quad (76)$$

Assumptions 1–3 are structural admissibility hypotheses for the idealized local bridge analysis, not automatic consequences of the interface kinematics or of an arbitrary neural implementation. Assumption 1 prescribes one bounded bulk extension of the interface normal field; the extension is not unique and is treated as part of the bridge data. Assumptions 2 and 3 provide alternative nonlinear controls: for the Allen–Cahn cubic used in the benchmark, the required one-sided estimate follows from the sign structure of the cubic term whenever the pairing is meaningful, while 3 is retained for more general locally controlled nonlinearities. The fully discrete detector/bridge/restart implementation is connected to these continuous assumptions through the flagged-step defect conditions used in Theorem 2.

For the Allen–Cahn cubic

$$\mathcal{N}(v) = \frac{1}{\varepsilon^2} v(1 - v^2), \quad (77)$$

For the sign convention used in this manuscript, pairing the Allen–Cahn cubic (77) with v gives a positive quadratic contribution and a nonpositive quartic contribution.

Thus the weak one-sided bound in Assumption 2 holds with $L_N = \varepsilon^{-2}$, whenever the cubic term is meaningful as a V' contribution.

This is a one-sided growth estimate, not a global Lipschitz assertion for the full cubic. The bounded-regime or truncation convention is retained only for more general nonlinearities or for local Lipschitz statements outside this one-sided estimate.

For a generally locally Lipschitz nonlinearity, a bounded-regime alternative gives a local constant of the form

$$L_N \leq \frac{1 + R^2}{\varepsilon^2}. \quad (78)$$

Let

$$y(t) = \|v_\theta(t)\|_H^2. \quad (79)$$

The theorem is therefore formulated as an a priori estimate on an admissible bridge trajectory satisfying the stated hypotheses.

The estimate is also conditional on Assumption 1 and on the full weak bridge equation. Theorem 1 concerns an idealized weak bridge trajectory satisfying (72) for every test function in V . The finite-dimensional Dirac–Frenkel parameter dynamics of Section 2.4.3 enter the global analysis through the discrete propagation framework when their residual or flagged-step defect is bounded separately.

Theorem 1. *Continuous weak energy estimate for the idealized local bridge problem. Let $I_n = (t_n, t_{n+1})$. Assume $v_\theta \in L^2(I_n; V)$, $\partial_t v_\theta \in L^2(I_n; V')$, $g \in L^2(I_n; V')$, and let the bounded bulk transport extension w be prescribed according to Assumption 1, in particular $w \in L^\infty(I_n; L^\infty(\Omega)^d)$. Suppose that the weak bridge Equation (72) holds for all $\psi \in V$ and for almost every $t \in I_n$, that the diffusion form satisfies (74), and that the nonlinear term satisfies either Assumption 2 or Assumption 3. Define*

$$C_w = \frac{c_V^2 \|w\|_{L^\infty(\Omega)}^2}{\alpha}, \quad \lambda = 2(\beta + C_w + L_N), \quad C_f = \frac{2}{\alpha}. \quad (80)$$

Then, for almost every $t \in (t_n, t_{n+1})$,

$$y'(t) \leq \lambda y(t) + C_f \|g(t)\|_{V'}^2. \quad (81)$$

Consequently, for every $t \in [t_n, t_{n+1}]$,

$$y(t) \leq e^{\lambda(t-t_n)} y(t_n) + C_f \int_{t_n}^t e^{\lambda(t-s)} \|g(s)\|_{V'}^2 ds. \quad (82)$$

Proof. Take the admissible test function $\psi = v_\theta(t)$ in (72); this is allowed because $v_\theta(t) \in V$ for almost every t . Using the Lions–Magenes identity (71),

$$\frac{1}{2} y'(t) + a(v_\theta, v_\theta) = \langle \mathcal{N}(v_\theta), v_\theta \rangle_{V', V} + \langle g, v_\theta \rangle_{V', V} - b_w(v_\theta, v_\theta). \quad (83)$$

By Lemma 1,

$$|b_w(v_\theta, v_\theta)| \leq \frac{\alpha}{4} \|v_\theta\|_V^2 + C_w \|v_\theta\|_H^2. \quad (84)$$

By Young's inequality applied to the forcing term,

$$|\langle g, v_\theta \rangle_{V', V}| \leq \frac{\alpha}{4} \|v_\theta\|_V^2 + \frac{1}{\alpha} \|g\|_{V'}^2. \quad (85)$$

The nonlinear term satisfies

$$\langle \mathcal{N}(v_\theta), v_\theta \rangle_{V', V} \leq L_N \|v_\theta\|_H^2. \quad (86)$$

The Gårding inequality gives

$$a(v_\theta, v_\theta) \geq \alpha \|v_\theta\|_V^2 - \beta \|v_\theta\|_H^2. \quad (87)$$

Substitution of (84)–(87) into (83) yields

$$\frac{1}{2} y'(t) + \frac{\alpha}{2} \|v_\theta\|_V^2 \leq (\beta + C_w + L_N) y(t) + \frac{1}{\alpha} \|g(t)\|_{V'}^2. \quad (88)$$

Dropping the nonnegative term on the left and multiplying by 2 gives (81), with λ as defined in (80). The integral estimate (82) follows from Gronwall's inequality. $\square$

Remark 1. The proof of Theorem 1 uses the weak transport bound (67)–(68) from Lemma 1. The resulting energy estimate is valid at the Sobolev level for the admissible bridge state and bounded transport extension, so no pointwise differentiability or classical spatial derivative of the bridge state is needed.

Remark 2. Assumption 1 and the boundary-admissible manifold are structural hypotheses of the idealized continuous bridge. Theorem 1 applies to the lifted, projected, or constrained bridge variable in V and to a trajectory satisfying the full weak bridge equation. Finite-dimensional projected parameter paths are connected to this estimate through residual control or through the flagged-step defect hypotheses of Theorem 2.

Remark 3. The weak Dirac–Frenkel system (34)–(40) tests tangent directions of the neural manifold, whereas Equation (72) is posed for every test function in V . This separation clarifies the different roles of the projected parameter dynamics, the idealized full weak bridge estimate, and the flagged/restart defect hypotheses used to transfer local correction quality to the global fully discrete march.

Remark 4. For the Allen–Cahn cubic used in the benchmark, the proof uses the direct one-sided weak estimate displayed above, so L_N may be chosen as ε^{-2} under the stated sign convention. The bounded-regime or truncated-nonlinearity wording is retained only to cover more general local nonlinearities for which a separate one-sided bound is assumed rather than derived globally.

Remark 5. Practical interpretation of Assumption 3. For the reported Allen–Cahn benchmark, the continuous estimate may use the direct one-sided Allen–Cahn bound stated above. For a more general nonlinearity where that direct bound is unavailable, Assumption 3 is an explicit admissibility hypothesis, or the effect of a stated limiter or truncation. A particular fully discrete implementation must then provide the corresponding admissibility or defect control through its own analysis or diagnostics.

2.5.4. Conditional Discrete Error Propagation

We now connect the bridge estimate to a macroscopic discrete propagation framework. The discrete theorem is formulated through explicit one-step hypotheses for smooth steps and flagged/restart steps. In this theorem, the flagged set includes both physical event windows and any restart-recovery steps on which the algorithm uses BDF1 or BDF2 before returning to IMEX-BDF3.

Let the uniform time grid be

$$t_n = n\tau, \quad n = 0, 1, \dots, N, \quad N\tau = T. \quad (89)$$

Let $u(t_n)$ denote the exact solution sampled at t_n , let U^n denote the numerical approximation, and define

$$E_n = \|u(t_n) - U^n\|_H. \quad (90)$$

Let J_f denote the extended flagged set. It is the union of the event/detector set and the restart-recovery set. Thus smooth IMEX-BDF3 behavior is assumed only for $n \notin J_f$.

More explicitly, if J_{det} is the set of detected event indices and two recovery steps are used after each accepted event, then

$$J_f = J_{\text{det}} \cup R(J_{\text{det}}), \quad R(J_{\text{det}}) = \{n + r : |n \in J_{\text{det}}, |r \in \{1, 2\}\} \cap \{0, 1, \dots, N - 1\}.$$

In the oracle reference, the detected set is induced by the exact jump times. In the manufactured-benchmark diagnostic realization, the analogous flagged set is produced by the event-local diagnostic detector described in Section 3.4. This definition keeps the smooth IMEX-BDF3 hypothesis separated from the deliberately order-reduced BDF1/BDF2 recovery steps.

Assume that, in the considered refinement regime, the constants entering the following one-step estimates are independent of the time-step size. Specifically, suppose that there exist nonnegative constants κ_s, κ_j, C_s , an order p , and local flagged-step defects $\delta_n(\tau)$ such that the following one-step hypotheses hold.

Smooth regime, $n \notin J_f$:

$$E_{n+1} \leq (1 + \kappa_s \tau) E_n + C_s \tau^{p+1}. \quad (91)$$

Flagged or restart regime, $n \in J_f$:

$$E_{n+1} \leq (1 + \kappa_j \tau) E_n + \delta_n(\tau). \quad (92)$$

These estimates are hypotheses of the theorem. Theorem 2 is a discrete Gronwall-type propagation statement: once the smooth-step and flagged/restart one-step bounds are

verified for a particular implementation, it describes how the corresponding local defects accumulate in the global error.

Define A_n and ℓ_n as follows. If $n \notin J_f$, set

$$A_n = 1 + \kappa_s \tau, \quad \ell_n = C_s \tau^{p+1}. \quad (93)$$

If $n \in J_f$, set

$$A_n = 1 + \kappa_j \tau, \quad \ell_n = \delta_n(\tau). \quad (94)$$

Then both regimes are summarized by

$$E_{n+1} \leq A_n E_n + \ell_n. \quad (95)$$

For integers $0 \leq m \leq n$, define

$$\Phi_{m,n} = \prod_{r=m}^{n-1} A_r, \quad \Phi_{n,n} = 1. \quad (96)$$

Theorem 2. *Conditional discrete propagation under explicitly assumed flagged-step defect control. Under the one-step assumptions (91)–(92),*

$$E_n \leq \Phi_{0,n} E_0 + \sum_{m=0}^{n-1} \Phi_{m+1,n} \ell_m. \quad (97)$$

If

$$K = \max\{\kappa_s, \kappa_j\}, \quad (98)$$

then

$$\Phi_{m,n} \leq \exp(K(t_n - t_m)). \quad (99)$$

Consequently,

$$E_n \leq e^{KT} E_0 + \sum_{m=0}^{n-1} e^{K(t_n - t_{m+1})} \ell_m. \quad (100)$$

If the initial data are exact, $E_0 = 0$, then

$$E_n \leq e^{KT} \left(C_s T \tau^p + \sum_{m \in J_f} \delta_m(\tau) \right). \quad (101)$$

If

$$\sum_{m \in J_f} \delta_m(\tau) \leq C_J \tau^p, \quad (102)$$

then

$$E_n \leq e^{KT} (C_s T + C_J) \tau^p. \quad (103)$$

A sufficient condition is that $|J_f| \leq J_M$ and that, for every $m \in J_f$,

$$\delta_m(\tau) \leq C_\delta \tau^p. \quad (104)$$

In that case,

$$E_n \leq e^{KT} (C_s T + J_M C_\delta) \tau^p. \quad (105)$$

Proof. Iterating the scalar recurrence (95) gives (97). Since $A_r \leq 1 + K\tau \leq e^{K\tau}$, (99) follows, and substitution gives (100). Splitting the sum into smooth and flagged indices gives (101). The remaining statements follow from (102), the uniform bound on $|J_f|$, and (104). $\square$

Remark 6. *Theorem 2 is a conditional propagation result. It states that smooth-step macroscopic order is retained when the accumulated event and restart defects have the same asymptotic size as the smooth-step truncation contribution. The theorem is therefore useful as a bookkeeping and design criterion: it identifies the defect control that a detector/bridge/restart implementation must satisfy.*

Remark 7. *The conditional estimate clarifies the accuracy requirement for a practical BDF1-to-BDF2 restart. A plain BDF1 post-event recovery generally produces only an $O(\Delta t^2)$ state defect. Therefore, retaining the $p = 3$ smooth-step order requires the accepted bridge endpoint and the restart reconstruction to satisfy the stated $O(\Delta t^3)$ accumulated flagged-defect condition. If the accumulated flagged or restart defects are only $O(\Delta t^q)$, the theorem gives the corresponding lower effective order q .*

Thus, in the absence of a verified third-order flagged-defect bound, the practical BDF1/BDF2 restart is best interpreted as a controlled defect-reduction and history-repair mechanism. The numerical results are used to demonstrate reduction of the dominant history-contamination error and improvement of finite-step error constants under selective detection.

3. Results: Numerical Verification and Manufactured-Benchmark Diagnostic Assessment

This section examines the proposed framework on a manufactured one-dimensional Allen–Cahn benchmark for which the interface motion, jump times, and exact solution are known. The computational backbone is an IMEX-BDF3 discretization, and the benchmark is designed so that a moving diffuse interface crosses the Eulerian grid under a prescribed piecewise-constant velocity law. This setting isolates the temporal interface-crossing effect discussed in Sections 2–4 and permits a direct comparison between the uncorrected multi-step backbone, an event-aligned oracle restart reference, and a manufactured-benchmark diagnostic realization.

The numerical study is organized around three realizations. The first is the baseline IMEX-BDF3 scheme, which advances the state without dedicated event-local correction. The second is an oracle restart realization, which uses exact event information and serves as a controlled reference for the mechanism identified in the theoretical analysis. The third is a manufactured-benchmark diagnostic realization, which combines detector-triggered flagged steps, local corrected candidates, acceptance logic, and fallback selection. This organization allows the numerical results to distinguish mechanism-level support from implementation-level behavior.

The numerical evidence is used as a diagnostic test of the ingredients required by the conditional propagation theorem: event localization, local bridge accuracy, boundary-consistent endpoint sampling, and restart-consistent history reconstruction in the manufactured benchmark. These experiments provide benchmark support for the proposed mechanism and clarify the conditions under which a detector-driven implementation can approach the hypotheses of Theorem 2.

Throughout this section, the manufactured benchmark uses the spatial domain $\Omega = [-1, 1]$, final time $T = 0.8$, interface-width parameter $\varepsilon = 0.04$, diffusivity $\nu = 1$, initial interface position $x_\Gamma(0) = x_0 = -0.25$, jump times 0.2, 0.4, and 0.6, and piecewise velocities 0.8, -0.4 , 0.8, and -0.4 . The practical study is reported over the step-size sequence $\Delta t \in \{8 \times 10^{-4}, 4 \times 10^{-4}, 2 \times 10^{-4}, 10^{-4}\}$.

All computations and figure-generation workflows in this section were carried out with custom Python scripts using Python 3.11.9, with NumPy, SciPy, and Matplotlib, for the IMEX-BDF3 backbone, the oracle restart reference, the detector-driven corrected variants, numerical-data generation, and figure production. The exact package information is provided in the requirements/environment files in the Supplementary Materials. The accompanying code package includes the source scripts, README instructions, requirements file, environment file, and figure-generation workflow needed to regenerate the numerical values reported in the tables and code-generated figures. Code availability is stated in the Supplementary Materials and Data Availability Statement.

3.1. Manufactured Benchmark Specification

For reproducibility, the manufactured initial-boundary-value problem is stated explicitly. The prescribed interface velocity is

$$c(t) = \begin{cases} 0.8, & 0 \leq t < 0.2, \\ -0.4, & 0.2 \leq t < 0.4, \\ 0.8, & 0.4 \leq t < 0.6, \\ -0.4, & 0.6 \leq t \leq 0.8, \end{cases} \quad (106)$$

and the interface location is

$$x_{\Gamma}(t) = -0.25 + \int_0^t c(s) ds. \quad (107)$$

The exact manufactured profile used to evaluate errors is

$$u_{\text{ex}}(x, t) = \tanh\left(\frac{x - x_{\Gamma}(t)}{\sqrt{2}\varepsilon}\right). \quad (108)$$

The forced Allen–Cahn equation solved in the benchmark is defined by substituting this exact profile into the differential equation:

$$\partial_t u = \nu \partial_{xx} u + \varepsilon^{-2} u(1 - u^2) + f(x, t), \quad f = \partial_t u_{\text{ex}} - \nu \partial_{xx} u_{\text{ex}} - \varepsilon^{-2} u_{\text{ex}}(1 - u_{\text{ex}}^2). \quad (109)$$

For the reported validation $\nu = 1$; the general ν form in (109) is retained to show how the manufactured forcing is constructed from the prescribed exact solution.

The initial and boundary data are imposed from the exact solution:

$$u(x, 0) = u_{\text{ex}}(x, 0), \quad u(-1, t) = u_{\text{ex}}(-1, t), \quad u(1, t) = u_{\text{ex}}(1, t). \quad (110)$$

The spatial grid and error norm are also fixed. With $n_x = 1025$ unless otherwise stated, the grid is

$$x_i = -1 + ih, \quad h = \frac{2}{n_x - 1}, \quad i = 0, \dots, n_x - 1, \quad (111)$$

the Laplacian is approximated on the interior nodes by the centered second-order finite difference, and boundary values are supplied by the exact Dirichlet data above. The reported discrete L^2 -error is

$$E_2(t_n) = \left[h \sum_{i=0}^{n_x-1} |U_i^n - u_{\text{ex}}(x_i, t_n)|^2 \right]^{1/2}. \quad (112)$$

The exact nonhomogeneous Dirichlet boundary values are inserted through the boundary-vector contribution in the interior finite-difference system. Equivalently, at the continuous weak level, this corresponds to the lifted-variable convention introduced in Section 2.1.

The initial BDF3 history is initialized by exact sampling of u_{ex} at the required starting time levels. The oracle restart reference uses the known event locations only to rebuild clean post-event history; the detector-driven realization uses the event-local diagnostic detector and acceptance logic described above.

3.2. Global Error Comparison Across the Tested Time-Step Sizes

The first level of assessment concerns the final-time and maximal-in-time L^2 -errors of the three realizations. These quantities provide the clearest global view of whether event-local correction changes the macroscopic behavior of the time integrator.

Table 4 reports these final-time and maximal-in-time L^2 -errors together with apparent pairwise final-error slopes across the tested step sizes.

Table 4. Final-time and maximal-in-time L^2 -errors for the baseline IMEX-BDF3 method, the oracle restart reference, and the manufactured-benchmark diagnostic realization over the tested step sizes $\Delta t \in \{8 \times 10^{-4}, 4 \times 10^{-4}, 2 \times 10^{-4}, 10^{-4}\}$. The last three columns report consecutive-step apparent pairwise final-error slopes, not asymptotic convergence orders. The oracle restart provides a mechanism-level reference based on exact event information, while the manufactured-benchmark diagnostic realization combines detector-triggered correction, acceptance logic, and fallback selection.

Δt	Baseline Final L^2	Oracle Final L^2	Certified Final L^2	Baseline Max L^2	Oracle Max L^2	Certified Max L^2	Baseline Slope	Oracle Slope	Certified Slope
8×10^{-4}	1.131×10^{-1}	1.032×10^{-1}	1.554×10^{-1}	1.554×10^{-1}	1.374×10^{-1}	1.554×10^{-1}	—	—	—
4×10^{-4}	1.173×10^{-3}	2.90×10^{-5}	3.53×10^{-5}	1.525×10^{-3}	4.3×10^{-5}	1.75×10^{-4}	6.591	11.795	12.105
2×10^{-4}	6.14×10^{-4}	2.97×10^{-5}	4.55×10^{-5}	8.30×10^{-4}	4.8×10^{-5}	8.8×10^{-5}	0.934	−0.031	−0.368
10^{-4}	3.30×10^{-4}	2.98×10^{-5}	4.59×10^{-5}	4.55×10^{-4}	4.9×10^{-5}	7.4×10^{-5}	0.896	−0.005	−0.013

Table 4 separates the numerical behavior into three patterns. The baseline IMEX-BDF3 error decreases under refinement but remains affected by event-history contamination. The oracle restart reference is much smaller at the finer step sizes, supporting the interpretation that event-aligned history reconstruction suppresses the dominant interface-induced defect. The manufactured-benchmark diagnostic realization is worse at the coarsest step, where the event window is not sufficiently resolved, but improves on the baseline at the three finer step sizes. The displayed slopes are diagnostic pairwise slopes over a short refinement sequence and are interpreted cautiously when the coarsest step is unresolved or when the finest steps are affected by an error floor.

The convergence pattern reported in Table 4 is visualized in Figure 3.

Figure 3 presents the same information as Table 4 in graphical form. The baseline curve decreases under refinement but remains well above the oracle reference at the finer steps. The oracle curve drops rapidly and then stabilizes at a low level. The manufactured-benchmark diagnostic curve is clearly lower than the baseline over part of the tested range, but it turns upward at the finest step. Read together, Table 4 and Figure 3 show that event-aligned correction addresses the relevant mechanism, while the present detector-driven realization does not retain the same level of improvement uniformly across the full refinement range. Because a terminal-error plot does not indicate when this deterioration develops, the next subsection turns to the full error histories at the finest tested step size.




Figure 3. Final L^2 -error versus time-step size Δt for the baseline IMEX-BDF3 method, the oracle restart reference, and the manufactured-benchmark diagnostic realization. The figure compares the global terminal accuracy of the uncorrected multistep backbone, the event-aligned oracle reference, and the manufactured-benchmark diagnostic pipeline over the tested refinement range.

3.3. Error History at the Finest Tested Step Size

Global final-time errors do not reveal where the deterioration occurs in time. The second diagnostic therefore examines the full L^2 -error histories at the finest stored step size, $\Delta t = 10^{-4}$, together with the jump markers and the flagged intervals produced by the diagnostic detector.

Figure 4 reveals the temporal structure of error accumulation. The baseline IMEX-BDF3 history exhibits visible rises at the event times, followed by elevated post-event levels. The oracle restart remains at a much lower level, with narrow event-associated features and no comparable cumulative growth. This behavior is consistent with the interpretation that an important source of contamination is the transfer of multistep history across temporal interface events.

The manufactured-benchmark diagnostic history begins at a low error level and remains competitive during the early part of the interval, but it then rises after the first event and remains elevated over the later part of the simulation. The shaded flagged neighborhoods remain sparse and event-centered after the detector correction, so the later elevation cannot be attributed to a loss of detector locality. Instead, it points to the remaining interaction among local correction quality, restart recovery, and event resolution near the flagged windows. This temporal behavior motivates the detector-level and correction-quality analyses in the next subsections.




Figure 4. Time history of the L^2 -error at the finest tested step size $\Delta t = 10^{-4}$ for the baseline IMEX-BDF3 method, the oracle restart reference, and the manufactured-benchmark diagnostic realization.

Dashed vertical lines indicate the three jump times $t = 0.2, 0.4$, and 0.6 . Light-blue shaded bands indicate the event-centered flagged intervals detected by the event detector.

3.4. Detector Performance and Restoration of Locality

The behavior seen in Figure 4 indicates that the principal implementation difficulty is not event detection in the sense of sensitivity but event localization in the sense of keeping the flagged set sparse and local. Table 5 makes this point explicit.

Table 5. Detector diagnostics for the event-local diagnostic detector. Reported quantities include the number of oracle event steps, flagged steps, true positives, false positives, false negatives, precision, recall, F1-score, overflag ratio, and mean localization error in time-step units. The event-local diagnostic detector retains full event recall while flagging only a small event-local set.

Δt	Oracle Events	Flagged Steps	True Pos.	False Pos.	False Neg.	Prec.	Rec.	F1	Overflag Ratio	Mean Loc. Err.
8×10^{-4}	6	9	6	3	0	0.667	1.000	0.800	0.500	0.0
4×10^{-4}	6	9	6	3	0	0.667	1.000	0.800	0.500	0.0
2×10^{-4}	6	9	6	3	0	0.667	1.000	0.800	0.500	0.0
1×10^{-4}	6	9	6	3	0	0.667	1.000	0.800	0.500	0.0

In the detector convention used here, each physical velocity-jump time contributes two oracle event-step indices, corresponding to the event and restart-window bookkeeping used by the implementation. Thus the three physical jump times $t = 0.2, 0.4$, and 0.6 generate six oracle event steps in Table 5.

With $t_n = n\Delta t$ and jump times $\tau_j \in \{0.2, 0.4, 0.6\}$, the oracle event mask is $\chi_n^{\text{evt}} = 1$ when $t_n < \tau_j \leq t_{n+1}$ or $\tau_j - t_n$ is below the alignment tolerance for at least one j . The manufactured-benchmark diagnostic detector used in Tables 5 and 6 first defines the event-local diagnostic set

$$\mathcal{J}_{\text{diag}} = \mathcal{J}_{\text{jump}} \cup \mathcal{J}_{\text{speed}},$$

where

$$\mathcal{J}_{\text{jump}} = \{n : \chi_{\text{jump}}(n) = 1\}, \quad \mathcal{J}_{\text{speed}} = \{n : \chi_{\text{speed}}(n) = 1\}.$$

Table 6. Correction statistics for the manufactured-benchmark diagnostic realization with the event-local diagnostic detector. For each tested step size, the table reports the number of flagged steps, restart count, accepted and non-accepted corrections, and the distribution of selected correction tags. The results show that the diagnostic pipeline intervenes on a sparse flagged set and that the selected local correction type changes with the time-step size.

Δt	Flagged Steps	Restart Count	Accepted Corr.	Rejected Corr.	Accepted Neural	Accepted Subcycle	Accepted Front
8×10^{-4}	9	12	9	0	0	9	0
4×10^{-4}	9	12	6	3	6	0	0
2×10^{-4}	9	12	6	3	6	0	0
1×10^{-4}	9	12	6	3	2	4	0

The flagged set used by the diagnostic realization is then obtained by one-step forward padding,

$$\mathcal{J}_{\text{flag}} = \mathcal{J}_{\text{diag}} \cup \{n + 1 : n \in \mathcal{J}_{\text{diag}}\},$$

with indices outside the time grid discarded. The current-state-only geometric trigger is disabled in this event-local validation because it flags nearly every step and destroys locality. This benchmark diagnostic detector uses the manufactured jump/speed-change information to isolate the event-local correction mechanism. Predictive deployment would require a residual-based, geometric, or a posteriori event indicator that does not use manufactured solution data.

Table 5 shows that the event-local diagnostic detector does not miss the benchmark events. Recall is 1.0 for all tested step sizes, the number of false negatives is zero, and the mean localization error is zero in step units. In contrast with the broad geometric trigger, the flagged-step count is nine rather than thousands, while the precision remains 0.667 across the tested grid. The remaining practical issue is therefore the accuracy of the local correction on the flagged windows. The same detector behavior is shown graphically in Figure 5.




Figure 5. Detector diagnostics for the event-local diagnostic detector. The upper panel compares the number of flagged steps with the number of oracle event steps. The lower panel reports precision, recall, and overflag ratio. The event-local diagnostic detector retains full recall while keeping the flagged set sparse.

Figure 5 confirms visually that the flagged set remains close to the oracle event set rather than spreading across the full trajectory.

3.5. Acceptance Statistics, Restarts, and Fallback Behavior

Once a step has been flagged, the diagnostic realization constructs the four local endpoint candidates just described and applies the acceptance and fallback logic neural $\rightarrow$ subcycle $\rightarrow$ front $\rightarrow$ IMEX. The diagnostic question is therefore whether the remaining finest-step error floor is associated with local candidate quality, restart recovery, or sharp-interface resolution, rather than with broad nearly every-step activation of the geometric trigger. The implementation compares candidate updates by residual-based, reference-distance, exact-distance audit, and norm-growth quantities; the exact-distance ratio is reported diagnostically and is not used as a theoretical assumption.

Table 6 reports the corresponding restart, acceptance, non-acceptance, and fallback statistics.

The restart count is not identical to the number of flagged detector indices. It is the audit-log count of restart-mode actions, including the BDF1/BDF2 recovery entries required

after accepted event windows. Therefore the restart count can exceed the flagged-step count; in Table 6, nine flagged detector indices produce twelve restart/recovery entries.

Table 6 shows that the diagnostic pipeline acts only on a small event-local set, so the local bridge and restart machinery no longer operate over nearly the full trajectory. The same acceptance and fallback behavior is visualized in Figure 6.




Figure 6. Acceptance, restart, fallback, and selected-correction statistics for the manufactured-benchmark diagnostic realization across the tested step sizes. The upper panel reports accepted and non-accepted corrections together with restart counts. The lower panel summarizes the selected accepted-correction tags. The figure shows that the corrected diagnostic pipeline remains localized to the flagged event windows.

Figure 6 complements Table 6 by showing accepted and non-accepted corrections, restart counts, and selected correction tags across the tested steps. The selected correction type changes with refinement, with subcycle corrections appearing at the finest step. The figure clarifies that the practical difficulty is no longer systematic overflagging but the residual accuracy of the local correction and restart sequence on the event-centered windows.

3.6. Local Event-Window Analysis

Because the mechanism identified in Sections 2–4 is intrinsically event-local, the global diagnostics above should be complemented by local views around the jump times $t = 0.2$, $t = 0.4$, and $t = 0.6$. The event-window zooms provide the most direct numerical picture of how the three realizations behave near the temporal interfaces.

These event-window zooms are shown in Figure 7.




Figure 7. Local L^2 -error histories near the three jump times $t = 0.2$, $t = 0.4$, and $t = 0.6$ at the finest tested step size $\Delta t = 10^{-4}$. The figure compares the baseline IMEX-BDF3 method, the oracle restart

reference, and the manufactured-benchmark diagnostic realization in narrow event-centered windows and highlights the local evolution of the error immediately before and after each interface crossing.

Figure 7 shows local L^2 -error histories near the three jump times at the finest tested step size. Near the first event at $t = 0.2$, the manufactured-benchmark diagnostic realization is initially very small and remains competitive immediately after the crossing. This indicates that the practical machinery can produce locally useful corrections. At the same time, the diagnostic curve rises away from its post-event level, whereas the oracle restart stays near the low-error regime.

The second and third panels show the same mechanism more clearly. By the time the later events are reached, the manufactured-benchmark diagnostic curve no longer retains the low-error character of the oracle reference and becomes comparable to, or above, the baseline in the local event windows. These zooms refine the global interpretation: the diagnostic realization can be locally useful, especially near the first event, but repeated event-centered interventions remain sensitive to local correction quality, restart recovery, and event resolution.

3.7. Time-Step Heatmaps of Relative Behavior

The event-window zooms are local in time, and the convergence curves are global in the final-time sense. Figure 8 gives a complementary view by examining the full $(t, \Delta t)$ -dependence of the relative behavior of the locally corrected realizations against the baseline solver.




Figure 8. Time-step heatmaps of relative behavior with respect to the baseline solver. The upper panel shows the oracle relative-error behavior, the middle panel shows the diagnostic relative-error behavior, and the lower panel shows the flagged-step density map after detector correction. Dashed vertical lines mark the prescribed jump times $t = 0.2$, $t = 0.4$, and $t = 0.6$. Negative values in the relative-error panels indicate improvement over the baseline IMEX-BDF3 backbone. In the lower panel, white denotes flagged-step density 0 (unflagged), and black denotes flagged-step density 1 (flagged).

Figure 8 displays heatmaps of the oracle and diagnostic relative-error behavior together with the flagged-step density. The upper panel is predominantly negative, indicating that the oracle restart improves on the baseline over most of the time-resolution domain, especially near event neighborhoods. This provides broad visual support for the mechanism-level interpretation.

The middle panel, corresponding to $\log_{10}(\text{diagnostic}/\text{baseline})$, is more mixed. Substantial negative regions are present, particularly at coarser and intermediate resolutions, confirming that the manufactured-benchmark diagnostic realization improves on the baseline over a significant portion of the domain. Positive regions also appear, especially toward finer resolution and later times, and these are the regions in which the diagnostic realization loses its advantage.

The lower panel shows the flagged-step density and provides a visual counterpart to Table 5 and Figure 5. The nearly white background indicates that activation remains event-local. Thus, the remaining diagnostic error is more consistently associated with local-correction and sharp-interface resolution limits than with systematic overflagging.

3.8. Local Acceptance Quality over Accepted Corrections

The next question concerns the accepted corrections themselves. Specifically, when a local correction is accepted by the diagnostic pipeline, does it improve the surrogate quantities monitored by the acceptance layer, and how closely does that improvement align with the exact local error measure? The event-acceptance data permit a direct aggregate answer.

Table 7 summarizes these local acceptance-quality measures over the accepted corrections.

Table 7. Mean local quality ratios over accepted corrections. Ratios below one indicate improvement relative to the corresponding uncorrected IMEX candidate. The reported quantities summarize the residual ratio, reference-distance ratio, and exact-distance ratio together with the number of accepted corrections and the subset accepted on oracle-event steps. These data show that accepted corrections generally improve local residual and reference-distance measures, while the exact-distance ratio can remain sensitive at coarse resolution.

Δt	Accepted Corr.	Accepted on Oracle Steps	Mean Residual Ratio	Mean Ref.-Dist. Ratio	Mean Exact-Dist. Ratio
8×10^{-4}	9	6	0.229315	0.018885	3.775669
4×10^{-4}	6	3	0.094218	0.063462	0.727088
2×10^{-4}	6	3	0.230167	0.174843	0.877615
1×10^{-4}	6	3	0.173746	0.143375	0.941023

The remaining finest-step error floor is therefore better understood as a consequence of applying locally reasonable corrections on demanding event-centered windows, where sharp-interface resolution, candidate quality, and restart-history reconstruction interact.

The variation of these acceptance ratios across the tested step sizes is plotted in Figure 9.

Figure 9 displays the same information as Table 7 in graphical form. The residual-ratio and reference-distance curves show local improvement under acceptance, whereas the exact-distance curve is more variable. This distinction is informative: internal acceptance quantities improve consistently, while repeated event interventions still require coordination with exact local accuracy and global history recovery. A broader geometric view of the same trend is obtained next from three-dimensional error surfaces over time and step size.




Figure 9. Local acceptance-quality diagnostics over accepted corrections. The plot reports the residual ratio, the reference-distance ratio, and the exact-distance ratio across the tested step sizes. Ratios below one indicate improvement relative to the corresponding uncorrected IMEX candidate and provide a local view of the acceptance layer used by the manufactured-benchmark diagnostic realization.

3.9. Three-Dimensional Absolute Error Surfaces

A further global diagnostic is obtained by viewing the absolute error as a surface over $(t, \Delta t)$. This representation does not replace the two-dimensional convergence and history plots, but it provides a compact view of the geometry of the error landscape for the three realizations.

Figure 10 displays three-dimensional surfaces of $\log_{10}(L^2 \text{error})$ for the baseline IMEX-BDF3 method, the oracle restart reference, and the manufactured-benchmark diagnostic realization. The baseline surface shows the expected event-associated rises and sustained elevated regions after the jumps. The oracle surface lies markedly lower and exhibits narrower event-associated features, which is consistent with the mechanism-level conclusion that event-aligned restarting suppresses the principal contamination induced by carrying multistep history across the interfaces.




Figure 10. Three-dimensional surfaces of $\log_{10}(L^2 \text{error})$ over $(t, \Delta t)$ for the baseline IMEX-BDF3 method, the oracle restart reference, and the manufactured-benchmark diagnostic realization. The surfaces provide a global view of the error landscape across time and resolution and highlight the contrast between the uncorrected backbone, the event-aligned oracle reference, and the manufactured-benchmark diagnostic pipeline.

The manufactured-benchmark diagnostic surface occupies an intermediate regime. Large regions are lower than the baseline, whereas the oracle reference remains lower over

the full domain. Elevated regions appear toward finer resolution and later portions of the interval, in agreement with the deterioration seen in Figure 4 and in the event-window zooms. This figure therefore shows that the diagnostic realization captures part of the mechanism over part of the domain. The next figure sharpens this comparison by plotting the same behavior directly relative to the baseline solver.

3.10. Three-Dimensional Relative-Error Surfaces Against the Baseline Solver

The most informative three-dimensional diagnostic is the relative behavior of the locally corrected realizations against the baseline solver. The next figure therefore suppresses the common absolute scale and focuses directly on gain or loss relative to the uncorrected IMEX-BDF3 backbone.

Figure 11 displays $\log_{10}(\text{oracle}/\text{baseline})$ and $\log_{10}(\text{diagnostic}/\text{baseline})$ as surfaces over $(t, \Delta t)$. The oracle-relative surface is dominated by negative regions, which means that the oracle restart is better than the baseline over most of the domain and often by a substantial margin. This is one of the clearest global visual confirmations of the framework mechanism: once the interface event is aligned with the restart, the contamination associated with inherited multistep history is sharply reduced.

The diagnostic-relative surface is qualitatively different. Negative regions are present and are not negligible, which confirms that the diagnostic realization outperforms the baseline over a substantial part of the time-resolution plane. Positive regions are also present, especially at finer resolution and later times. The figure therefore condenses the main practical conclusion of the section into a single visual summary: the manufactured-benchmark diagnostic realization is often beneficial, but its performance relative to the baseline remains sensitive to detector selectivity and the preservation of locality. These relative surfaces prepare the aggregate interpretation collected in the concluding subsection of the numerical study.




Figure 11. Three-dimensional relative-error surfaces against the baseline IMEX-BDF3 solver. The left panel shows $\log_{10}(\text{oracle}/\text{baseline})$ and the right panel shows $\log_{10}(\text{diagnostic}/\text{baseline})$ over $(t, \Delta t)$. Negative values indicate improvement relative to the baseline solver, while positive values indicate deterioration.

3.11. Overall Numerical Interpretation

Taken together, the numerical evidence supports two conclusions. First, the oracle restart consistently outperforms the baseline in the final-time errors, time histories, event-

window zooms, heatmaps, and relative three-dimensional surfaces. This supports the mechanism-level interpretation that inherited multistep history across temporal interfaces is a dominant source of the observed degradation. Second, the manufactured-benchmark diagnostic realization improves on the baseline over the resolved part of the tested range, with performance governed by local correction quality, restart recovery, detector selectivity, and sharp-interface resolution. The accepted-correction statistics show that the acceptance layer can improve residual and reference-distance quantities, while the remaining error floor indicates that local improvement must also be preserved through re-entry into the global multistep history.

3.12. Interface-Thickness Sensitivity, Diagnostic Cost, and Sharp-Interface Stress Test

To assess interface-thickness dependence, the manufactured-benchmark diagnostic realization was also tested at $\varepsilon = 0.04$, $\varepsilon = 0.03$, and $\varepsilon = 0.02$. At the finest common time step $\Delta t = 10^{-4}$, the diagnostic realization improves the baseline for all three ε values, but the improvement factor decreases as the interface becomes sharper.

Across these interface-thickness tests, the event-local diagnostic detector remains tied to the manufactured jump/speed-change windows with one-step padding and does not reintroduce the broad geometric trigger responsible for excessive activation in the broad-trigger diagnostic. Therefore, the reduced improvement factor at smaller interface width is attributed mainly to sharper event-window resolution, local candidate accuracy, and the fixed local solver budget rather than to deteriorating detector selectivity.

Table 8 summarizes the interface-thickness sensitivity at the finest common time step.

Table 8. Interface-thickness sensitivity at the finest common time step. The diagnostic improvement factor is computed as the baseline final -error divided by the diagnostic final -error. Values followed by $\times$ denote dimensionless fold-change factors.

ε	Baseline Final L^2	Oracle Final L^2	Diagnostic Final L^2	Diagnostic Improvement Factor
0.04	3.3003×10^{-4}	2.9751×10^{-5}	4.5932×10^{-5}	$7.2\times$
0.03	4.1945×10^{-4}	5.5461×10^{-5}	9.0480×10^{-5}	$4.6\times$
0.02	6.4532×10^{-4}	1.4076×10^{-4}	2.4735×10^{-4}	$2.6\times$

The $\varepsilon = 0.02$ case was then examined with a refined time-step ladder. The coarse step $\Delta t = 8 \times 10^{-4}$ produced overflow/NaN values and is therefore outside the stable resolved regime for this fixed grid and splitting. For $\Delta t \leq 10^{-4}$ the diagnostic correction improves the baseline, but both oracle and diagnostic errors plateau, indicating a residual floor not removed by simply increasing the local micro-solver budget.

The refined sharp-interface stress test and the associated cost ratio are reported in Table 9.

Table 9. Refined $\varepsilon = 0.02$ stress test and diagnostic cost ratio relative to the baseline IMEX-BDF3 wall time. The plateau of the oracle and diagnostic errors indicates that the remaining discrepancy is no longer dominated only by the time-step size.

Δt	Baseline Final L^2	Oracle Final L^2	Diagnostic Final L^2	Diagnostic Cost Ratio
1×10^{-4}	6.4532×10^{-4}	1.4076×10^{-4}	2.4735×10^{-4}	$3.63\times$
5×10^{-5}	4.4996×10^{-4}	1.4109×10^{-4}	2.5138×10^{-4}	$3.10\times$
2.5×10^{-5}	3.5100×10^{-4}	1.4116×10^{-4}	2.5199×10^{-4}	$2.79\times$

The cost model is local: the diagnostic realization has the cost of the baseline IMEX march plus a correction cost proportional to the number of flagged steps and to the local

bridge/subcycle budget. The dimensionless cost ratio reported in Table 9 is interpreted through the following decomposition.

The overhead is therefore linked to detector selectivity. When the number of flagged steps remains bounded for a fixed number of physical events, the additional work remains localized; when activation grows proportionally to the total number of time steps, both cost and accuracy advantages may be lost. The present cost ratios are diagnostic overhead measurements for the current implementation and motivate leaner event indicators and lower-cost bridge candidates for future practical solvers.

$$C_{\text{tot}} \approx C_{\text{IMEX}} + N_f (C_{\text{bridge}} + C_{\text{cert}} + C_{\text{restart}}) \quad (113)$$

$$R_{\text{cost}} = \frac{C_{\text{tot}}}{C_{\text{IMEX}}} \quad (114)$$

Thus, a value such as $3.63 \times$ in Table 9 means that the diagnostic realization required about 3.63 times the baseline wall time. Such overhead is acceptable only if the correction is needed for accuracy and remains localized, and it motivates future work on cheaper event indicators and lower-cost bridge candidates.

3.13. Comparative Positioning Against Classical Alternatives

The numerical comparison in this paper is benchmark-based. The baseline IMEX-BDF3 march measures the effect of carrying contaminated multistep history through an interface event. The event-aligned restart reference isolates the mechanism by using exact event information and clean BDF1-to-BDF2-to-BDF3 reconstruction. The manufactured-benchmark diagnostic realization tests whether a local bridge and restart pipeline can recover part of that benefit when the event set is supplied by the benchmark detector. Classical adaptive IMEX, multirate, or subcycling strategies remain important external alternatives. The results show that, on the tested moving-interface benchmark, a selective event-local correction can outperform the unmodified IMEX-BDF3 baseline when the crossing window is resolved and the detector remains sparse. The public code repository provides the complete scripts needed to reproduce this baseline/oracle/diagnostic comparison.

4. Discussion and Limitations

The discussion centers on the localized history effect isolated in the analysis and numerical tests. When a moving interface crosses the grid, the short history used by an IMEX-BDF3 backbone may no longer represent a single smooth temporal regime, allowing an event-local defect to propagate through later updates. The proposed framework is therefore best understood as a local history-repair strategy embedded in a classical march.

This interpretation allows the present work to be positioned more precisely relative to existing time-integration and scientific machine-learning literatures. In the classical stiff-integration literature, several different lines address order loss or related accuracy difficulties, but they do so by mechanisms different from the one isolated here. Self-consistent IMEX formulations reduce inaccuracies introduced by decoupled treatment of nonlinear implicit–explicit blocks by embedding the explicit contribution inside the implicit nonlinear solve more coherently [28]. High-stage-order IMEX general linear methods reduce certain stiff-order limitations through stage-order design and improved internal consistency [30–32]. Residual-balanced IMEX decompositions preserve order even under approximate implicit solution by redirecting the remaining residual to the explicit side in a controlled way [29]. IMEX-MRI-SR methods develop restart-capable stage structures and order conditions in multirate stiff settings [26]. These are all highly relevant because they show that temporal accuracy in stiff settings can be preserved by carefully redesigning the

time integrator. However, the present paper is not proposing another globally redesigned stiff integrator. It instead identifies a localized event-sensitive defect and asks whether it can be corrected locally while leaving the classical backbone untouched elsewhere.

On the machine-learning side, the proposed bridge differs from global PINN, operator-learning, and reduced-order frameworks, whose primary goals are global approximation, repeated-query acceleration, or compression of solution manifolds [38–63]. The closest conceptual neighbors are sequential-in-time neural Galerkin, local deep Galerkin, adaptive-basis Galerkin, and related causal neural trial-space methods [64–70]. The present work adopts this local-in-time perspective but assigns the neural manifold a narrower task: producing a restart-compatible bridge on flagged intervals where the surrounding IMEX history is least reliable.

The detector-driven study shows that mechanism-level improvement depends on the quality of the practical localization and correction pipeline. The manufactured-benchmark diagnostic realization improves on the baseline once event windows are sufficiently resolved, with its advantage governed by selective localization, controlled local correction, and restart-consistent re-entry. The accepted-correction statistics confirm that the correction layer is active: accepted candidates improve residual-based and reference-distance measures, and the pipeline uses both direct neural corrections and fallback subcycle updates. The next algorithmic priority is therefore better coordination among detection, correction, and restart recovery.

This also explains the conditional form of the theoretical results. The continuous bridge analysis and the discrete propagation theorem identify conditions under which accumulated event/restart defects of matching order remain consistent with the smooth-step order under the stated defect assumptions. The numerical study is consistent with this structure in the controlled restart setting, while the practical implementation shows which ingredients remain most important for future algorithmic work: selective event localization, controlled local correction, and consistent restart-history reconstruction.

The detector-driven experiments show that localized correction can reduce the dominant history-contamination error when event localization is selective and the accepted flagged-step defects remain small enough for the conditional propagation estimate to be relevant.

The present paper formulates and analyzes a localized hybrid DNG-IMEX correction principle for interface-induced temporal order reduction and examines it on a controlled manufactured benchmark. The analysis clarifies a mechanism-level question, and the numerical study evaluates both a controlled reference setting and a diagnostic detector-driven realization.

The comparative scope is focused. Against the unmodified IMEX-BDF3 baseline, the event-aligned and detector-driven results quantify the benefit of controlling history contamination. Against broader adaptive IMEX, multirate, or production subcycling methods, the present evidence is qualitative and diagnostic. The cost ratios in Section 3.12 show that the current local correction pipeline can be significantly more expensive than the baseline, so practical efficiency will require more selective detectors and cheaper accepted candidates.

A principal limitation concerns detector transferability. The reported manufactured-benchmark diagnostic detector is event-local because the manufactured jump/speed-change windows are available in the test problem. For applications in which events are unknown or poorly localized, the hybrid correction would require a residual-based, geometric, or a posteriori event indicator that does not use manufactured solution data. Low detector selectivity may reduce or eliminate the computational benefit of local correction, even when the bridge itself remains stable, and may also increase the overhead beyond what is acceptable for practical multidimensional computations.

Section 2.5 shows that accumulated flagged/restart defects of matching order remain consistent with smooth-step order only under the stated one-step hypotheses. The controlled restart experiments support this mechanism-level interpretation, while the practical study emphasizes the continuing importance of detector selectivity.

The numerical evidence is benchmark-based. The manufactured one-dimensional Allen–Cahn setting is appropriate for isolating interface events, controlling the exact solution, and studying post-event history effects with precision. Broader testing remains important for multidimensional phase-field problems, more complex geometries, nonuniform meshes, and settings in which the relevant event set is not prescribed in advance. The numerical section therefore serves as a controlled validation of the proposed mechanism and implementation choices.

These points indicate natural directions for further work. On the analytical side, it would be valuable to strengthen the connection between event localization, bridge residual control, and restart-history consistency within a unified global estimate. On the algorithmic side, improved localization criteria, more selective activation rules, and richer local bridge constructions merit further study. On the numerical side, the same framework should be examined on broader benchmark families and on higher-dimensional problems. These directions define the next steps toward stronger interface-aware diagnostic and predictive realizations.

5. Conclusions

This paper introduced a DNG-IMEX framework for stiff parabolic problems with localized temporal interfaces and examined its analytical and numerical behavior on a manufactured Allen–Cahn benchmark with prescribed interface-velocity jumps. The central question was whether the degradation observed in a classical multistep IMEX backbone can be understood as an event-local history effect and whether a short correction strategy can improve that behavior without globally modifying the backbone.

The analytical contribution consists of corrected interface kinematics, a Sobolev-level weak transport estimate that avoids classical C^1 assumptions on the bridge state, a continuous weak energy estimate for an idealized local bridge problem, and a conditional discrete propagation theorem under explicitly assumed flagged/restart defect control.

The analytical statement is a conditional compatibility result for the diagnostic restart mechanism. If the flagged or restart defects are only second order because of a plain BDF1 recovery, the same propagation argument gives the corresponding lower order. Establishing matching-order flagged defects for a fully predictive detector is a natural topic for future analysis.

The numerical study supports this interpretation at two levels. In the controlled event-aligned setting, restart with exact post-event history reconstruction yields substantially smaller errors than the baseline IMEX-BDF3 march and is consistent with the proposed mechanism of interface-induced history contamination. In the manufactured-benchmark diagnostic setting, the event-local diagnostic detector retains full recall with a sparse flagged set, and the manufactured-benchmark diagnostic realization improves the baseline at sufficiently resolved time steps. The stress test shows that sharper interfaces require stronger event resolution and can reveal a fixed-grid or local-correction error floor. These results identify the mechanism and the conditions that a future predictive implementation must satisfy.

The paper contributes an analytical framework and a diagnostic numerical study for interface-induced temporal order reduction. By identifying a localized correction principle and providing benchmark evidence for the history-contamination mechanism, the study clarifies the coupled roles of event localization, local bridge accuracy, and restart-history consistency in future interface-aware realizations.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/math14111948/s1>. The supplementary package includes the numerical data, custom Python scripts, requirements file, environment file, README and manifest instructions, generated numerical tables, reference outputs, high-resolution figure files, and figure-generation workflow used to reproduce the reported numerical results and figures.

Funding: This research was funded by the Deanship of Scientific Research at Northern Border University, Arar, Kingdom of Saudi Arabia, through project number “NBU-FPEJ-2026-1475-02”.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The author extends their appreciation to the Deanship of Scientific Research at Northern Border University, Arar, KSA, for funding this research work through the project number “NBU-FPEJ-2026-1475-02”.

Conflicts of Interest: The author declares no conflicts of interest. The funder had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Furzeland, R.M.; Verwer, J.G.; Zegeling, P.A. A numerical study of three moving-grid methods for one-dimensional partial differential equations which are based on the method of lines. *J. Comput. Phys.* **1990**, *89*, 349–388. [\[CrossRef\]](#)
2. Hawken, D.F. Review of adaptive-grid techniques for solution of partial differential equations. *Prog. Aerosp. Sci.* **1987**, *24*, 29–49. [\[CrossRef\]](#)
3. Du, Q.; Feng, X. The phase field method for geometric moving interfaces and their numerical approximations. *Handb. Numer. Anal.* **2020**, *21*, 425–508.
4. Liu, Z.L.; Huang, Y. Comparison of the Interface-Capturing Behavior between LSM and MCLS. *Key Eng. Mater.* **2017**, *730*, 533–540. [\[CrossRef\]](#)
5. Mullins, M.; Kamil, H.; Fahsi, A.; Soulaïmani, A. Physics-informed neural networks for solving moving interface flow problems using the level set approach. *Phys. Fluids* **2025**, *37*, 107124. [\[CrossRef\]](#)
6. Gassner, G.J.; Winters, A.R. A novel robust strategy for discontinuous Galerkin methods in computational fluid mechanics: Why? When? What? Where? *Front. Phys.* **2021**, *8*, 500690. [\[CrossRef\]](#)
7. Shu, C.W. High order WENO and DG methods for time-dependent convection-dominated PDEs: A brief survey of several recent developments. *J. Comput. Phys.* **2016**, *316*, 598–613. [\[CrossRef\]](#)
8. Xia, Y.; Xu, Y.; Shu, C.-W. Efficient time discretization for local discontinuous Galerkin methods. *Discret. Contin. Dyn. Syst.-Ser. B* **2007**, *8*, 677–693. [\[CrossRef\]](#)
9. Haynes, R.D. Recent Advances in Schwarz Waveform Moving Mesh Methods—A New Moving Subdomain Method. In *Domain Decomposition Methods in Science and Engineering XIX*; Huang, Y., Kornhuber, R., Widlund, O., Xu, J., Eds.; Lecture Notes in Computational Science and Engineering; Springer: Berlin/Heidelberg, Germany, 2011; Volume 78, pp. 253–260.
10. De La Hoz, F.; Vellido, F. Numerical simulations of time-dependent partial differential equations. *J. Comput. Appl. Math.* **2016**, *295*, 175–184. [\[CrossRef\]](#)
11. Luan, V.T. Efficient exponential Runge–Kutta methods of high order: Construction and implementation. *BIT Numer. Math.* **2021**, *61*, 535–560. [\[CrossRef\]](#)
12. Narayanamurthi, M.; Sandu, A. Efficient implementation of partitioned stiff exponential Runge–Kutta methods. *Appl. Numer. Math.* **2020**, *152*, 141–158. [\[CrossRef\]](#)
13. Xu, Z.; Cai, W.; Song, Y.; Wang, Y. Explicit high-order energy-preserving exponential time differencing method for nonlinear Hamiltonian PDEs. *Appl. Math. Comput.* **2021**, *404*, 126208. [\[CrossRef\]](#)
14. Myers, C.; Palmer, T.; Palmer, C. A hybrid Finite Volume-Smoothed Particle Hydrodynamics approach for shock capturing applications. *Comput. Methods Appl. Mech. Eng.* **2023**, *417*, 116412. [\[CrossRef\]](#)
15. Nguyen, V.D.; Kirchhelle, C.; Abdollahi, A.; García Grajales, J.A.; Li, D.; Benatia, K.; Gorbunov, K.; Bielle, S.; Goriely, A.; Jérusalem, A. A numerical framework coupling finite element and meshless methods in sequential and parallel simulations. *Finite Elem. Anal. Des.* **2023**, *219*, 103927. [\[CrossRef\]](#)
16. Bock, F.E.; Keller, S.; Huber, N.; Klusemann, B. Hybrid modelling by machine learning corrections of analytical model predictions towards high-fidelity simulation solutions. *Materials* **2021**, *14*, 1883. [\[CrossRef\]](#)

17. Chang, X.; Bodnar, B.; Grashorn, J.; Hild, F.; Gatuingt, F.; Roux, S. Simulation-guided PGD-stereocorrelation for hybrid testing. *Comput. Methods Appl. Mech. Eng.* **2025**, *435*, 117672. [\[CrossRef\]](#)
18. Chinesta, F.; El Khaldi, F.; Cueto, E. Hybrid twin: An intimate alliance of knowledge and data. In *The Digital Twin*; Crespi, N., Drobot, A.T., Minerva, R., Eds.; Springer International Publishing: Cham, Switzerland, 2023; pp. 279–298.
19. Chinesta, F.; Cueto, E.; Champaney, V.; Ghnatios, C.; Ammar, A.; Hascoët, N.; González, D.; Alfaro, I.; Di Lorenzo, D.; Pasquale, A.; et al. *A Gentle Introduction to Data, Learning, and Model Order Reduction: Techniques and Twinning Methodologies*; Studies in Big Data; Springer: Berlin/Heidelberg, Germany, 2025; Volume 174, pp. 1–227.
20. Jeon, J.; Lee, J.; Vinuesa, R.; Kim, S.J. Residual-based physics-informed transfer learning: A hybrid method for accelerating long-term CFD simulations via deep learning. *Int. J. Heat Mass Transf.* **2024**, *220*, 124900. [\[CrossRef\]](#)
21. Kumar, R.; Ezhilarasi, D. A state-of-the-art survey of model order reduction techniques for large-scale coupled dynamical systems. *Int. J. Dyn. Control* **2023**, *11*, 900–916. [\[CrossRef\]](#)
22. Novello, P.; Poëtte, G.; Lugato, D.; Peluchon, S.; Congedo, P.M. Accelerating hypersonic reentry simulations using deep learning-based hybridization (with guarantees). *J. Comput. Phys.* **2024**, *498*, 112700. [\[CrossRef\]](#)
23. Sancarlos, A.; Cameron, M.; Le Peuvedic, J.M.; Groulier, J.; Duval, J.L.; Cueto, E.; Chinesta, F. Learning stable reduced-order models for hybrid twins. *Data-Centric Eng.* **2021**, *2*, e10. [\[CrossRef\]](#)
24. Torregrosa, S.; Champaney, V.; Ammar, A.; Herbert, V.; Chinesta, F. Hybrid twins based on optimal transport. *Comput. Math. Appl.* **2022**, *127*, 12–24. [\[CrossRef\]](#)
25. Torregrosa, S.; Champaney, V.; Ammar, A.; Herbert, V.; Chinesta, F. Predicting high-fidelity data from coarse-mesh computational fluid dynamics corrected using hybrid twins based on optimal transport. *Mech. Ind.* **2024**, *25*, 31. [\[CrossRef\]](#)
26. Fish, A.C.; Reynolds, D.R.; Roberts, S.B. Implicit–explicit multirate infinitesimal stage-restart methods. *J. Comput. Appl. Math.* **2024**, *438*, 115534. [\[CrossRef\]](#)
27. Frank, J.; Hundsdorfer, W.; Verwer, J.G. On the stability of implicit-explicit linear multistep methods. *Appl. Numer. Math.* **1997**, *25*, 193–205. [\[CrossRef\]](#)
28. Kadioglu, S.Y. Analysis of the self-consistent IMEX method for tightly coupled non-linear systems. *J. Comput. Appl. Math.* **2017**, *322*, 148–160. [\[CrossRef\]](#)
29. Rodrigues, S.B.; Braga, G.B.F.; de Medeiros, M.A.F. The residual balanced IMEX decomposition for singly-diagonally-implicit schemes. *Appl. Numer. Math.* **2025**, *208*, 58–78. [\[CrossRef\]](#)
30. Zhang, H.; Sandu, A.; Blaise, S. Partitioned and implicit–explicit general linear methods for ordinary differential equations. *J. Sci. Comput.* **2014**, *61*, 119–144. [\[CrossRef\]](#)
31. Zhang, H.; Sandu, A.; Blaise, S. High order implicit-explicit general linear methods with optimized stability regions. *SIAM J. Sci. Comput.* **2016**, *38*, A1430–A1453. [\[CrossRef\]](#)
32. Zhang, H.; Sandu, A. Application of implicit-explicit general linear methods to reaction-diffusion problems. *AIP Conf. Proc.* **2015**, *1648*, 150017.
33. Cassandra, C.G. The event-driven paradigm for control, communication and optimization. *J. Control Decis.* **2014**, *1*, 3–17. [\[CrossRef\]](#)
34. Khulief, Y.A. Numerical modelling of impulsive events in mechanical systems. *Int. J. Model. Simul.* **2010**, *30*, 80–86.
35. Dehghan, M.; Abbaszadeh, M. Variational multiscale element free Galerkin (VMEFG) and local discontinuous Galerkin (LDG) methods for solving two-dimensional Brusselator reaction–diffusion system with and without cross-diffusion. *Comput. Methods Appl. Mech. Eng.* **2016**, *300*, 770–797. [\[CrossRef\]](#)
36. Mamonov, A.V.; Olshanskii, M.A. A priori analysis of a tensor ROM for parameter dependent parabolic problems. *SIAM J. Numer. Anal.* **2025**, *63*, 239–261. [\[CrossRef\]](#)
37. Quarteroni, A.; Rozza, G.; Manzoni, A. Certified reduced basis approximation for parametrized partial differential equations and applications. *J. Math. Ind.* **2011**, *1*, 3. [\[CrossRef\]](#)
38. Biswas, A.; Tian, J.; Ulusoy, S. Error estimates for deep learning methods in fluid dynamics. *Numer. Math.* **2022**, *151*, 753–777. [\[CrossRef\]](#)
39. Cen, J.; Zou, Q. Deep finite volume method for partial differential equations. *J. Comput. Phys.* **2024**, *517*, 113307. [\[CrossRef\]](#)
40. Dong, H.; Zhang, J.; Yang, J.; Jiao, Y.; Zhang, S. Dual-Activation PINNs: Frequency–Gradient Synergy via Growing Cosine Unit and Non-Monotonic Cubic Unit. In *Proceedings of the 2nd International Conference on Intelligent Computing and Data Analysis*; Association for Computing Machinery: New York, NY, USA, 2025; pp. 902–906.
41. Guo, J.; Zhu, H.; Yang, Y.; Guo, C. Advances in physics-informed neural networks for solving complex partial differential equations and their engineering applications: A systematic review. *Eng. Appl. Artif. Intell.* **2025**, *161*, 112044. [\[CrossRef\]](#)
42. Jagtap, A.D.; Kawaguchi, K.; Em Karniadakis, G. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2020**, *476*, 20200334. [\[CrossRef\]](#)
43. Kiyani, E.; Shukla, K.; Urbán, J.F.; Darbon, J.; Karniadakis, G.E. Optimizing the optimizer for physics-informed neural networks and Kolmogorov–Arnold networks. *Comput. Methods Appl. Mech. Eng.* **2025**, *446*, 118308. [\[CrossRef\]](#)

44. Lyu, L.; Zhang, Z.; Chen, M.; Chen, J. MIM: A deep mixed residual method for solving high-order partial differential equations. *J. Comput. Phys.* **2022**, *452*, 110930. [\[CrossRef\]](#)
45. Ngueabou, Y.V.; Olonijou, S.D. Integrating Spectral Methods with Neural Network Architectures: A Review of Hybrid Approaches to Solving Differential Equation. *Arch. Comput. Methods Eng.* **2025**, *33*, 5643–5699. [\[CrossRef\]](#)
46. Pratama, D.A.; Bakar, M.A.; Ismail, N.B.; Mashuri, M. ANN-based methods for solving partial differential equations: A survey. *Arab J. Basic Appl. Sci.* **2022**, *29*, 233–248. [\[CrossRef\]](#)
47. Ren, Z.; Zhou, S.; Liu, D.; Liu, Q. Physics-informed neural networks: A review of methodological evolution, theoretical foundations, and interdisciplinary frontiers toward next-generation scientific computing. *Appl. Sci.* **2025**, *15*, 8092. [\[CrossRef\]](#)
48. Shin, Y.; Darbon, J.; Karniadakis, G.E. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. *Commun. Comput. Phys.* **2020**, *28*, 2042–2074. [\[CrossRef\]](#)
49. Wang, T.; Liu, G.; Li, E.; Xu, X. Adaptive deep physics-informed neural network with dual-nested activation for solving complex partial differential equations. *Comput. Methods Appl. Mech. Eng.* **2025**, *444*, 118125. [\[CrossRef\]](#)
50. Yan, D.; He, L. DP-PINN+: A Dual-Phase PINN learning with automated phase division. *J. Comput. Sci.* **2025**, *90*, 102637. [\[CrossRef\]](#)
51. Zhao, Z.; Wang, Y.; Zhang, W.; Ba, Z.; Sun, L. Physics-informed neural networks in heat transfer-dominated multiphysics systems: A comprehensive review. *Eng. Appl. Artif. Intell.* **2025**, *157*, 111098. [\[CrossRef\]](#)
52. Akbari, S.; Pawar, S.; San, O. Numerical assessment of a nonintrusive surrogate model based on recurrent neural networks and proper orthogonal decomposition: Rayleigh–Bénard convection. *Int. J. Comput. Fluid Dyn.* **2022**, *36*, 599–617. [\[CrossRef\]](#)
53. Brivio, S.; Fresca, S.; Manzoni, A. PTPI-DL-ROMs: Pre-trained physics-informed deep learning-based reduced order models for nonlinear parametrized PDEs. *Comput. Methods Appl. Mech. Eng.* **2024**, *432*, 117404. [\[CrossRef\]](#)
54. Brivio, S.; Fresca, S.; Manzoni, A. Handling geometrical variability in nonlinear reduced order modeling through continuous geometry-aware DL-ROMs. *Comput. Methods Appl. Mech. Eng.* **2025**, *442*, 117989. [\[CrossRef\]](#)
55. Chen, W.; Wang, Q.; Hesthaven, J.S.; Zhang, C. Physics-informed machine learning for reduced-order modeling of nonlinear problems. *J. Comput. Phys.* **2021**, *446*, 110666. [\[CrossRef\]](#)
56. Cicci, L.; Fresca, S.; Manzoni, A. Deep-HyROMnet: A deep learning-based operator approximation for hyper-reduction of nonlinear parametrized PDEs. *J. Sci. Comput.* **2022**, *93*, 57. [\[CrossRef\]](#)
57. Cicci, L.; Fresca, S.; Manzoni, A.; Quarteroni, A. Efficient approximation of cardiac mechanics through reduced-order modeling with deep learning-based operator approximation. *Int. J. Numer. Methods Biomed. Eng.* **2024**, *40*, e3783. [\[CrossRef\]](#)
58. Valenzuela del Río, J.E.; Ramamurthy, A. Differentiable partial differential equation solvers and their application to quasi real-time inverse dynamical problems on thermal management. In *AIAA SCITECH 2024 Forum*; AIAA: Reston, VA, USA, 2024; p. 1628.
59. Dutta, S.; Rivera-Casillas, P.; Farthing, M.W. Neural ordinary differential equations for data-driven reduced order modeling of environmental hydrodynamics. *CEUR Workshop Proc.* **2021**, *2964*, 209.
60. Franco, N.R.; Fraulin, D.; Manzoni, A.; Zunino, P. On the latent dimension of deep autoencoders for reduced order modeling of PDEs parametrized by random fields. *Adv. Comput. Math.* **2024**, *50*, 96. [\[CrossRef\]](#)
61. Gowrachari, H.; Demo, N.; Stabile, G.; Rozza, G. Non-intrusive model reduction of advection-dominated hyperbolic problems using neural network shift augmented manifold transformation. *Comput. Fluids* **2025**, *300*, 106758. [\[CrossRef\]](#)
62. Nelsen, N.H.; Stuart, A.M. Operator learning using random features: A tool for scientific computing. *SIAM Rev.* **2024**, *66*, 535–571. [\[CrossRef\]](#)
63. Puri, V.; Prakash, A.; Kara, L.B.; Zhang, Y.J. SNF-ROM: Projection-based nonlinear reduced order modeling with smooth neural fields. *J. Comput. Phys.* **2025**, *532*, 113957. [\[CrossRef\]](#)
64. Berman, J.; Schwerdtner, P.; Peherstorfer, B. Neural Galerkin schemes for sequential-in-time solving of partial differential equations with deep networks. In *Handbook of Numerical Analysis*; Elsevier: Amsterdam, The Netherlands, 2024; Volume 25, pp. 389–418.
65. Fischer, C.; Bai, Y.; Fu, K.; McClanahan, N.; Kimn, J.H.; Doom, J.J. The Local Deep Galerkin Method Applied to the (2+1)-Dimensional Navier-Stokes and Cahn-Hilliard Equations. In *AIAA SCITECH 2026 Forum*; AIAA: Reston, VA, USA, 2026; p. 1935.
66. Gu, Y.; Ng, M.K. Deep adaptive basis Galerkin method for high-dimensional evolution equations with oscillatory solutions. *SIAM J. Sci. Comput.* **2022**, *44*, A3130–A3157. [\[CrossRef\]](#)
67. Kast, M.; Hesthaven, J.S. Positional embeddings for solving PDEs with evolutionary deep neural networks. *J. Comput. Phys.* **2024**, *508*, 112986. [\[CrossRef\]](#)
68. Kumar, H.; Yadav, N. Approximate Solution of Hyperbolic Telegraphic Equation Using Modified Deep Galerkin Method. *Proc. Natl. Acad. Sci. India Sect. A Phys. Sci.* **2025**, *95*, 525–539. [\[CrossRef\]](#)
69. Wen, Y.; Vanden-Eijnden, E.; Peherstorfer, B. Coupling parameter and particle dynamics for adaptive sampling in Neural Galerkin schemes. *Phys. D Nonlinear Phenom.* **2024**, *462*, 134129. [\[CrossRef\]](#)

70. Yang, X.; Ran, M.; Pu, H. Weak neural Galerkin method for nonlinear time-dependent partial differential equations. *Phys. Lett. A* **2025**, *554*, 130757. [[CrossRef](#)]
71. Li, Z.; Zheng, H.; Kovachki, N.; Jin, D.; Chen, H.; Liu, B.; Azizzadenesheli, K.; Anandkumar, A. Physics-informed neural operator for learning partial differential equations. *ACM/IMS J. Data Sci.* **2024**, *1*, 1–27. [[CrossRef](#)]
72. Chassagneux, J.-F.; Chen, J.; Frikha, N.; Zhou, C. A learning scheme by sparse grids and Picard approximations for semilinear parabolic PDEs. *IMA J. Numer. Anal.* **2023**, *43*, 3109–3168. [[CrossRef](#)]
73. Allen, S.M.; Cahn, J.W. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metall.* **1979**, *27*, 1085–1095. [[CrossRef](#)]
74. Lions, J.-L.; Magenes, E. *Non-Homogeneous Boundary Value Problems and Applications*; Springer-Verlag: Berlin, Germany; Heidelberg, Germany; New York, NY, USA, 1972; Volume I–II.
75. Ascher, U.M.; Ruuth, S.J.; Wetton, B.T.R. Implicit-explicit methods for time-dependent partial differential equations. *SIAM J. Numer. Anal.* **1995**, *32*, 797–823. [[CrossRef](#)]
76. Ascher, U.M.; Ruuth, S.J.; Spiteri, R.J. Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations. *Appl. Numer. Math.* **1997**, *25*, 151–167. [[CrossRef](#)]
77. Tyson, W.C.; Roy, C.J. A higher-order error estimation framework for finite-volume CFD. *J. Comput. Phys.* **2019**, *394*, 632–657. [[CrossRef](#)]
78. Strassmann, D.; Reinert, D.; Dipankar, A. Truncation error analysis for the piecewise parabolic method. *Q. J. R. Meteorol. Soc.* **2025**, *151*, e5074. [[CrossRef](#)]
79. Cheng, Q.; Liu, C.; Shen, J. A new interface capturing method for Allen-Cahn type equations based on a flow dynamic approach in Lagrangian coordinates, I. One-dimensional case. *J. Comput. Phys.* **2020**, *419*, 109509. [[CrossRef](#)]
80. Storvik, E.; Bringedal, C. Sharp-interface limit of the Cahn-Hilliard-Biot equations. *Appl. Math. Lett.* **2025**, *166*, 109522. [[CrossRef](#)]
81. Dirac, P.A.M. *The Principles of Quantum Mechanics*; Oxford University Press: Oxford, UK, 1930.
82. Frenkel, J. *Wave Mechanics: Advanced General Theory*; Clarendon Press: Oxford, UK, 1934.
83. Osher, S.; Sethian, J.A. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *J. Comput. Phys.* **1988**, *79*, 12–49. [[CrossRef](#)]
84. Gårding, L. Dirichlet's problem for linear elliptic partial differential equations. *Math. Scand.* **1953**, *1*, 55–72. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.