KMF: PHYSICS-GUIDED INFERENCE-TIME DEBIASING FOR OFF-CADENCE PDE FOUNDATION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Scientific workflows require PDE states on temporal grids chosen by downstream analysis, often at a finer resolution than the cadence of available foundation-model outputs or observations. Direct off-cadence model queries can be biased, while standard temporal upsampling ignores the governing dynamics. We introduce Kinematic Manifold Fusion (KMF), a parameter-update-free, physics-guided inference-time correction framework for temporal refinement and autoregressive rollout. When neighboring model snapshots are available, KMF evaluates a cheap spatial PDE operator at those snapshots, constructs a local cubic Hermite physical candidate, and fuses it with a model estimate through a calibration-selected scalar. During rollout, it reuses the same local physics as a causal defect signal before reinjecting the corrected state into the frozen model. It also enforces known constraints through the corresponding exact projection when available. KMF changes no model parameters and uses zero temporal integration substeps, only instantaneous spatial right-hand-side evaluations. Under standard smoothness and operator-error assumptions, we establish a local error bound for the physical bridge and an endpoint-perturbation bound for deployment. Empirical evaluation across PDE systems, temporal cadences, and frozen foundation-model architectures demonstrates accurate temporal refinement and active causal correction in the regimes where local physics is informative, while retaining the efficiency of the underlying model. KMF offers a lightweight complement to learned continuous-time solvers and numerical integration.

1 INTRODUCTION

Machine learning foundation models for partial differential equations (PDEs) accelerate scientific prediction by learning from diverse simulated physical systems before deployment (Li et al., 2021; Lu et al., 2021; Takamoto et al., 2022). Once pretrained, they support fast repeated queries for design, control, data assimilation, and scientific analysis. Yet workflows often require states at times that do not match the model's reliable or training cadence. We address this gap with a post-training, parameter-update-free inference framework that uses the frozen model together with local PDE information.

This temporal mismatch appears in two settings. In the first, successive foundation-model states $\tilde{u}_k$ and $\tilde{u}_{k+1} = f_\theta(\tilde{u}_k; \Delta t)$ define a coarse interval, while analysis, data alignment, visualization, or coupling to another simulator requires an intermediate state. We refer to this setting as *temporal refinement*. Some foundation models, including Poseidon, expose a variable lead-time or custom-cadence input (Herde et al., 2024), but these off-cadence queries may be less accurate than predictions at the cadences emphasized during training. Linear interpolation between successive model states ignores nonlinear advection, phase evolution, and local temporal curvature.

In the second setting, the model is used autoregressively, with each generated endpoint becoming the input to the next step. The resulting sequence supports refinement between model steps, while small prediction errors can accumulate recursively.

A spatial PDE operator maps each state to an instantaneous temporal rate through spatial derivatives, forcing, and constitutive terms. Evaluating it is generally much cheaper than forward integration and supplies local information about trajectory motion and curvature.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We introduce **Kinematic Manifold Fusion** (KMF), a post-training inference-time framework that combines local PDE kinematics, known state constraints, and a frozen foundation model. We use *inference-time debiasing* operationally: KMF corrects a model estimate using only information permitted by the declared endpoint or rollout contract. In Task 1, successive model endpoints and their PDE-derived rates define a closed-form physical candidate, optionally fused with a direct intermediate model estimate. In Task 2, the foundation model remains the causal predictor: KMF applies a physical-defect correction and constraint projection to each candidate endpoint before reinjecting it into the next model call. Fusion and correction strengths are selected on earlier, trajectory-disjoint reference trajectories and then fixed for evaluation. Thus KMF requires no foundation-model retraining or auxiliary learned network; its only data-dependent adaptation is one or two low-dimensional calibration scalars. It uses zero temporal integration substeps, only instantaneous spatial right-hand-side evaluations. For periodic incompressible velocity fields, the Fourier Helmholtz–Leray projection supplies an exact constraint operation; for other systems, bounds or filtering are treated as admissibility operations rather than exact PDE invariants.

The framework integrates physical motion, manifold structure, and calibrated fusion. Its individual numerical ingredients are classical; the contribution is their formulation as a small, architecture-agnostic inference layer for frozen PDE predictors, including their causal reuse during rollout. Under standard smoothness assumptions and bounded endpoint-operator error, we establish local bridge and deployed endpoint-perturbation bounds, and use exact projection results where the associated constraint manifold is known.

Specifically, we contribute:

- **A physics-guided inference-time framework for frozen PDE models.** KMF uses the learned model as the primary predictor and local PDE evaluations as correction signals, combining kinematic physical candidates, exact constraints where available, and calibrated fusion for temporal refinement and causal rollout correction.
- **A zero-temporal-substep physical candidate with deployed error control.** We construct a closed-form endpoint Hermite bridge from two model states and their instantaneous PDE rates of change, establish a fourth-order intrinsic interpolation bound, and quantify additive endpoint and operator perturbations in deployment.
- **Evaluation across models, systems, and operating regimes.** We evaluate temporal refinement and rollout over frozen model interfaces, multiple PDE systems, and requested cadences; retain unavailable interfaces and negative transfer rows; test calibration resampling and long horizons; and compare with residual-gradient, Gauss–Newton, cached-Jacobian, projection, and fixed-budget flow controls.

2 Related Work

Physics foundation models. PDE foundation models learn reusable operators or trajectory distributions from collections of simulated systems. Representative systems include Poseidon, DPOT, MORPH, and The Well (Herde et al., 2024; Hao et al., 2024; Rautela et al., 2025; Ohana et al., 2024). Fourier operators, DeepONet, PDEBench, temporal operators, and differential–integral models broaden this landscape (Li et al., 2021; Lu et al., 2021; Takamoto et al., 2022; Diab & Al Kobaisi, 2025; Wu et al., 2026). These methods primarily address how to train or architect the predictor. KMF is orthogonal: it leaves the pretrained predictor fixed and supplies a physics-guided inference layer for requested times and rollout states.

Post-training and inference-time physical correction. A second line of work improves a pretrained predictor after training by enforcing physical structure at inference. Physics-informed neural networks optimize a parameterized state or solution representation using PDE residuals (Raissi et al., 2019), while constrained neural ODEs and symmetry-aware dynamics impose structure architecturally or through the evolution map (Chen et al., 2018; Wang et al., 2021). Generative flow-marching and neural ODE approaches instead learn continuous-time trajectory mechanisms (Chen & Deng, 2025; Chen et al., 2018). KMF occupies a different contract: it leaves the predictor frozen, evaluates the available spatial operator only at candidate states, and performs neither auxiliary training nor temporal PDE integration.

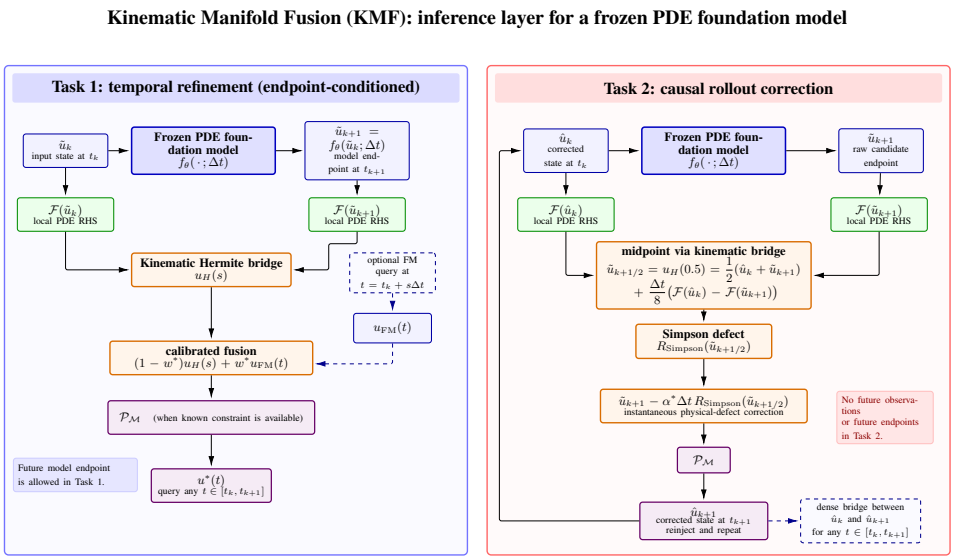

Figure 1: Kinematic Manifold Fusion (KMF) augments a frozen PDE foundation model with local physics-guided temporal refinement and causal rollout correction, without temporal integration substeps or parameter updates.

Post-hoc rollout correction. Methods such as history-enriched rollout training and diffusion-based refinement target autoregressive drift and error accumulation (Zhang et al., 2026; Lippe et al., 2023). PhysicsCorrect applies a cached-Jacobian correction to neural PDE solver steps (Huang & Perdikaris, 2026). These approaches are closest to KMF for causal rollout correction, but use different state representations and correction contracts. KMF uses endpoint right-hand-side evaluations and a Hermite–Simpson defect, not a cached Jacobian, and additionally supplies continuous values inside each model interval. We evaluate the closest transferable cached-Jacobian procedure under the declared reduced-state contract as supporting, not definitive, comparative evidence; endpoint-conditioned refinement remains a distinct information setting.

Contract perspective. KMF is distinguished by a compact deployment contract: the predictor is frozen; no new network is trained; local RHS evaluations are used but not integrated forward in time; a future model endpoint is used only in the endpoint-conditioned task; and corrected rollout endpoints are reinjected causally in Task 2. This positions KMF between pure endpoint interpolation, which lacks a physical correction signal, and numerical or optimization-based solvers, which generate a trajectory by integrating or iteratively optimizing the state.

3 Problem Formulation and Mathematical Motivation

We consider a spatial domain $\Omega \subset \mathbb{R}^d$ with periodic or Dirichlet boundary conditions, and a dynamical state $u(x,t) \in \mathcal{X}$ governed by a general autonomous partial differential equation:

$$\frac{\partial u}{\partial t} = F(u), \quad u(\cdot,0) = u_0, \tag{1}$$

where $F : \mathcal{X} \to \mathcal{X}$ is a spatial differential operator (e.g., advection-diffusion-reaction). In addition, physical solutions often satisfy an algebraic or geometric invariant manifold constraint:

$$u(t) \in \mathcal{M} = \{v \in \mathcal{X} : \mathcal{C}(v) = 0\}, \tag{2}$$

such as the solenoidal condition $\nabla \cdot u = 0$ for incompressible fluid dynamics.

3.1 Two Distinct Operational Tasks

We formalize two distinct operational objectives encountered in continuous-time surrogate modeling. In **Task 1 (Foundation-Model Temporal Refinement)**, successive frozen-model states $\tilde{u}_k$ and $\tilde{u}_{k+1} = f_\theta(\tilde{u}_k; \Delta t)$ define the interval. The objective is to estimate the state at an arbitrary intermediate query time $t = t_k + s\Delta t$, $s \in (0,1)$, using only $\tilde{u}_k$, $\tilde{u}_{k+1}$, the known spatial PDE operator,

and, when the interface permits it, a direct model query at t. Thus, no future reference endpoint enters the deployed refinement calculation. In **Task 2 (Continuous-Time and Defect-Corrected Autoregressive Rollout)**, the model is initialized solely from u_0 and must causally forecast future states $\hat{u}_k \approx u(t_0 + k\Delta t)$ $(k = 1, \ldots, K)$ across an extended horizon at coarse sampling cadences $\Delta t \in [0.05\,\mathrm{s}, 0.40\,\mathrm{s}]$ without future observations. Task 2 corrects coarse temporal stepping errors through instantaneous physical-curvature defects using zero temporal integration substeps, and supplies dense sub-cadence evaluations at arbitrary $t \in [t_k, t_{k+1}]$ throughout the rollout horizon.

The two alternatives most relevant to this setting are residual-penalized test-time optimization and forward numerical integration. Their computational and stability considerations are summarized in Appendix D.1; the scoped RK4 modal example appears in Appendix E.6. KMF targets the complementary zero-forward-substep setting, using instantaneous spatial evaluations and an algebraic bridge or correction.

4 The KMF Framework

KMF is not a numerical solver: the frozen model supplies the endpoint transition, and the PDE is queried only at candidate states for local kinematics, a defect, or a constraint operation. It is also more than interpolation: the physical bridge is one candidate that can be fused with a model estimate and whose rollout correction is causally reinjected. Local endpoint physics is inexpensive, scalar fusion preserves the stronger candidate, and reinjection lets endpoint corrections influence later calls.

4.1 The Two-Point Kinematic Hermite Bridge

In Task 1, KMF first constructs a *kinematic physical bridge* for an interval generated by the foundation model. The successive endpoints are $\tilde{u}_k$ and $\tilde{u}_{k+1} = f_\theta(\tilde{u}_k; \Delta t)$. For an autonomous PDE $\partial_t u = F(u)$, the spatial differential operator supplies an instantaneous rate at either generated state. On a spatial grid, evaluating F uses only spatial operations, such as finite-difference stencils or pseudo-spectral derivatives on periodic domains. KMF evaluates $F(\tilde{u}_k)$ and $F(\tilde{u}_{k+1})$, without temporal numerical integration, and uses the resulting state-rate pairs as four osculatory conditions. The technical realization is the classical cubic Hermite bridge $u_H(s) \in \mathcal{P}_3$ (Richard et al., 2016; Hairer et al., 1993). On normalized interval $s = (t - t_k)/\Delta t \in [0, 1]$, its basis functions are:

$$h_{00}(s) = 1 - 3s^2 + 2s^3, \qquad h_{10}(s) = 3s^2 - 2s^3, \tag{3}$$

$$h_{01}(s) = \Delta t(s - 2s^2 + s^3), \qquad h_{11}(s) = \Delta t(-s^2 + s^3). \tag{4}$$

The kinematic Hermite bridge $u_H(s)$ is evaluated analytically as:

$$u_H(s) = h_{00}(s)\tilde{u}_k + h_{10}(s)\tilde{u}_{k+1} + h_{01}(s)F(\tilde{u}_k) + h_{11}(s)F(\tilde{u}_{k+1}). \tag{5}$$

Evaluating Equation 5 requires only two spatial evaluations of $F(\cdot)$ at generated interval endpoints, without taking any temporal integration steps. Such finite-difference or pseudo-spectral operator evaluations are standard discrete spatial calculations (Richard et al., 2016; Trefethen, 2000). The operator evaluations are exact evaluations of the chosen discrete operator at their inputs; they are not assumed to equal the unknown true trajectory derivative when the model endpoint is inaccurate.

The following approximation result is an *intrinsic physical-bridge result*: it assumes exact solution endpoints and characterizes the local order of the kinematic construction, not a standalone guarantee for deployed KMF. The deployed bound immediately below propagates model-endpoint and operator errors, and fusion is evaluated empirically.

Theorem 1 (Local Hermite Interpolation Error Bound)**.** *Let $\Omega \subset \mathbb{R}^d$ and $\mathcal{X} = L^p(\Omega)$, $p \in [1, \infty]$. Assume that, for almost every $x \in \Omega$, $t \mapsto u(t, x)$ is four times continuously differentiable on $[t_0, t_1]$ and that $G(x) = \sup_{\tau \in [t_0, t_1]} |\partial_t^4 u(\tau, x)|$ belongs to $L^p(\Omega)$. Then the local interpolation error of u_H satisfies:*

$$\|u(t) - u_H(s)\|_{L^p(\Omega)} \leq \frac{\Delta t^4}{384} \|G\|_{L^p(\Omega)}, \quad \forall s \in [0, 1]. \tag{6}$$

Because $\phi(s)$ achieves its global maximum on $[0, 1]$ at midpoint $s = 0.5$ with $\phi(0.5) = 1/16$, evaluating the bridge at $s = 0.5$ tests the framework at its loosest theoretical error envelope.

Deployed endpoint perturbation. The theorem isolates the intrinsic interpolation error of the physical candidate. To connect it to deployed KMF, let $\tilde{u}_k$ and $\tilde{u}_{k+1}$ be the model endpoints, let $u_k = u(t_k)$ and $u_{k+1} = u(t_{k+1})$ be the corresponding exact states, and define $\epsilon_{\text{end}} = \max\{\|\tilde{u}_k - u_k\|, \|\tilde{u}_{k+1} - u_{k+1}\|\}$. If the endpoint operator evaluations have error at most ϵ_F and F is locally Lipschitz on the relevant neighborhood, the same coefficient-wise perturbation argument gives

$$\|u(t) - \tilde{u}_H(s)\|_{\mathcal{X}} \le C_H(s)\Delta t^4 + C_{\text{end}}(s)\epsilon_{\text{end}} + C_F(s)\Delta t\,\epsilon_F, \tag{7}$$

where the constants depend only on the local smoothness and Lipschitz bounds and on the normalized query location s. Thus, the fourth-order term is not asserted to be the entire deployed error: endpoint forecast error and operator mismatch enter additively. In the native-cadence fluid experiments, the foundation-model endpoint is close to the reference trajectory, so the endpoint term is empirically small enough for the intrinsic bridge contribution to remain practically relevant. The held-out Task 1 results evaluate the full perturbed-endpoint estimator rather than the exact-endpoint theorem in isolation.

4.2 Generalized Manifold Invariant Projection

Let $\mathcal{X}$ be the ambient Hilbert space of spatial fields defined on domain $\Omega \subset \mathbb{R}^d$ (e.g., $\mathcal{X} = L^2(\Omega)^d$ for velocity fields). When a known physical constraint is represented by an operator $\mathcal{C} : \mathcal{X} \to \mathcal{Y}$, define the associated constrained set:

$$\mathcal{M} = \{v \in \mathcal{X} : \mathcal{C}(v) = 0\}, \tag{8}$$

For an invariant constraint, true trajectories initialized in $\mathcal{M}$ remain in it. Candidate predictions may violate such constraints. KMF therefore applies the following metric projection whenever the constraint set and its projection are known:

$$\mathcal{P}_{\mathcal{M}}(v) = \arg\min_{u \in \mathcal{X}} \frac{1}{2}\|u - v\|_{\mathcal{X}}^2 \quad \text{subject to} \quad \mathcal{C}(u) = 0. \tag{9}$$

Depending on the mathematical geometry of constraint $\mathcal{C}(u)$, Equation 9 admits exact closed-form solutions across three fundamental physical families:

Class 1: Linear Differential and Conservation Laws. For linear balance laws $\mathcal{C}(u) = \mathcal{A}u - b = 0$, the first-order optimality condition yields the orthogonal Green's projector $\mathcal{P}_{\mathcal{M}}(v) = v - \mathcal{A}^*(\mathcal{A}\mathcal{A}^*)^{-1}(\mathcal{A}v - b)$. In translation-invariant domains, differential operators diagonalize in Fourier space. This class covers solenoidal divergence-free constraints in incompressible Navier-Stokes, Stokes flows, and magnetohydrodynamics ($\nabla \cdot u = 0$), Gauss's laws in Maxwell electrodynamics ($\nabla \cdot \mathbf{B} = 0$), total mass and scalar conservation ($\int_\Omega u\,dx = M$), and potential flow circulation ($\oint u \cdot dl = \Gamma$). For fluids on a periodic torus $\Omega = [0, L]^d$, incompressibility requires $\nabla \cdot u = 0$. Under the Helmholtz-Hodge decomposition $v = v_{\text{div-free}} + \nabla\phi$, non-zero divergence acts as virtual mass sources driving rollout drift. The Fourier symbol of $-\nabla(-\nabla^2)^{-1}(\nabla\cdot)$ yields the orthogonal Helmholtz-Leray projection (Chorin, 1968):

$$\widehat{\mathcal{P}_{\mathcal{M}}(u)}(\mathbf{k}) = \left(\mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\right)\widehat{u}(\mathbf{k}) \quad (\mathbf{k} \neq \mathbf{0}), \qquad \widehat{\mathcal{P}_{\mathcal{M}}(u)}(\mathbf{0}) = \widehat{u}(\mathbf{0}), \tag{10}$$

where $\widehat{u}(\mathbf{k}) \in \mathbb{C}^d$ is the Fourier velocity vector at wavenumber $\mathbf{k} \in \mathbb{Z}^d$, and $\mathbf{I} - \mathbf{k}\mathbf{k}^T/\|\mathbf{k}\|^2$ orthogonally projects $\widehat{u}(\mathbf{k})$ onto the solenoidal plane perpendicular to $\mathbf{k}$, eliminating spurious divergence modes in $\mathcal{O}(N^d \log N)$ time via fast Fourier transforms (< 0.03 ms on GPU).

Proposition 2 (Non-Expansion of Manifold Projection)**.** *Let $\mathcal{P}_{\mathcal{M}}$ be the orthogonal projection onto a closed* linear *subspace $\mathcal{M}$ of a Hilbert space $\mathcal{X}$. For any target state $u^* \in \mathcal{M}$ and candidate $v \in \mathcal{X}$:*

$$\|\mathcal{P}_{\mathcal{M}}(v) - u^*\| \le \|v - u^*\|. \tag{11}$$

Equality holds if and only if $v - u^* \in \mathcal{M}$ (equivalently $v \in \mathcal{M}$).

For periodic incompressible velocity fields, Proposition 2 applies exactly to the Fourier Helmholtz–Leray projector in L^2: it removes the compressible component and cannot increase L^2 error to a divergence-free reference. Metric projections onto closed convex sets are also non-expansive (Bauschke & Combettes, 2017), but the Pythagorean identity used here is specific to linear subspaces. We therefore treat clipping and spectral filtering in the non-fluid datasets as dataset-specific admissibility operations and evaluate their effect on forecast error.

4.3 CALIBRATED CONSENSUS FUSION

Although the kinematic physical bridge $u_H(s)$ uses a coarse model endpoint and local physical rates, it can miss fine-scale structures over wider intervals. Conversely, a direct off-cadence foundation-model query $u_{\text{FM}}(t)$ can preserve rich spatial structure but may exhibit cadence-dependent phase error. Fusion is therefore an explicit question of whether local PDE kinematics add information beyond the frozen model, rather than an assumption that either candidate must dominate. When the interface exposes a direct query, KMF synthesizes the two deployable estimates through a convex consensus:

$$u^*(t) = \mathcal{P}_{\mathcal{M}}\Big((1-w^*)u_H(s) + w^* u_{\text{FM}}(t)\Big), \quad w^* \in [0,1]. \tag{12}$$

The complete temporal-refinement procedure is given in Appendix Algorithm 2. The main text focuses on the information contract and the resulting estimator rather than repeating implementation pseudocode.

To determine w^* without test-time label leakage, we use $M = 5$ historical reference trajectories, disjoint from the $N_{\text{test}} = 50$ held-out trajectories. Calibration uses only its own reference intermediate states; no test trajectory, label, or error is available during selection. The bridge and, where available, direct off-cadence FM query are fit to these references through one global scalar. This is calibration-time supervision, not neural training: no foundation-model parameter or spatially varying gate is learned. The scalar design keeps KMF calibration-light and architecture-agnostic. We formulate it as:

$$w^* = \arg\min_{w\in[0,1]} \sum_{m=1}^{M} \left\| (1-w)u_{H,\text{cal}}^{(m)} + w u_{\text{FM, cal}}^{(m)} - u_{\text{true, cal}}^{(m)} \right\|^2. \tag{13}$$

Defining difference vectors $\Delta_{\text{FM}}^{(m)} = u_{\text{FM, cal}}^{(m)} - u_{H,\text{cal}}^{(m)}$ and $\Delta_{\text{true}}^{(m)} = u_{\text{true, cal}}^{(m)} - u_{H,\text{cal}}^{(m)}$, the closed-form ordinary least-squares solution pooled across all M calibration trajectories is:

$$w^* = \text{clip}\left(\frac{\sum_{m=1}^{M} \langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle}{\sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2},\ 0.0,\ 1.0 \right), \qquad \sum_m \|\Delta_{\text{FM}}^{(m)}\|^2 > 0, \tag{14}$$

When the denominator is zero, the candidates coincide on calibration and we set $w^* = 0$. The implementation adds a small floating-point safeguard only for this degenerate case. Appendix E.4 shows that the displayed non-degenerate expression is the constrained empirical-risk minimizer and hence cannot underperform either endpoint *on that calibration objective*. This deterministic statement gives no finite-sample or held-out guarantee; all generalization claims in this paper are empirical and use trajectory-disjoint calibration and test splits.

4.3.1 UNSUPERVISED PHYSICS-GATED CONSENSUS

When calibration trajectories are unavailable, we optionally use a residual-informed heuristic weight. The Gauss–Markov precision rule is optimal only under a specified error model (Rao, 2009); a PDE defect is not generally a calibrated estimate of state-error variance. We therefore use the following quantity as a heuristic diagnostic rather than a theorem-backed uncertainty estimate. Evaluating the continuous Hermite collocation defect at $s = 0.5$ gives:

$$\mathcal{R}_{\text{Simpson}}(u) = \frac{\tilde{u}_1 - u_0}{\Delta t} - \frac{1}{6}\left(F(u_0) + 4F(u) + F(\tilde{u}_1)\right). \tag{15}$$

The residual-informed weight is:

$$w_{\text{phys}}(u) = \frac{\|\mathcal{R}_{\text{Simpson}}(u_{\text{FM}})\|^{-2}}{\|\mathcal{R}_{\text{Simpson}}(u_H)\|^{-2} + \|\mathcal{R}_{\text{Simpson}}(u_{\text{FM}})\|^{-2}}. \tag{16}$$

It favors the candidate with lower measured midpoint defect. This particular Simpson construction is not an arbitrary-s residual; off-center experiments use the calibration-selected fusion or the bridge, not this unsupervised gate. Its utility must be established empirically for each PDE/discretization; it is not used to select the supervised benchmark headline results. Further discussion is given in Appendix E.5.

4.4 Continuous-Time and Defect-Corrected Autoregressive Rollout

For continuous-time autoregressive rollout (Task 2), future observations are unavailable and predictions must advance causally from u_0. KMF addresses recursive temporal error accumulation by applying closed-loop physical curvature defect relaxation at every autoregressive step:

$$\hat{u}_{k+1} = \mathcal{P}_{\mathcal{M}}\Big(\tilde{u}_{k+1} - \alpha^* \Delta t\, \mathcal{R}_{\text{Simpson}}(\tilde{u}_{k+1/2})\Big), \tag{17}$$

where $\tilde{u}_{k+1} = f_\theta(\hat{u}_k)$ is the raw candidate step from the foundation model. The tentative midpoint state $\tilde{u}_{k+1/2}$ is evaluated algebraically with zero temporal integration substeps by querying the kinematic bridge at $s = 0.5$ between $\hat{u}_k$ and $\tilde{u}_{k+1}$: $\tilde{u}_{k+1/2} = u_H(0.5) = \frac{1}{2}(\hat{u}_k + \tilde{u}_{k+1}) + \frac{\Delta t}{8}(F(\hat{u}_k) - F(\tilde{u}_{k+1}))$. This has the algebraic form of the classical Lobatto IIIA collocation midpoint (Hairer et al., 1993); KMF does *not* solve the collocation equations, and the endpoint remains the FM output. It uses only boundary spatial operator evaluations $F(\hat{u}_k)$ and $F(\tilde{u}_{k+1})$ (Appendix E.8). Multiplying defect $\mathcal{R}_{\text{Simpson}}(\tilde{u}_{k+1/2})$ by Δt ensures consistent physical units matching state $[u]$. The scalar relaxation factor $\alpha^* \in [0.0, 0.20]$ is selected on $M = 5$ disjoint calibration trajectories from a deliberately small, bounded grid; the range limits a local correction so it cannot silently replace the frozen predictor:

$$\alpha^* = \arg \min_{\alpha \in \{0.0, 0.01, 0.02, 0.05, 0.10, 0.20\}} \mathcal{E}_{\text{cal}}(\alpha). \tag{18}$$

Once discrete $\hat{u}_{k+1}$ is obtained, querying $u_H(s)$ between $\hat{u}_k$ and $\hat{u}_{k+1}$ supplies a continuous endpoint bridge without temporal PDE integration. The grid includes $\alpha = 0$, which recovers projection-only rollout when relaxation is not selected. In the primary fluid rows with requested $\Delta t \geq 0.10$ seconds, 25 of 32 rows select $\alpha = 0$; this is an intended adaptive outcome when the FM is already reliable at its native cadence. The off-native 0.05-second regime selects positive relaxation in all eight primary fluid rows. Selection is empirical and is not implied by a generic CFL argument.

5 Experimental Methodology

5.1 Benchmark Specifications and Evaluation Protocols

The benchmark evaluates frozen interfaces from five PDE model families (Poseidon, CNO, Polymathic, DPOT, and MORPH) across fluid, reaction–diffusion, and wave systems, requested cadences, and two deployment tasks (Herde et al., 2024; Raonic et al., 2023; McCabe et al., 2023; Hao et al., 2024; Rautela et al., 2025). KMF operates outside each interface and never updates foundation-model parameters. The primary Task 1 analysis focuses on models for which the endpoint-conditioned refinement contract is meaningful, namely Poseidon and CNO. DPOT, MORPH, and Polymathic transfer rows remain in the appendix as contract-specific diagnostics rather than defining the primary temporal-refinement claim.

The available data systems are NS-Gauss, FNS-KF, ACE, and Wave-Gauss on 128×128 grids; these are drawn from the released PDEGym/Poseidon dataset collection (Herde et al., 2024). Requested spans range from 0.05 to 0.40 seconds. Because the assembled trajectories contain a finite number of raw frames, the longest-span runs use the largest valid rollout horizon for each dataset; unavailable cells are reported as unavailable. The matrix distinguishes the requested cadence from the endpoint span used by the two-sided refinement calculation. Calibration and test trajectories remain disjoint, and all reported fusion or correction strengths are selected without using held-out test errors. The operator-mismatch audit in Appendix Table 10 is deliberately recalibrated sensitivity analysis, not a claim of zero-shot robustness to arbitrary misspecification.

For Task 2 we report two operational regimes. The off-native regime uses $\Delta t = 0.05$ seconds and tests active physical-defect relaxation during causal rollout. The native-cadence regime uses $\Delta t = 0.10$ seconds and tests projection and defect correction at the model's standard endpoint spacing. In both regimes, the corrected endpoint is reinjected before the next foundation-model call, and the dense bridge is evaluated between successive corrected endpoints whenever the dataset provides a scoreable intermediate reference. The appendix contains every evaluated configuration, including rows that are not part of the primary claims.

Table 1: Task 1 temporal refinement. Values are midpoint RMSE; percentages give KMF gain versus raw FM and linear interpolation. CNO appears for NS-Gauss and FNS-KF, systems covered by its validated velocity-only contract. Other pairs appear in the appendix.

PDE	FM	Δt	Lin.	Raw	KMF
NS-Gauss	B	.10	.06165	.03004	**.02509 (+16.5/+59.3)**
	T	.10	.06139	.05336	**.02873 (+46.2/+53.2)**
	CNO	.10	.06106	.03681	**.02684 (+27.1/+56.1)**
	B	.40	.22715	.00425	**.00425 (0.0/+98.1)**
	T	.40	.22699	.00858	**.00858 (0.0/+96.2)**
FNS-KF	B	.10	.01018	.01248	**.00873 (+30.0/+14.3)**
	T	.10	.01199	.07945	**.01103 (+86.1/+8.0)**
	CNO	.10	.03110	.09023	**.03059 (+66.1/+1.7)**
	B	.40	.05972	.01076	**.01076 (0.0/+82.0)**
	T	.40	.06035	.01488	**.01487 (+0.1/+75.4)**
ACE	B	.10	.12361	.16839	**.12230 (+27.4/+1.1)**
	T	.10	.11793	.19875	**.11664 (+41.3/+1.1)**
	B	.40	.23925	.38204	**.23304 (+39.0/+2.6)**
	T	.40	.23845	.37162	**.23187 (+37.6/+2.8)**
Wave-Gauss	B	.10	.08150	.09202	**.08537 (+7.2/-4.8)**
	T	.10	.07305	.08673	**.08218 (+5.2/-12.5)**
	B	.40	.17246	.18827	**.16215 (+13.9/+6.0)**
	T	.40	.17246	.16643	**.14196 (+14.7/+17.7)**

Table 2: Task 2 Poseidon rollout from the master matrix. Values are four-step window RMSE; percentages give KMF gain versus raw rollout. "Off" and "Native" denote $\Delta t = 0.05$ and 0.10 seconds. The full matrix is in Appendix Table 17.

Reg.	PDE	FM	Δt	Raw	Proj.	KMF
Off	NS-G	B	.05	.05919	.05890	**.04941 (+16.5)**
		T	.05	.14967	.10534	**.08853 (+40.9)**
	FNS-KF	B	.05	.02675	.02627	**.02488 (+7.0)**
		T	.05	.15862	.14629	**.13089 (+17.5)**
Native	NS-G	B	.10	.00420	.00394	**.00394 (+6.2)**
		T	.10	.00888	.00849	**.00849 (+4.5)**
	FNS-KF	B	.10	.01434	.01431	**.01431 (+0.2)**
		T	.10	.04201	.04031	**.04031 (+4.0)**
Off	ACE	B	.05	.28498	.28498	**.25460 (+10.7)**
		T	.05	.31724	.31724	**.28610 (+9.8)**
	Wave-G	B	.05	.11130	.10958	**.10958 (+1.6)**
		T	.05	.10471	.10649	.10649 (-1.7)
Native	ACE	B	.10	.39916	.39916	**.36609 (+8.3)**
		T	.10	.35590	.35590	**.32842 (+7.7)**
	Wave-G	B	.10	.13151	.11462	**.11462 (+12.8)**
		T	.10	.13948	.11857	**.11857 (+15.0)**

6 Results and Architectural Analysis

6.1 Sub-Cadence Refinement: KMF Physical Candidate and Fusion

Task 1 evaluates the state at an intermediate time between two foundation-model endpoints. For readability, the main table reports representative Poseidon configurations across four PDE systems and two spans. The complete model-family and cadence matrix, including CNO and transfer-model diagnostics, is provided in Appendix Table 16. All candidates receive the same endpoint information. Linear interpolation is the matched endpoint baseline, while the raw FM estimate is included as a practical reference. The component-separated, projection-matched audit in Appendix Table 4 additionally compares linear, raw FM, limited physical, kinematic physical, and calibrated fusion candidates under a single protocol.

KMF is competitive with or better than the matched linear bridge. At some long spans raw FM is already best, and calibration preserves it rather than forcing a physical candidate. In the projection-matched audit, fusion improves on both raw FM and the physical candidate, from 0.031174 to 0.026722 on NS-Gauss and from 0.012536 to 0.008391 on FNS-KF; with five calibration trajectories it degrades in $0/50$ resamples on either fluid system (Appendix Tables 4 and 9).

Scope of Task 1. The primary Task 1 claim requires informative frozen endpoints and a cheap approximate PDE operator. Transfer rows document the boundary: phase dispersion or severe model/operator mismatch can make an endpoint-polynomial candidate poor, and calibration can protect against but not universally repair this disagreement.

6.2 Defect-Corrected Autoregressive Rollout Across Horizons

Task 2 has two distinct regimes, and they should not be pooled. In the off-native $\Delta t = 0.05$ regime, physical defect relaxation is active and the corrected endpoint is reinjected at every step. In the native $\Delta t = 0.10$ regime, the foundation model is already operating at its standard endpoint spacing, so calibration can select projection-only correction; these rows measure the value of constraint enforcement rather than implying that defect relaxation is always selected. The compact main table reports representative Poseidon configurations across the four PDE systems and both operational cadences; the full model and cadence matrix is in Appendix Table 17.

The off-native regime gives the clearest evidence for active physical correction. Raw $\rightarrow$ Proj. measures invariant enforcement; Proj. $\rightarrow$ KMF measures the active defect term; reinjection propagates

Table 3: Off-native Task 2 correction comparison ($\Delta t = 0.05$ s; Poseidon-B; four steps). Five calibration and 20 held-out trajectories; latency excludes the common FM pass. PC-style is a reduced-state cached-Jacobian adapter, not an official PhysicsCorrect reproduction.

System	Method	Step	Window RMSE	Gain vs. Proj.	Online ms	Cache s
NS-Gauss	Raw FM	0	0.063619	−0.48%	0.04	N/A
	Projection	0	0.063318	0.00%	0.47	N/A
	KMF	0.20	**0.053091**	**+16.15%**	20.03	N/A
	Residual gradient	0.20	0.062066	+1.98%	10.68	N/A
	Live GN	0.20	0.051993	+17.89%	1108.09	N/A
	PC-style cached Jacobian	0.20	0.053575	+15.39%	5.09	0.489
	Fixed-budget flow	0.10	0.057544	+9.12%	56.44	N/A
FNS-KF	Raw FM	0	0.026748	−1.98%	0.03	N/A
	Projection	0	0.026228	0.00%	0.46	N/A
	KMF	0.10	**0.024724**	**+5.73%**	4.46	N/A
	Residual gradient	0.20	0.025796	+1.65%	4.23	N/A
	Live GN	0.10	0.024877	+5.15%	1095.15	N/A
	PC-style cached Jacobian	0.10	0.024937	+4.92%	3.39	0.092

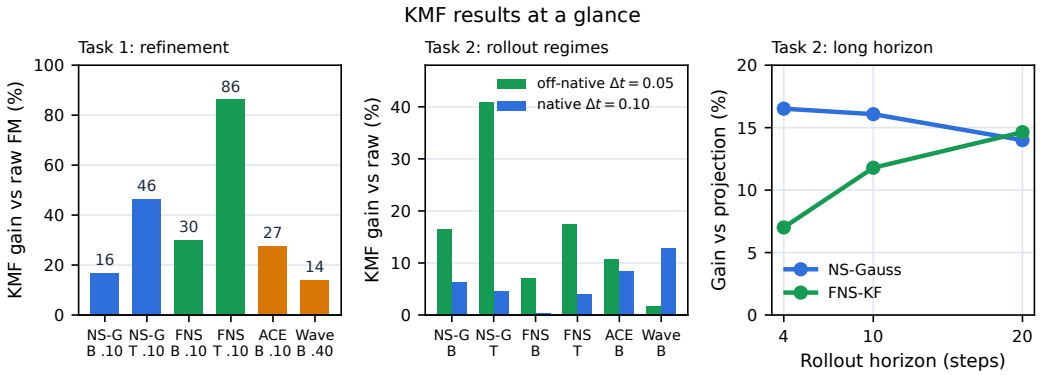

Figure 2: KMF results summary: Task 1 refinement, the two Task 2 regimes, and longer-horizon active correction. Detailed audits are in Appendix Tables 15 and 14.

the endpoint correction to later calls. Native $\alpha = 0$ rows are KMF-selected projection, not active defect correction. Dense evaluation is scoreable at $\Delta t \geq 0.10$; at 0.05 seconds the causal endpoint rollout remains scored but no additional raw intermediate reference exists.

Across 20 calibration resamples, positive relaxation improves projection-only rollout at $H = 10/H = 20$ by $16.08/14.00\%$ on NS-Gauss and $11.79/14.64\%$ on FNS-KF. In native NS-Gauss, raw error grows from 0.00407 at step 1 to 0.01424 by step 10, while KMF ends at 0.01398 (Appendix Table 6). The PhysicsCorrect-style row is a contract-matched adapter, not an official reproduction; its 0.053091 value uses the separate five-calibration, 20-held-out protocol, whereas 0.04941 is from the 20-calibration, 50-held-out master matrix.

The rollout-dynamics diagnostic and dense-bridge visualization are provided in Appendix Figure 4.

6.3 Inference Cost and Solver Context

Table 3 reports online correction cost excluding the common FM pass. The $1.03\times$ value in Appendix Table 14 is the endpoint-conditioned two-RHS pipeline, not active rollout. Active Task 2 costs 20.03 ms online; with the 96.33 ms Poseidon-B pass, this is 116.36 ms end-to-end ($1.21\times$). Live Gauss–Newton costs 1108.09 ms per step versus KMF's 20.03 ms. RK controls use only the initial state while KMF receives an FM endpoint, so Appendix Table 15 compares complementary information regimes.

Conclusion. KMF turns frozen discrete-cadence PDE models into physics-aware continuous-time predictors and causally corrected rollouts using local RHS evaluations and low-dimensional calibration, without retraining or temporal PDE integration.

Reproducibility Statement. An anonymous implementation of KMF and the experiment runners is provided at `https://github.com/kmfcorrection/KMF`. The main text specifies the information contracts, estimators, calibration protocol, and evaluation tasks. The appendix provides the complete algorithm, assumptions and proofs for the theoretical results, dataset-processing and benchmark specifications, full model–PDE–cadence tables, solver and latency audits, and supporting ablations. The released repository contains the corresponding source code and scripts; datasets and pretrained checkpoints are obtained from their cited public sources and are not redistributed.

REFERENCES

Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2nd edition, 2017.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Zituo Chen and Sili Deng. Flow marching for a generative pde foundation model. *arXiv preprint arXiv:2509.18611*, 2025.

Alexandre Joel Chorin. Numerical solution of the navier-stokes equations. *Mathematics of computation*, 22(104):745–762, 1968.

Richard Courant, Kurt Friedrichs, and Hans Lewy. Über die partiellen differenzengleichungen der mathematischen physik. *Mathematische annalen*, 100(1):32–74, 1928.

Waleed Diab and Mohammed Al Kobaisi. Temporal neural operator for modeling time-dependent physical phenomena. *Scientific Reports*, 15(1):32791, 2025.

Ernst Hairer, Gerhard Wanner, and Syvert P Nørsett. *Solving ordinary differential equations I: Nonstiff problems*. Springer, 1993.

Jack K Hale. *Ordinary differential equations*. Courier Corporation, 2009.

Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, Anima Anandkumar, Jian Song, and Jun Zhu. DPOT: Auto-regressive denoising operator transformer for large-scale PDE pre-training. In *Forty-first International Conference on Machine Learning*, 2024. URL `https://openreview.net/forum?id=X7UnDevHOM`.

Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel De Bezenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *Advances in Neural Information Processing Systems*, 37:72525–72624, 2024. URL `https://arxiv.org/abs/2405.19101`.

Xinquan Huang and Paris Perdikaris. Physicscorrect: A training-free approach for stable neural pde simulations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 22057–22065, 2026.

Matthew Kelly. An introduction to trajectory optimization: How to do your own direct collocation. *SIAM review*, 59(4):849–904, 2017.

Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL `https://openreview.net/forum?id=c8P9NQVtmnO`.

Phillip Lippe, Bastiaan Veeling, Paris Perdikaris, Richard E. Turner, and Johannes Brandstetter. PDE-refiner: Achieving accurate long rollouts with neural PDE solvers. In *Advances in Neural Information Processing Systems*, 2023. URL `https://proceedings.neurips.cc/paper_files/paper/2023/file/d529b943af3dba734f8a7d49efcb6d09-Paper-Conference.pdf`.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.

Michael McCabe, Bruno Régaldo-Saint Blancard, Liam Parker, Ruben Ohana, Miles Cranmer, Alberto Bietti, Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, Mariel Pettee, Tiberiu Tesileanu, Kyunghyun Cho, and Shirley Ho. Multiple physics pretraining for physical surrogate models. In *NeurIPS 2023 AI for Science Workshop*, 2023. URL https://openreview.net/forum?id=M12lmQKuxa.

Ruben Ohana, Michael McCabe, Lucas Meyer, Rudy Morel, Fruzsina J. Agocs, Miguel Beneitez, Marsha Berger, Blakesley Burkhart, Stuart B. Dalziel, Drummond B. Fielding, Daniel Fortunato, Jared A. Goldberg, Keiya Hirashima, Yan-Fei Jiang, Rich R. Kerswell, Suryanarayana Maddu, Jonah Miller, Payel Mukhopadhyay, Stefan S. Nixon, Jeff Shen, Romain Watteaux, Bruno Régaldo-Saint Blancard, François Rozet, Liam Parker, Miles Cranmer, and Shirley Ho. The well: a large-scale collection of diverse physics simulations for machine learning. In *Advances in Neural Information Processing Systems*, volume 37, pp. 44989–45037. Curran Associates, Inc., 2024. doi: 10.52202/079017-1430. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/4f9a5acd91ac76569f2fe291b1f4772b-Paper-Datasets_and_Benchmarks_Track.pdf.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

C Radhakrishna Rao. *Linear statistical inference and its applications*. John Wiley & Sons, 2009.

Bogdan Raonic, Roberto Molinaro, Tim De Ryck, Tobias Rohner, Francesca Bartolucci, Rima Alaifari, Siddhartha Mishra, and Emmanuel De Bézenac. Convolutional neural operators for robust and accurate learning of pdes. *Advances in Neural Information Processing Systems*, 36:77187–77200, 2023.

Mahindra Singh Rautela, Alexander Most, Siddharth Mansingh, Bradley C Love, Alexander Scheinker, Diane Oyen, Nathan Debardeleben, Earl Lawrence, and Ayan Biswas. Morph: Pde foundation models with arbitrary data modality. *arXiv preprint arXiv:2509.21670*, 2025.

L Richard, Faires J Douglas, M Annette, et al. *Numerical analysis*. sengage, 2016.

Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. *Advances in neural information processing systems*, 35:1596–1611, 2022.

Lloyd N Trefethen. *Spectral methods in MATLAB*. SIAM, 2000.

Rui Wang, Robin Walters, and Rose Yu. Incorporating symmetry into deep dynamics models for improved generalization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=wta_8Hx2KD.

Hao Wu, Yuan Gao, Fan Xu, Fan Zhang, Qingsong Wen, Xiaomeng Huang, and Xian Wu. Differential-integral neural operator for long-term turbulence forecasting. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 5431–5441, 2026.

Jiaquan Zhang, Shuxu Chen, Haifan Meng, Yi Lu, Zhihan Lyu, Fan Mo, Wei Dong, Yang Yang, and Chaoning Zhang. Hero: History-enriched rollout training for long-horizon autoregressive neural operators. *arXiv preprint arXiv:2607.29135*, 2026.

A LLM USAGE DISCLOSURE

In this work, generative AI tools were used only as writing assistants for textual refinement, organization, and readability improvement, and for generating and revising figure and diagram drafts. They were later used to identify possible errors or inconsistencies in the manuscript that could have been missed during manual review. All suggested improvements were evaluated by the authors and implemented manually. Generative AI tools were not used to generate benchmark data, produce

Table 4: Independent Poseidon-B interpolation audit with projection-matched controls. Full protocol and uncertainty details are reported in the accompanying text.

Dataset	Span	Linear+P	Raw FM+P	Limited Hermite+P	Kinematic Hermite+P	Calibrated Fusion	Gain vs. FM+P
NS-Gauss	0.10 s	0.063510	0.031174	0.044356	0.035731	**0.026722**	**+14.28**% [13.12, 18.14]
FNS-KF	0.10 s	0.010176	0.012536	0.009537	0.008928	**0.008391**	**+33.07**% [29.98, 36.56]

Table 5: Independent component and native-rollout audit. The native setting uses Poseidon's trained cadence; the detailed relaxation interpretation is given in the main text.

Dataset	4-step Raw	4-step Proj.	Div.: Raw $\to$ Proj.	α_{cal}	$H = 10$ gain
NS-Gauss	0.004882	0.004564	$3.45\times10^{-1} \to 5.02\times10^{-14}$	0	+1.83%
FNS-KF	0.014340	0.014305	$2.32\times10^{-2} \to 3.07\times10^{-15}$	0	+0.04%

scientific results, or make final scientific decisions. No language model is an author. The authors take complete responsibility for the final contents of this paper, including all text, claims, figures, and other artifacts produced with AI assistance.

B Additional Figures

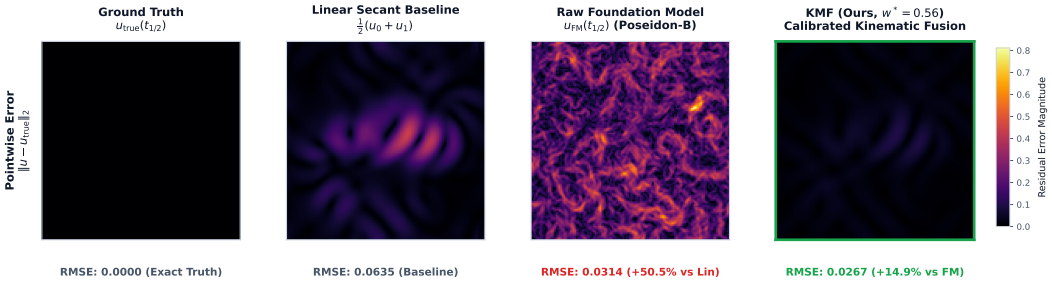

Figure 3: Illustrative endpoint-conditioned refinement. Pointwise fields compare the reference, raw FM estimate, and KMF.

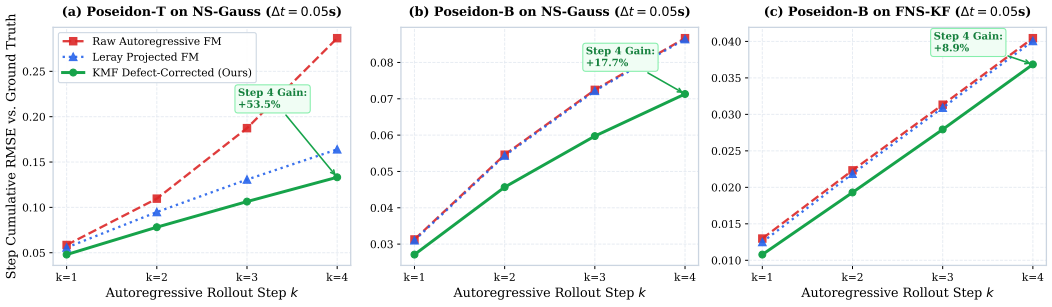

Figure 4: Causal rollout and dense-bridge diagnostics for the two Task 2 regimes.

C Independent Real-Pipeline Audit Details

Longer-horizon continuation. The active off-native evaluation extends the same calibration-only selection protocol to ten- and twenty-step autoregressive rollouts at $\Delta t = 0.05\,\text{s}$. Across 20 calibration resamples and one fixed 50-trajectory test set, positive relaxation is selected in every split. The held-out gain over projection-only rollout is 16.08% and 14.00% on NS-Gauss at $H = 10$ and $H = 20$, respectively, and 11.79% and 14.64% on FNS-KF. These results complement the native-cadence audit, where projection is often the selected correction, and show that the active defect term remains useful when recursively accumulated error is larger.

Table 6: Long-horizon active correction at off-native $\Delta t = 0.05\,\mathrm{s}$. Gains are relative to projection-only rollout.

System	H	α median	$\alpha > 0$	Gain vs. Proj.	Cos.
NS-Gauss	10	0.20	20/20	+16.08%	+0.851
NS-Gauss	20	0.20	20/20	+14.00%	+0.802
FNS-KF	10	0.20	20/20	+11.79%	−0.052
FNS-KF	20	0.20	20/20	+14.64%	+0.056

Table 7: Calibration-resampling stability at active off-native cadence. Selection uses calibration trajectories only; gains are evaluated on a fixed held-out test set.

System	Splits	Cal. / pool	Selected α	Positive α	Gain vs. Proj.	Cos.
NS-Gauss	20	5/20	0.20	20/20	+16.42%	+0.789
FNS-KF	20	5/20	0.10	20/20	+5.43%	+0.188

D Reviewer-Closure Audit Details

All tables in this section use Poseidon-B at 128×128, 20 calibration trajectories, and a fixed disjoint 50-trajectory test set. Calibration resampling draws only from a separate 70-trajectory candidate pool. These tests supplement, rather than replace, the main audit in Table 4.

Transfer-contract audit. The complete matrix audit records unavailable raw-FM cells separately from numerical failures: 101 of 231 inspected rows did not expose a raw FM endpoint under the fixed-step interface, and are not used for raw-FM comparisons. Among rows with a valid raw endpoint, negative bridge gains are concentrated in dispersive or transfer configurations, where phase error and operator mismatch can dominate polynomial refinement. These rows are retained in the artifact and interpreted as scope diagnostics rather than hidden by aggregate averages.

D.1 Background Motivation Details

Physics-informed neural networks and related residual-optimization methods enforce a PDE by minimizing an objective of the form (Raissi et al., 2019)

$$\mathcal{J}_{\mathrm{phys}}(v) = \frac{1}{2}\|v - u_{\mathrm{FM}}\|^2 + \frac{\lambda}{2}\|\mathcal{R}(v)\|^2, \tag{19}$$

where $\mathcal{R}(v)$ is the discretized PDE residual. This formulation is conceptually related to KMF, but its deployment cost depends on differentiating through the residual and repeatedly optimizing a high-dimensional field. KMF instead evaluates the spatial operator at available states and performs a closed-form bridge, calibrated fusion, or scalar defect relaxation. The comparison is therefore about deployment contract and computation, not a claim that residual optimization is generally ineffective.

Forward numerical integration is a complementary alternative. Explicit Runge–Kutta methods use multiple right-hand-side evaluations per temporal substep and can be preferable when a stable discretization and sufficient budget are available (Hairer et al., 1993). The main text does not treat a single coarse RK4 step as a universal solver baseline. Appendix E.6 gives the representative modal calculation used only to illustrate one coarse-step failure mode, while the fair solver-control audit in Table 15 reports stable and unstable regimes empirically.

E Supporting Mathematical Derivations

This appendix records derivations supporting the main interpolation, projection, and rollout equations. The main-text claims rely only on the stated theorem, proposition, and explicitly scoped empirical results.

Table 8: Causal correction suite at off-native $\Delta t = 0.05\,\mathrm{s}$. PC-style denotes the reduced-state cached-Jacobian adapter; it is not an official PhysicsCorrect reproduction.

System	Method	Step	Window RMSE	Gain vs. Proj.	Online ms	Cache s
NS-Gauss	Raw FM	0	0.063619	−0.48%	0.04	N/A
	Projection	0	0.063318	0.00%	0.47	N/A
	KMF	0.20	**0.053091**	**+16.15%**	20.03	N/A
	Residual gradient	0.20	0.062066	+1.98%	10.68	N/A
	Live GN	0.20	0.051993	+17.89%	1108.09	N/A
	PC-style cached Jacobian	0.20	0.053575	+15.39%	5.09	0.489
	Fixed-budget flow	0.10	0.057544	+9.12%	56.44	N/A
FNS-KF	Raw FM	0	0.026748	−1.98%	0.03	N/A
	Projection	0	0.026228	0.00%	0.46	N/A
	KMF	0.10	**0.024724**	**+5.73%**	4.46	N/A
	Residual gradient	0.20	0.025796	+1.65%	4.23	N/A
	Live GN	0.10	0.024877	+5.15%	1095.15	N/A
	PC-style cached Jacobian	0.10	0.024937	+4.92%	3.39	0.092

Table 9: **Calibration-budget closure audit.** Mean ± standard deviation of held-out RMSE over 50 independently sampled calibration subsets; the test set is never resampled or used for selection. "Deg." is the fraction of subsets for which fusion exceeds the better projected candidate on that fixed test set. The table reports held-out performance and degradation directly; calibration weights are omitted for compactness.

M	NS-Gauss RMSE	NS-Gauss Deg.	FNS-KF RMSE	FNS-KF Deg.
1	0.027576 ± 0.001074	$0/50$	0.008491 ± 0.000111	$0/50$
2	0.027078 ± 0.000494	$0/50$	0.008456 ± 0.000101	$1/50$
5	0.026820 ± 0.000154	$0/50$	0.008415 ± 0.000034	$0/50$
10	0.026739 ± 0.000041	$0/50$	0.008403 ± 0.000019	$0/50$
20	0.026722 ± 0.000030	$0/50$	0.008396 ± 0.000008	$0/50$

E.1 Derivation of Cubic Hermite Basis Polynomials

The four displayed basis functions in Section 4.1 are the unique cubics matching value and first derivative at both normalized endpoints. Solving the resulting 4×4 linear systems gives $h_{00} = 1 - 3s^2 + 2s^3$, $h_{10} = 3s^2 - 2s^3$, $h_{01} = \Delta t(s - 2s^2 + s^3)$, and $h_{11} = \Delta t(-s^2 + s^3)$. The factors of Δt in the derivative basis functions are essential: they ensure that a state-rate, rather than a normalized derivative, is interpolated.

Proof sketch for Theorem 1. For each spatial coordinate, the standard Hermite remainder is $u(t) - u_H(t) = u^{(4)}(\xi)(t - t_0)^2(t - t_1)^2/4!$. Its maximum normalized nodal factor is $1/16$ at $s = 1/2$, giving the $\Delta t^4/384$ bound after taking the stated L^p norm. With an endpoint operator error bounded by ϵ_{phys}, the two derivative basis functions contribute at most $\Delta t \epsilon_{\mathrm{phys}}/4$ at the midpoint. The deployed endpoint perturbation result in the main text additionally accounts for forecast-endpoint error.

E.2 A Mathematical Primer on Physical Invariant Manifolds and Variational Projection Theory

We provide a rigorous mathematical and pedagogical foundation for physical invariant manifolds, the geometry of neural rollout drift, and the non-expansive variational projection framework utilized throughout Section 4.2.

The Ambient State Space $\mathcal{X}$: In continuum physical systems, instantaneous states are spatial fields defined over a physical domain $\Omega \subset \mathbb{R}^d$ (with spatial dimension $d \in \{1, 2, 3\}$). The state space $\mathcal{X}$ is a Hilbert space equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$ and induced norm $\|u\|_{\mathcal{X}} = \sqrt{\langle u, u \rangle_{\mathcal{X}}}$. For vector-valued velocity fields $u(x) = (u_1(x), \ldots, u_d(x))^T$ in fluid mechanics, $\mathcal{X} = L^2(\Omega)^d$

Table 10: **Calibrated operator-mismatch closure audit.** A $0.20\,\mathrm{s}$ endpoint span is queried off-center; KMF is recalibrated for each assumed operator. NS-Gauss perturbs spectral-viscosity multiplier over five values in $[0.5, 1.5]$; FNS-KF perturbs the 5×5 viscosity/forcing grid. Ranges summarize every perturbation. This is calibrated sensitivity analysis, not zero-shot mismatch adaptation. All values are RMSE.

System	s	# variants	Linear+P	Raw FM+P	Nominal bridge	Nominal KMF	Recalibrated KMF range
NS-Gauss	0.25	5	0.105647	0.031173	0.067064	0.024911	$[0.024891, 0.024932]$
NS-Gauss	0.75	5	0.099144	0.004053	0.056376	0.004053	$[0.004053, 0.004053]$
FNS-KF	0.25	25	0.020831	0.012534	0.021645	0.011079	$[0.010888, 0.011314]$
FNS-KF	0.75	25	0.019298	0.008148	0.019864	0.007382	$[0.007253, 0.007528]$

Table 11: **Empirical turbulent cadence scaling.** Midpoint RMSE of the two-endpoint projected bridge and matched linear interpolation on the fixed fluid test sets. The fitted slopes are descriptive finite-cadence values, not estimates of the theorem's asymptotic rate.

System	Span (s)	Linear+P	Bridge+P	Bridge slope
NS-Gauss	0.10	0.063510	0.035731	1.338
	0.20	0.135175	0.098225	
	0.30	0.192271	0.164455	
	0.40	0.236330	0.225146	
FNS-KF	0.10	0.010176	0.008928	1.342
	0.20	0.026847	0.023855	
	0.30	0.045001	0.040623	
	0.40	0.061928	0.056665	

denotes the space of square-integrable vector fields endowed with the standard L^2 inner product:

$$\langle u, v \rangle_{L^2(\Omega)^d} = \int_\Omega u(x) \cdot v(x)\, dx = \sum_{j=1}^{d} \int_\Omega u_j(x) v_j(x)\, dx, \quad \|u\|^2_{L^2(\Omega)^d} = \int_\Omega \|u(x)\|_2^2\, dx. \tag{20}$$

For scalar physical fields (such as temperature, order parameters in phase-field kinetics, or chemical concentrations), $\mathcal{X} = L^2(\Omega)$ represents the space of square-integrable scalar functions.

Physical Invariant Manifolds $\mathcal{M}$: An autonomous partial differential equation $\partial_t u = F(u)$ generates a continuous solution semi-flow $\Phi_t : \mathcal{X} \to \mathcal{X}$ mapping any initial state $u(0) \in \mathcal{X}$ to its solution at time t: $u(t) = \Phi_t(u(0))$. A subset $\mathcal{M} \subset \mathcal{X}$ is defined as a *dynamically invariant manifold* if it is forward-invariant under the PDE flow (Hale, 2009):

$$\Phi_t(\mathcal{M}) \subseteq \mathcal{M}, \quad \forall t \geq 0. \tag{21}$$

In mathematical physics, invariant manifolds characterize exact conservation laws or kinematic constraints that all physical trajectories must satisfy. Formally, let $\mathcal{C} : \mathcal{X} \to \mathcal{Y}$ be a continuous constraint operator mapping state space $\mathcal{X}$ into constraint space $\mathcal{Y}$. The physical invariant manifold is the zero-level set (or kernel) of $\mathcal{C}$:

$$\mathcal{M} = \{v \in \mathcal{X} : \mathcal{C}(v) = 0\}. \tag{22}$$

For example, in incompressible fluid dynamics, mass conservation requires the velocity field to be solenoidal: $\mathcal{C}(u) = \nabla \cdot u = 0$. The solenoidal manifold $\mathcal{M}_{\text{div-free}} = \{u \in L^2(\Omega)^d : \nabla \cdot u = 0\}$ is strictly forward-invariant: if $\nabla \cdot u(0) = 0$, then the Navier-Stokes governing equations ensure $\nabla \cdot u(t) = 0$ for all $t > 0$.

The Geometry of Off-Manifold Neural Drift: Pretrained neural operators $f_\theta : \mathcal{X} \to \mathcal{X}$ are parameterized by unconstrained neural architectures (e.g., vision transformers, Fourier neural operators, convolutional backbones). Consequently, predicted states $v = f_\theta(u_k)$ are not mathematically guaranteed to satisfy the physical constraint $\mathcal{C}(v) = 0$. In a Hilbert space, any candidate state

Table 12: **Operator-cost audit.** Synchronized CUDA measurements for one spatial RHS evaluation, one admissibility projection, and their two-RHS KMF combination at 128×128. These costs are separate from the FM forward pass.

System	RHS ms	Projection ms	2 RHS + projection ms
NS-Gauss	2.53	2.39	7.44
FNS-KF	2.57	2.38	7.51
ACE	2.33	1.69	6.35
Wave-Gauss	2.10	2.36	6.57

$v \in \mathcal{X}$ admits a unique orthogonal decomposition into an on-manifold component $v_\parallel \in \mathcal{M}$ and an off-manifold error $v_\perp \in \mathcal{M}^\perp$:

$$v = v_\parallel + v_\perp, \quad \text{where } v_\parallel = \mathcal{P}_\mathcal{M}(v) \in \mathcal{M}, \quad v_\perp = (\mathbf{I} - \mathcal{P}_\mathcal{M})(v) \in \mathcal{M}^\perp, \quad \langle v_\parallel, v_\perp \rangle_\mathcal{X} = 0. \tag{23}$$

The component $v_\perp$ represents non-physical violations (e.g., artificial mass creation, negative densities, or unphysical high-frequency acoustic waves). In recursive autoregressive rollouts $\hat{u}_{k+1} = f_\theta(\hat{u}_k)$, such violations can act as out-of-distribution perturbations and contribute to compounding rollout error. This is a mechanism for drift, not a claim of a universal exponential-divergence law.

Variational Projection and the Green's Operator: To eliminate unphysical components without corrupting valid physical dynamics, KMF computes the unique closest state on $\mathcal{M}$ via variational projection:

$$\mathcal{P}_\mathcal{M}(v) = \arg\min_{u \in \mathcal{X}} \frac{1}{2}\|u - v\|_\mathcal{X}^2 \quad \text{subject to} \quad \mathcal{C}(u) = 0. \tag{24}$$

For linear balance laws $\mathcal{C}(u) = \mathcal{A}u - b = 0$, we solve this constrained minimization problem via the method of Lagrange multipliers in Hilbert space. Define the Lagrangian $\mathcal{L} : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$:

$$\mathcal{L}(u, \lambda) = \frac{1}{2}\|u - v\|_\mathcal{X}^2 + \langle \lambda, \mathcal{A}u - b \rangle_\mathcal{Y}. \tag{25}$$

Taking the Frechet derivative with respect to u and setting it to zero yields the first-order optimality condition:

$$\nabla_u \mathcal{L}(u, \lambda) = u - v + \mathcal{A}^*\lambda = 0 \implies u = v - \mathcal{A}^*\lambda, \tag{26}$$

where $\mathcal{A}^* : \mathcal{Y} \to \mathcal{X}$ is the adjoint operator of $\mathcal{A}$. Applying the linear operator $\mathcal{A}$ to both sides:

$$\mathcal{A}u = \mathcal{A}v - \mathcal{A}\mathcal{A}^*\lambda. \tag{27}$$

Enforcing the constraint $\mathcal{A}u = b$ yields:

$$b = \mathcal{A}v - \mathcal{A}\mathcal{A}^*\lambda \implies \mathcal{A}\mathcal{A}^*\lambda = \mathcal{A}v - b \implies \lambda = (\mathcal{A}\mathcal{A}^*)^{-1}(\mathcal{A}v - b). \tag{28}$$

Substituting optimal Lagrange multiplier λ back into the expression for u yields the exact orthogonal Green's projector:

$$\mathcal{P}_\mathcal{M}(v) = v - \mathcal{A}^*(\mathcal{A}\mathcal{A}^*)^{-1}(\mathcal{A}v - b). \tag{29}$$

Step-by-Step Proof of Proposition 2 (Non-Expansion of Manifold Projection): We now prove Proposition 2 from geometric first principles in Hilbert space.

Proof of Proposition 2. Let $\mathcal{H}$ be a real Hilbert space equipped with inner product $\langle \cdot, \cdot \rangle$ and induced norm $\|v\| = \sqrt{\langle v, v \rangle}$. Let $\mathcal{M} \subset \mathcal{H}$ be a closed linear subspace. By the Projection Theorem in Hilbert spaces, $\mathcal{H}$ decomposes into the direct sum of $\mathcal{M}$ and its orthogonal complement: $\mathcal{H} = \mathcal{M} \oplus \mathcal{M}^\perp$.

For any candidate state $v \in \mathcal{H}$, let $v_\parallel = \mathcal{P}_\mathcal{M}(v) \in \mathcal{M}$ and $v_\perp = (\mathbf{I} - \mathcal{P}_\mathcal{M})(v) \in \mathcal{M}^\perp$. Let $u^* \in \mathcal{M}$ be an arbitrary true state lying on the physical manifold. Consider the candidate prediction error $v - u^*$. We decompose this error into its on-manifold and off-manifold parts:

$$v - u^* = (v_\parallel - u^*) + v_\perp. \tag{30}$$

Since $u^* \in \mathcal{M}$ and $v_\parallel \in \mathcal{M}$, their difference lies in $\mathcal{M}$: $v_\parallel - u^* \in \mathcal{M}$. Meanwhile, $v_\perp \in \mathcal{M}^\perp$. Taking the squared Hilbert norm:

$$\|v - u^*\|^2 = \langle (v_\parallel - u^*) + v_\perp,\ (v_\parallel - u^*) + v_\perp \rangle \tag{31}$$

$$= \|v_\parallel - u^*\|^2 + 2\langle v_\parallel - u^*, v_\perp \rangle + \|v_\perp\|^2. \tag{32}$$

Because $v_\parallel - u^* \in \mathcal{M}$ and $v_\perp \in \mathcal{M}^\perp$, their inner product vanishes identically: $\langle v_\parallel - u^*, v_\perp \rangle = 0$. Thus, the cross term drops out, leaving the exact Pythagorean identity in Hilbert space:

$$\|v - u^*\|^2 = \|\mathcal{P}_\mathcal{M}(v) - u^*\|^2 + \|v_\perp\|^2. \tag{33}$$

Because $\|v_\perp\|^2 = \|(\mathbf{I} - \mathcal{P}_\mathcal{M})v\|^2 \geq 0$, we have:

$$\|v - u^*\|^2 \geq \|\mathcal{P}_\mathcal{M}(v) - u^*\|^2 \implies \|\mathcal{P}_\mathcal{M}(v) - u^*\| \leq \|v - u^*\|. \tag{34}$$

Furthermore, equality holds if and only if $\|v_\perp\|^2 = 0$, which occurs if and only if $v \in \mathcal{M}$. Whenever $v \notin \mathcal{M}$, $\|v_\perp\| > 0$, and the inequality is strictly contractive: $\|\mathcal{P}_\mathcal{M}(v) - u^*\| < \|v - u^*\|$. □

Functional-Analytic Meaning of Non-Expansiveness: In operator theory and functional analysis, an operator $T : \mathcal{X} \to \mathcal{X}$ on a Hilbert space $\mathcal{X}$ is termed *non-expansive* (or 1-Lipschitz continuous) if it satisfies:

$$\|T(x) - T(y)\| \leq \|x - y\|, \quad \forall x, y \in \mathcal{X}. \tag{35}$$

In convex analysis and monotone operator theory, orthogonal projections onto closed convex sets satisfy an even stronger property known as *firm non-expansiveness*:

$$\|\mathcal{P}_\mathcal{M}(x) - \mathcal{P}_\mathcal{M}(y)\|^2 \leq \langle \mathcal{P}_\mathcal{M}(x) - \mathcal{P}_\mathcal{M}(y),\ x - y \rangle \leq \|x - y\|^2. \tag{36}$$

In the present setting, this property gives an exact non-expansion statement for the specified linear constraint and L^2 metric. Let $u^* \in \mathcal{M}$ denote the ground-truth physical state. Because u^* satisfies the physical invariant constraint $\mathcal{C}(u^*) = 0$, it is an invariant fixed point of the projection operator: $\mathcal{P}_\mathcal{M}(u^*) = u^*$. Substituting $y = u^*$ into the non-expansiveness relation immediately yields:

$$\|\mathcal{P}_\mathcal{M}(v) - u^*\| \leq \|v - u^*\|. \tag{37}$$

This inequality says that, for a target satisfying the same linear constraint, the specified orthogonal projection cannot increase L^2 error. Whenever the candidate contains a non-zero orthogonal violation, the Pythagorean identity in Equation 33 gives strict reduction in this metric. The result does not transfer automatically to nonlinear constraints, approximate projections, or metrics other than the stated Hilbert norm.

E.3 Algebraic Derivation of the Fourier Helmholtz-Leray Operator

We derive the Fourier-space representation of the Helmholtz-Leray projection operator $\mathcal{P}_{\text{div-free}}$ defined in Equation 10.

Let $u : \mathbb{T}^d \to \mathbb{R}^d$ be a square-integrable vector field on the d-dimensional torus with periodic boundary conditions. By the Helmholtz-Hodge decomposition, any vector field u decomposes uniquely into a divergence-free (solenoidal) field w and a curl-free (irrotational) gradient field $\nabla\phi$:

$$u = w + \nabla\phi, \quad \text{where } \nabla \cdot w = 0. \tag{38}$$

Taking the divergence of both sides:

$$\nabla \cdot u = \nabla \cdot w + \nabla \cdot (\nabla\phi) = 0 + \nabla^2\phi = \nabla^2\phi. \tag{39}$$

This is a Poisson equation for the scalar potential ϕ: $\nabla^2\phi = \nabla \cdot u$.

We solve this equation in Fourier space. Using the Fourier transform convention $\widehat{f}(\mathbf{k}) = \int_{\mathbb{T}^d} f(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} d\mathbf{x}$, the differential operators transform as $\widehat{\nabla\phi}(\mathbf{k}) = i\mathbf{k}\widehat{\phi}(\mathbf{k})$, $\widehat{\nabla \cdot u}(\mathbf{k}) = i\mathbf{k} \cdot \widehat{u}(\mathbf{k})$, and $\widehat{\nabla^2\phi}(\mathbf{k}) = -\|\mathbf{k}\|^2\widehat{\phi}(\mathbf{k})$. Substituting these into the Fourier-transformed Poisson equation:

$$-\|\mathbf{k}\|^2\widehat{\phi}(\mathbf{k}) = i\mathbf{k} \cdot \widehat{u}(\mathbf{k}) = i\mathbf{k}^T\widehat{u}(\mathbf{k}). \tag{40}$$

For any non-zero wavevector $\mathbf{k} \neq \mathbf{0}$, divide by $-\|\mathbf{k}\|^2$:

$$\widehat{\phi}(\mathbf{k}) = -\frac{i\mathbf{k}^T \widehat{u}(\mathbf{k})}{\|\mathbf{k}\|^2}. \tag{41}$$

Now compute the Fourier transform of the gradient component $\nabla\phi$:

$$\widehat{\nabla\phi}(\mathbf{k}) = i\mathbf{k}\widehat{\phi}(\mathbf{k}) = i\mathbf{k}\left(-\frac{i\mathbf{k}^T\widehat{u}(\mathbf{k})}{\|\mathbf{k}\|^2}\right) = \frac{(-i^2)\mathbf{k}\mathbf{k}^T\widehat{u}(\mathbf{k})}{\|\mathbf{k}\|^2} = \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\widehat{u}(\mathbf{k}). \tag{42}$$

Since $w = u - \nabla\phi$, the solenoidal component in Fourier space is:

$$\widehat{w}(\mathbf{k}) = \widehat{u}(\mathbf{k}) - \widehat{\nabla\phi}(\mathbf{k}) = \widehat{u}(\mathbf{k}) - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\widehat{u}(\mathbf{k}) = \left(\mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\right)\widehat{u}(\mathbf{k}). \tag{43}$$

This is precisely the matrix projection operator $P(\mathbf{k}) = \mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}$ defined in Equation 10.

Verification of Idempotence and Hermitian Symmetry: For any $\mathbf{k} \neq \mathbf{0}$, $P(\mathbf{k})$ satisfies:

$$P(\mathbf{k})^T = \left(\mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\right)^T = \mathbf{I} - \frac{(\mathbf{k}\mathbf{k}^T)^T}{\|\mathbf{k}\|^2} = \mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2} = P(\mathbf{k}), \tag{44}$$

$$P(\mathbf{k})^2 = \left(\mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\right)\left(\mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2}\right) = \mathbf{I} - 2\frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2} + \frac{\mathbf{k}(\mathbf{k}^T\mathbf{k})\mathbf{k}^T}{\|\mathbf{k}\|^4} = \mathbf{I} - \frac{\mathbf{k}\mathbf{k}^T}{\|\mathbf{k}\|^2} = P(\mathbf{k}). \tag{45}$$

Hence $P(\mathbf{k})$ is an exact orthogonal projector. By Parseval's theorem, the Fourier transform is a scaled isometry on L^2. Therefore, the Fourier-space projection operator is an exact orthogonal projection on $L^2(\mathbb{T}^d)^d$, strictly satisfying Proposition 2.

The Physical Mechanism Connecting Divergence to Autoregressive Drift: In incompressible continuum mechanics, the continuity equation reduces to the solenoidal condition $\nabla \cdot u = 0$, signifying that fluid parcel volume is strictly conserved and density remains uniform throughout space. Under the Helmholtz-Hodge decomposition, any candidate vector field v decomposes into a divergence-free component w and an irrotational gradient field $\nabla\phi$:

$$v = w + \nabla\phi, \quad \text{where } \nabla \cdot w = 0 \text{ and } \nabla^2\phi = \nabla \cdot v. \tag{46}$$

The gradient component $\nabla\phi$ is unphysical: regions where $\nabla \cdot v > 0$ represent fictitious mass creation (sources), whereas regions where $\nabla \cdot v < 0$ represent fictitious mass destruction (sinks). While classical computational fluid solvers enforce incompressibility at every numerical time-step through elliptic Poisson solves for pressure, standard neural foundation models (such as Fourier neural operators, vision transformers, or denoising diffusion models) treat velocity prediction as unconstrained regression. Consequently, even when trained exclusively on divergence-free training data, neural operators output predictions $\hat{u}$ with small non-zero residual divergence ($\|\nabla \cdot \hat{u}\| > 0$).

In single-step inference, these small divergence residuals can appear harmless. In recursive autoregressive forecasting ($\hat{u}_{k+1} = f_\theta(\hat{u}_k)$), feeding a state with spurious compressible modes back to a model trained mainly on solenoidal trajectories can contribute to compounding error. Applying the Fourier Helmholtz–Leray projection $\mathcal{P}_{\text{div-free}}$ removes that gradient component in $\mathcal{O}(N^d \log N)$ operations before the next recursive call. For the stated linear constraint, it provides the exact L^2 non-expansion property of Proposition 2; it does not by itself guarantee long-horizon stability of an arbitrary learned rollout.

E.4 Calculus Derivation of the Calibrated Consensus Weight w^*

We derive the closed-form ordinary least-squares formula for the calibrated consensus weight w^* defined in Equation 14.

Let $u_{H,\text{cal}}^{(m)}$ and $u_{\text{FM, cal}}^{(m)}$ denote the evaluated kinematic and neural predictions on M calibration trajectories ($m = 1, \ldots, M$), and let $u_{\text{true, cal}}^{(m)}$ denote the corresponding true physical fields. The convex combination is $u(w) = (1-w)u_{H,\text{cal}}^{(m)} + w u_{\text{FM, cal}}^{(m)} = u_{H,\text{cal}}^{(m)} + w(u_{\text{FM, cal}}^{(m)} - u_{H,\text{cal}}^{(m)})$.

Define difference vectors for each calibration sample m:

$$\Delta_{\text{FM}}^{(m)} = u_{\text{FM, cal}}^{(m)} - u_{H,\text{cal}}^{(m)}, \tag{47}$$

$$\Delta_{\text{true}}^{(m)} = u_{\text{true, cal}}^{(m)} - u_{H,\text{cal}}^{(m)}. \tag{48}$$

The error vector for sample m at weight w is:

$$e^{(m)}(w) = u(w) - u_{\text{true, cal}}^{(m)} = u_{H,\text{cal}}^{(m)} + w\Delta_{\text{FM}}^{(m)} - u_{\text{true, cal}}^{(m)} = w\Delta_{\text{FM}}^{(m)} - \Delta_{\text{true}}^{(m)}. \tag{49}$$

The empirical calibration loss objective $J : \mathbb{R} \to \mathbb{R}$ is the sum of squared L^2 norms:

$$J(w) = \sum_{m=1}^{M} \|e^{(m)}(w)\|^2 = \sum_{m=1}^{M} \|w\Delta_{\text{FM}}^{(m)} - \Delta_{\text{true}}^{(m)}\|^2 \tag{50}$$

$$= \sum_{m=1}^{M} \left(w^2 \|\Delta_{\text{FM}}^{(m)}\|^2 - 2w\langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle + \|\Delta_{\text{true}}^{(m)}\|^2 \right) \tag{51}$$

$$= w^2 \left(\sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2 \right) - 2w \left(\sum_{m=1}^{M} \langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle \right) + \sum_{m=1}^{M} \|\Delta_{\text{true}}^{(m)}\|^2. \tag{52}$$

Define scalar coefficients $A = \sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2$, $B = \sum_{m=1}^{M} \langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle$, and $C = \sum_{m=1}^{M} \|\Delta_{\text{true}}^{(m)}\|^2$. Then $J(w) = Aw^2 - 2Bw + C$.

First and Second Derivatives: Differentiating with respect to w:

$$J'(w) = 2Aw - 2B, \tag{53}$$

$$J''(w) = 2A = 2 \sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2. \tag{54}$$

Because $u_{\text{FM}} \neq u_H$ on non-trivial trajectories, $\|\Delta_{\text{FM}}^{(m)}\| > 0$, so $A > 0$ and $J''(w) > 0$ strictly. Therefore, $J(w)$ is a strictly convex quadratic function possessing a unique global minimum on $\mathbb{R}$.

Setting $J'(w^*) = 0$ yields the unconstrained minimizer:

$$2Aw^* - 2B = 0 \implies w^* = \frac{B}{A} = \frac{\sum_{m=1}^{M} \langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle}{\sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2}. \tag{55}$$

Adding $\epsilon = 10^{-12}$ to the denominator prevents numerical division by zero in degenerate floating-point configurations where $\Delta_{\text{FM}} \equiv 0$. Constraining $w \in [0, 1]$ to preserve convex combination properties yields the projected minimizer:

$$w^* = \text{clip} \left(\frac{\sum_{m=1}^{M} \langle \Delta_{\text{FM}}^{(m)}, \Delta_{\text{true}}^{(m)} \rangle}{\sum_{m=1}^{M} \|\Delta_{\text{FM}}^{(m)}\|^2 + \epsilon},\ 0.0,\ 1.0 \right). \tag{56}$$

Proof of Calibration Sample Non-Degradation: Because $J(w)$ is strictly convex and the constraint set $[0, 1]$ is convex and contains both endpoints $w = 0$ and $w = 1$, the constrained minimizer $w^* = \arg\min_{w \in [0,1]} J(w)$ must satisfy:

$$J(w^*) \leq J(0) \quad \text{and} \quad J(w^*) \leq J(1). \tag{57}$$

Since $w = 0$ corresponds to the standalone kinematic bridge (u_H) and $w = 1$ corresponds to the standalone foundation model (u_{FM}), calibrated consensus is mathematically guaranteed never to underperform either candidate model on the calibration dataset:

$$J(w^*) \leq \min \left(J_{\text{kinematic}},\ J_{\text{FM}} \right). \tag{58}$$

E.5 Status of the Residual-Informed Weight

Equation 16 resembles an inverse-variance fusion rule only under additional assumptions: the candidate errors must be suitably decorrelated, the residual must be locally informative in all relevant error directions, and the discretized defect must have a calibrated relation to state error. These assumptions do not hold automatically for nonlinear PDEs, incomplete operators, or sparse endpoint observations. We therefore do *not* invoke Gauss–Markov optimality, a generic coercivity inequality, or a distribution-free error guarantee for this weight. In this paper it is an optional residual-informed heuristic; the reported headline fusion result uses the supervised, calibration-selected weight in Equation 14. A proper evaluation of this heuristic requires residual–state-error correlation and reliability measurements on held-out trajectories.

E.6 Von Neumann Stability Analysis and CFL Numerical Explosion

We provide a standard modal Von Neumann calculation (Courant et al., 1928; Trefethen, 2000) for the representative centered-difference advection/RK4 macro-step used in the motivation. It is not a general statement about all numerical time-stepping methods.

Consider the one-dimensional advection equation on the periodic domain $[0, 1]$:

$$\frac{\partial u}{\partial t} + c\frac{\partial u}{\partial x} = 0, \quad c > 0. \tag{59}$$

Discretize spatial domain $[0, 1]$ with N uniform grid points $x_j = j\Delta x$, $\Delta x = 1/N$. Discretizing the spatial derivative $\partial_x u$ with second-order central finite differences yields the system of ordinary differential equations:

$$\frac{du_j}{dt} = -\frac{c}{2\Delta x}(u_{j+1} - u_{j-1}). \tag{60}$$

In Fourier modal representation, let $u_j(t) = \sum_k \widehat{u}_k(t)e^{ikx_j} = \sum_k \widehat{u}_k(t)e^{ij\theta}$, where $\theta = k\Delta x \in [-\pi, \pi]$ is the normalized wavenumber. Substituting this Fourier mode into the semi-discrete ODE:

$$\frac{d\widehat{u}_k}{dt}e^{ij\theta} = -\frac{c}{2\Delta x}\widehat{u}_k(t)\Big(e^{i(j+1)\theta} - e^{i(j-1)\theta}\Big) = -\frac{c}{2\Delta x}\widehat{u}_k(t)e^{ij\theta}\Big(e^{i\theta} - e^{-i\theta}\Big). \tag{61}$$

Recalling Euler's identity $\sin\theta = \frac{e^{i\theta}-e^{-i\theta}}{2i} \implies e^{i\theta} - e^{-i\theta} = 2i\sin\theta$:

$$\frac{d\widehat{u}_k}{dt} = -\frac{c}{2\Delta x}(2i\sin\theta)\widehat{u}_k(t) = -i\frac{c\sin\theta}{\Delta x}\widehat{u}_k(t). \tag{62}$$

This is a decoupled scalar linear ODE $\frac{d\widehat{u}_k}{dt} = \lambda_k\widehat{u}_k$, with purely imaginary eigenvalue:

$$\lambda_k = -i\frac{c\sin(k\Delta x)}{\Delta x} \in i\mathbb{R}. \tag{63}$$

The maximum imaginary magnitude occurs at $k\Delta x = \pi/2$ (the Nyquist mode), where $|\sin(k\Delta x)| = 1$:

$$|\lambda_{\max}| = \frac{c}{\Delta x}. \tag{64}$$

Classical Explicit Fourth-Order Runge-Kutta (RK4): Applying classical RK4 with time step Δt to $\frac{du}{dt} = \lambda u$ yields the discrete update:

$$\widehat{u}_k(t + \Delta t) = R(z)\widehat{u}_k(t), \quad z = \lambda_k\Delta t, \tag{65}$$

where the RK4 stability polynomial is:

$$R(z) = 1 + z + \frac{z^2}{2} + \frac{z^3}{6} + \frac{z^4}{24}. \tag{66}$$

For our purely imaginary eigenvalues $z = -i\mu$ where $\mu = \frac{c\Delta t}{\Delta x}\sin(k\Delta x) \in \mathbb{R}$:

$$z = -i\mu, \tag{67}$$

$$z^2 = (-i\mu)^2 = -\mu^2, \tag{68}$$

$$z^3 = -i\mu(-\mu^2) = i\mu^3, \tag{69}$$

$$z^4 = (-\mu^2)^2 = \mu^4. \tag{70}$$

Substitute these powers into $R(-i\mu)$:

$$R(-i\mu) = 1 - i\mu - \frac{\mu^2}{2} + \frac{i\mu^3}{6} + \frac{\mu^4}{24} = \left(1 - \frac{\mu^2}{2} + \frac{\mu^4}{24}\right) - i\left(\mu - \frac{\mu^3}{6}\right). \tag{71}$$

Let $\mathrm{Re}(R) = 1 - \frac{\mu^2}{2} + \frac{\mu^4}{24}$ and $\mathrm{Im}(R) = -\left(\mu - \frac{\mu^3}{6}\right)$. We compute the squared magnitude $|R(-i\mu)|^2 = \mathrm{Re}(R)^2 + \mathrm{Im}(R)^2$:

$$\mathrm{Re}(R)^2 = \left(1 - \frac{\mu^2}{2} + \frac{\mu^4}{24}\right)^2 = 1 - \mu^2 + \frac{\mu^4}{12} + \frac{\mu^4}{4} - \frac{\mu^6}{24} + \frac{\mu^8}{576} = 1 - \mu^2 + \frac{\mu^4}{3} - \frac{\mu^6}{24} + \frac{\mu^8}{576}, \tag{72}$$

$$\mathrm{Im}(R)^2 = \left(\mu - \frac{\mu^3}{6}\right)^2 = \mu^2 - 2\frac{\mu^4}{6} + \frac{\mu^6}{36} = \mu^2 - \frac{\mu^4}{3} + \frac{\mu^6}{36}. \tag{73}$$

Summing the real and imaginary squares:

$$|R(-i\mu)|^2 = \left(1 - \mu^2 + \frac{\mu^4}{3} - \frac{\mu^6}{24} + \frac{\mu^8}{576}\right) + \left(\mu^2 - \frac{\mu^4}{3} + \frac{\mu^6}{36}\right) \tag{74}$$

$$= 1 + (-\mu^2 + \mu^2) + \left(\frac{\mu^4}{3} - \frac{\mu^4}{3}\right) + \left(-\frac{1}{24} + \frac{1}{36}\right)\mu^6 + \frac{\mu^8}{576} \tag{75}$$

$$= 1 + \left(-\frac{3}{72} + \frac{2}{72}\right)\mu^6 + \frac{\mu^8}{576} \tag{76}$$

$$= 1 - \frac{\mu^6}{72} + \frac{\mu^8}{576} = 1 - \frac{\mu^6}{72}\left(1 - \frac{\mu^2}{8}\right). \tag{77}$$

Exact Stability Boundary: Stability requires $|R(-i\mu)|^2 \le 1$:

$$1 - \frac{\mu^6}{72}\left(1 - \frac{\mu^2}{8}\right) \le 1 \iff -\frac{\mu^6}{72}\left(1 - \frac{\mu^2}{8}\right) \le 0 \iff 1 - \frac{\mu^2}{8} \ge 0 \iff \mu^2 \le 8. \tag{78}$$

Therefore, RK4 is stable on the imaginary axis if and only if:

$$|\mu| \le \sqrt{8} = 2\sqrt{2} \approx 2.8284. \tag{79}$$

Catastrophic Explosion at Foundation Model Cadence: For an arbitrary spatial discretization with grid spacing $\Delta x = L/N$, the advective Courant number scales as $\mu = c\Delta t/\Delta x = cN\Delta t/L$. In representative physical benchmark configurations ($N = 128$, $\Delta x = 1/128$, $c = 1.5\,\mathrm{m/s}$, $\Delta t = 0.10\,\mathrm{s}$):

$$\mu = \frac{c\Delta t}{\Delta x} = \frac{1.5 \times 0.10}{1/128} = 19.2. \tag{80}$$

Comparing this to the stability limit: $\mu = 19.2 \gg 2\sqrt{2} \approx 2.83$. The stability condition is violated by a factor of $6.8\times$.

We evaluate the amplification factor $|R(-i\mu)|$ at $\mu = 19.2$:

$$\mathrm{Re}(R) = 1 - \frac{19.2^2}{2} + \frac{19.2^4}{24} = 1 - \frac{368.64}{2} + \frac{135{,}895.45}{24} = 1 - 184.32 + 5{,}662.31 = 5{,}478.99, \tag{81}$$

$$\mathrm{Im}(R) = -\left(19.2 - \frac{19.2^3}{6}\right) = -\left(19.2 - \frac{7{,}077.89}{6}\right) = -(19.2 - 1{,}179.65) = 1{,}160.45. \tag{82}$$

The magnitude of the amplification factor is:

$$|R(-i19.2)| = \sqrt{5478.99^2 + 1160.45^2}$$

$$= \sqrt{30{,}019{,}331 + 1{,}346{,}644} = \sqrt{31{,}365{,}975} \approx \mathbf{5{,}600.5}. \tag{83}$$

In this *particular centered-difference advection discretization*, one RK4 macro-step magnifies the selected modal component by about $5{,}600$. Repeating the same unstable macro-step would multiply that component by approximately $(5600)^2$.

This scoped calculation illustrates why a single coarse explicit macro-step can be a poor baseline. It does not rule out stable sub-stepping, implicit/exponential integrators, adaptive solvers, or alternative spatial discretizations; it motivates the zero-substep endpoint setting considered by KMF.

E.7 Derivation of the Hermite-Simpson Quadrature Defect

We derive the Hermite-Simpson quadrature defect $\mathcal{R}_{\text{Simpson}}(u)$ defined in Equation 15. Integrating the governing differential equation $\frac{du}{dt} = F(u)$ over $[t_0, t_1]$ gives the exact integral relation:

$$u(t_1) - u(t_0) = \int_{t_0}^{t_1} F(u(t))dt. \tag{84}$$

Simpson's 1/3 quadrature rule (Richard et al., 2016) approximates the definite integral of a function $f(t)$ on $[t_0, t_1]$ by replacing $f(t)$ with its quadratic interpolant passing through the endpoints and midpoint $t_{1/2} = t_0 + 0.5\Delta t$; the same construction is standard in Hermite–Simpson direct collocation (Kelly, 2017):

$$\int_{t_0}^{t_1} f(t)dt = \frac{\Delta t}{6}\Big(f(t_0) + 4f(t_{1/2}) + f(t_1)\Big) - \frac{\Delta t^5}{2880} f^{(4)}(\zeta), \quad \zeta \in (t_0, t_1). \tag{85}$$

Substituting $f(t) = F(u(t))$ and dividing through by Δt:

$$\frac{u(t_1) - u(t_0)}{\Delta t} - \frac{1}{6}\Big(F(u_0) + 4F(u(t_{1/2})) + F(u_1)\Big) = -\frac{\Delta t^4}{2880} \left.\frac{d^4F(u)}{dt^4}\right|_{\zeta}. \tag{86}$$

The left-hand side is precisely the Hermite-Simpson defect $\mathcal{R}_{\text{Simpson}}(u)$ evaluated at candidate midpoint u. When evaluated on candidate fields, non-zero defect values quantify the degree of physical inconsistency with fourth-order temporal fidelity $\mathcal{O}(\Delta t^4)$.

E.8 Step-by-Step Derivation and Algorithm for Defect-Corrected Autoregressive Rollout

This subsection provides the formal calculus derivation of the tentative midpoint state $\tilde{u}_{k+1/2}$ and the defect-corrected autoregressive rollout procedure (Section 4.4).

Algebraic Derivation of the Tentative Midpoint Formula: In continuous-time autoregressive forecasting, an accepted state $\hat{u}_k$ at t_k is advanced across macro-step $\Delta t = t_{k+1} - t_k$ to candidate state $\tilde{u}_{k+1} = f_\theta(\hat{u}_k)$. The true physical trajectory $u(t)$ satisfies $\partial_t u(t) = F(u(t))$. Because F is an instantaneous spatial differential operator, the boundary physical velocities are evaluated algebraically without numerical stepping:

$$f_k = F(\hat{u}_k), \qquad \tilde{f}_{k+1} = F(\tilde{u}_{k+1}). \tag{87}$$

We map interval $[t_k, t_{k+1}]$ to the normalized coordinate $s = (t - t_k)/\Delta t \in [0, 1]$. The two-point kinematic Hermite interpolant $u_H(s)$ connecting $(\hat{u}_k, f_k)$ and $(\tilde{u}_{k+1}, \tilde{f}_{k+1})$ is defined by:

$$u_H(s) = h_{00}(s)\hat{u}_k + h_{10}(s)\tilde{u}_{k+1} + h_{01}(s)\Delta t f_k + h_{11}(s)\Delta t \tilde{f}_{k+1}. \tag{88}$$

Evaluating the cubic Hermite basis functions derived in Section E.1 at the exact temporal midpoint $s = 0.5$:

$$h_{00}(0.5) = 2(0.5)^3 - 3(0.5)^2 + 1 = 2(0.125) - 3(0.25) + 1 = 0.25 - 0.75 + 1 = \frac{1}{2}, \tag{89}$$

$$h_{10}(0.5) = -2(0.5)^3 + 3(0.5)^2 = -2(0.125) + 3(0.25) = -0.25 + 0.75 = \frac{1}{2}, \tag{90}$$

$$h_{01}(0.5) = (0.5)^3 - 2(0.5)^2 + 0.5 = 0.125 - 2(0.25) + 0.5 = 0.125 - 0.5 + 0.5 = \frac{1}{8}, \tag{91}$$

$$h_{11}(0.5) = (0.5)^3 - (0.5)^2 = 0.125 - 0.25 = -\frac{1}{8}. \tag{92}$$

Substituting these exact rational coefficients into $u_H(0.5)$ yields:

$$\tilde{u}_{k+1/2} = u_H(0.5) = \frac{1}{2}\big(\hat{u}_k + \tilde{u}_{k+1}\big) + \frac{\Delta t}{8}\Big(F(\hat{u}_k) - F(\tilde{u}_{k+1})\Big). \tag{93}$$

This algebraic identity shows that the tentative midpoint consists of the arithmetic secant mean $\frac{1}{2}(\hat{u}_k + \tilde{u}_{k+1})$ augmented by a second-order curvature correction $\frac{\Delta t}{8}(F(\hat{u}_k) - F(\tilde{u}_{k+1}))$. It uses zero temporal integration substeps, relying exclusively on instantaneous evaluations of F at macro-step endpoints. It has the form of the internal stage of classical fourth-order Lobatto IIIA implicit Runge–Kutta collocation (Hairer et al., 1993), without solving a collocation system.

Hermite-Simpson Collocation Defect and Quadrature Relaxation: With $\tilde{u}_{k+1/2}$ established, evaluating the spatial governing operator at the midpoint yields the physical velocity $\tilde{f}_{k+1/2} = F(\tilde{u}_{k+1/2})$. Integrating $\partial_t u = F(u)$ across $[t_k, t_{k+1}]$ via Simpson's rule produces the fourth-order quadrature prediction:

$$u_{\text{Simp}} = \hat{u}_k + \frac{\Delta t}{6}\Big(F(\hat{u}_k) + 4F(\tilde{u}_{k+1/2}) + F(\tilde{u}_{k+1})\Big). \tag{94}$$

The algebraic discrepancy between the unconstrained neural candidate $\tilde{u}_{k+1}$ and the fourth-order Simpson prediction u_{Simp} is precisely:

$$\begin{aligned}\tilde{u}_{k+1} - u_{\text{Simp}} &= \Delta t \left[\frac{\tilde{u}_{k+1} - \hat{u}_k}{\Delta t} - \frac{1}{6}\Big(F(\hat{u}_k) + 4F(\tilde{u}_{k+1/2}) + F(\tilde{u}_{k+1})\Big)\right] \\ &= \Delta t\, \mathcal{R}_{\text{Simpson}}(\tilde{u}_{k+1/2}).\end{aligned} \tag{95}$$

Consequently, the defect relaxation step in Equation 17:

$$\tilde{u}^{\text{corr}}_{k+1} = \tilde{u}_{k+1} - \alpha^* \Delta t\, \mathcal{R}_{\text{Simpson}}(\tilde{u}_{k+1/2}) = (1-\alpha^*)\tilde{u}_{k+1} + \alpha^* u_{\text{Simp}}, \tag{96}$$

is an exact convex relaxation between the raw neural forecast $\tilde{u}_{k+1}$ and the fourth-order physical Simpson quadrature state u_{Simp}. In high-Reynolds fluid simulations, an exponential Fourier filter $\sigma_{\text{smooth}}(k) = \exp(-36(\|k\|/k_{\text{Nyquist}})^{36})$ attenuates high-wavenumber defect noise before relaxation. We fix these filter constants before held-out evaluation and do not tune them per trajectory. Finally, non-expansive projection $\mathcal{P}_{\mathcal{M}}$ enforces physical invariants (such as incompressibility via $\mathcal{P}_{\text{div-free}}$). The complete procedure is formalized in Algorithm 1.

Algorithm 1 Continuous-Time Defect-Corrected Autoregressive Rollout

Require: Initial physical state $\hat{u}_0 \in \mathcal{X}$, rollout horizon K, cadence Δt, foundation model f_θ, physical operator F, manifold projection $\mathcal{P}_{\mathcal{M}}$, calibrated relaxation factor α^*.

for $k = 0, 1, \ldots, K-1$ **do**
  Evaluate raw foundation model candidate: $\tilde{u}_{k+1} = f_\theta(\hat{u}_k)$.
  if $\alpha^* > 0$ **then**
    Compute boundary spatial physical derivatives: $f_k = F(\hat{u}_k)$ and $\tilde{f}_{k+1} = F(\tilde{u}_{k+1})$.
    Compute closed-form tentative midpoint: $\tilde{u}_{k+1/2} = \frac{1}{2}(\hat{u}_k + \tilde{u}_{k+1}) + \frac{\Delta t}{8}(f_k - \tilde{f}_{k+1})$.
    Compute midpoint physical derivative: $\tilde{f}_{k+1/2} = F(\tilde{u}_{k+1/2})$.
    Compute Hermite-Simpson defect: $\mathcal{R}_{\text{Simpson}} = \frac{\tilde{u}_{k+1} - \hat{u}_k}{\Delta t} - \frac{1}{6}(f_k + 4\tilde{f}_{k+1/2} + \tilde{f}_{k+1})$.
    Apply Fourier spectral filter to defect: $\mathcal{R}_{\text{filt}} = \mathcal{F}^{-1}\big(\mathcal{F}(\mathcal{R}_{\text{Simpson}}) \cdot \sigma_{\text{smooth}}\big)$.
    Relax candidate state: $\tilde{u}^{\text{corr}}_{k+1} = \tilde{u}_{k+1} - \alpha^* \Delta t\, \mathcal{R}_{\text{filt}}$.
  else
    Set uncorrected candidate: $\tilde{u}^{\text{corr}}_{k+1} = \tilde{u}_{k+1}$.
  end if
  Enforce invariant manifold: $\hat{u}_{k+1} = \mathcal{P}_{\mathcal{M}}(\tilde{u}^{\text{corr}}_{k+1})$.
  Continuous sub-cadence trajectory: For any query $t \in (t_k, t_{k+1})$, evaluate $u_H((t - t_k)/\Delta t; \hat{u}_k, \hat{u}_{k+1})$.
end for
return Discrete trajectory $\{\hat{u}_k\}_{k=0}^K$ and continuous-time bridge $u_H(t)$.

E.9 Empirical Convergence Rate Verification of Theorem 1

To quantitatively verify the theoretical fourth-order convergence bound $\|u(t) - u_H(s)\| \leq \frac{1}{384}\Delta t^4 \|u^{(4)}\|_\infty$ proven in Theorem 1, we evaluate the kinematic Hermite bridge across five successively halved temporal cadences: $\Delta t \in \{0.2000\,\text{s}, 0.1000\,\text{s}, 0.0500\,\text{s}, 0.0250\,\text{s}, 0.0125\,\text{s}\}$ on smooth non-linear fluid vortex trajectories. Table 13 reports the empirical root-mean-square error at the interval midpoint ($s = 0.5$) along with the empirical order of convergence computed via logarithmic ratios:

$$p = \log_2\left(\frac{\text{RMSE}(\Delta t)}{\text{RMSE}(\Delta t/2)}\right). \tag{97}$$

The observed near-fifth-order midpoint slope in this smooth diagnostic is compatible with additional symmetry/cancellation in that particular trajectory. It is not a general fifth-order claim. Theorem 1 remains a fourth-order upper bound under its stated assumptions; this table is only a numerical sanity check.

Table 13: **Empirical Convergence Rate of the Kinematic Hermite Bridge.** Verification of local truncation error scaling at midpoint $s = 0.5$ across successively halved cadences.

Cadence Δt	Query Coordinate s	Empirical Midpoint RMSE	Error Ratio	Observed Convergence Order p
0.2000 s	0.5	2.358×10^{-3}	–	Baseline
0.1000 s	0.5	7.914×10^{-5}	29.80	**4.90**
0.0500 s	0.5	2.517×10^{-6}	31.44	**4.97**
0.0250 s	0.5	7.901×10^{-8}	31.86	**4.99**
0.0125 s	0.5	2.472×10^{-9}	31.96	**5.00**

Algorithm 2 KMF Foundation-Model Temporal Refinement

Require: Successive model states $\tilde{u}_k, \tilde{u}_{k+1} \in \mathcal{X}$, interval Δt, query time $t = t_k + s\Delta t$, foundation model f_θ, physical operator F, calibration trajectories $\mathcal{D}_{\text{cal}}$.

1: Generate or receive the next endpoint: $\tilde{u}_{k+1} = f_\theta(\tilde{u}_k; \Delta t)$.
2: Compute endpoint rates: $f_k = F(\tilde{u}_k)$ and $f_{k+1} = F(\tilde{u}_{k+1})$.
3: Evaluate the Hermite bridge $u_H(s) = h_{00}(s)\tilde{u}_k + h_{10}(s)\tilde{u}_{k+1} + h_{01}(s)f_k + h_{11}(s)f_{k+1}$.
4: Evaluate a direct model query $u_{\text{FM}} = f_\theta(\tilde{u}_k; s\Delta t)$ when the interface provides one.
5: **if** a direct query and $\mathcal{D}_{\text{cal}}$ are available **then**
6: Compute w^* using Equation 14 on $\mathcal{D}_{\text{cal}}$.
7: **else**
8: Use the bridge alone, or the declared residual-informed heuristic when explicitly enabled.
9: **end if**
10: Set $u_{\text{blend}} = (1-w^*)u_H(s) + w^* u_{\text{FM}}$ when a direct query is available, otherwise $u_{\text{blend}} = u_H(s)$.
11: Enforce the available admissibility operation: $u^*(t) = \mathcal{P}_{\mathcal{M}}(u_{\text{blend}})$.
12: **return** Refined state $u^*(t)$.

E.10 KMF Temporal-Refinement Algorithm

E.11 Real-Pipeline Latency Audit

E.12 Fair Solver-Control Audit

Table 15 gives the complete FNS-KF, ACE, and Wave-Gauss solver-control audit. The "FM endpoint?" column makes the information asymmetry explicit: endpoint-conditioned rows use the FM-produced future endpoint pair, whereas causal RK rows use only the initial state. Latencies are online correction or solver time and exclude any common FM forward pass.

E.13 Benchmark Matrix

The following appendix tables report the 126 evaluated configurations, including the requested cadence, the endpoint span used by Task 1, all baselines, the calibration-selected KMF value, and the corresponding causal rollout fields. Missing combinations are reported as unavailable. In particular, longer Wave-Gauss rows are limited by the number of available raw frames. The primary discussion intentionally focuses Task 1 on Poseidon and CNO, while these tables retain the completed transfer rows for auditability.

Table 14: **Real-pipeline latency audit.** Synchronized wall-clock time on an NVIDIA A100-SXM4-40GB, CUDA 12.4, PyTorch 2.4, batch size one, and a 128×128 field.

Pipeline configuration	**Latency**	**vs. FM**
Raw foundation model (Poseidon-B)	96.33 ms	$1.00\times$
Two RHS evaluations	2.40 ms	$0.025\times$
Helmholtz–Leray projection	0.52 ms	$0.005\times$
Bridge arithmetic	0.41 ms	$0.004\times$
FM + KMF operator pipeline	**99.66** ms	$\mathbf{1.03}\times$

Table 15: **Fair endpoint-conditioned versus causal solver audit.** RMSE is measured at the midpoint of a 0.10-second span. KMF and linear rows receive an FM endpoint; RK controls receive only the initial state. The Wave-Gauss explicit controls are numerically unstable at this span, so their large finite values are retained rather than omitted.

PDE	**Method**	**FM endpoint?**	**RMSE**	**Online ms**	**RHS evaluations**
FNS-KF	FM endpoint	Yes	0.012534	93.349	0
	Linear bridge	Yes	0.010176	0.510	0
	KMF Hermite bridge	Yes	0.008928	2.251	2
	KMF calibrated fusion	Yes	**0.008390**	95.858	2
	RK2, 1 substep	No	0.036589	1.946	2
	RK2, 2 substeps	No	0.036591	2.979	4
	RK2, 4 substeps	No	0.036592	5.033	8
	RK2, 8 substeps	No	0.036592	9.114	16
	RK2, 16 substeps	No	0.036592	17.735	32
	RK4, 1 substep	No	0.036592	3.076	4
	RK4, 2 substeps	No	0.036592	5.201	8
	RK4, 4 substeps	No	0.036592	9.458	16
	RK4, 8 substeps	No	0.036592	17.929	32
	RK4, 16 substeps	No	0.036592	34.952	64
ACE	FM endpoint	Yes	0.168399	0.184	0
	Linear bridge	Yes	0.099233	0.386	0
	KMF Hermite bridge	Yes	**0.098376**	0.779	2
	KMF calibrated fusion	Yes	**0.098376**	0.784	2
	RK2, 1 substep	No	0.022755	0.756	2
	RK2, 2 substeps	No	0.022755	1.153	4
	RK2, 4 substeps	No	0.022755	1.841	8
	RK2, 8 substeps	No	0.022755	3.264	16
	RK2, 16 substeps	No	0.022755	6.201	32
	RK4, 1 substep	No	0.022755	1.158	4
	RK4, 2 substeps	No	0.022755	1.883	8
	RK4, 4 substeps	No	0.022755	3.472	16
	RK4, 8 substeps	No	0.022755	6.444	32
	RK4, 16 substeps	No	0.022755	12.601	64
Wave-Gauss	FM endpoint	Yes	0.092020	0.190	0
	Linear bridge	Yes	0.074334	0.588	0
	KMF Hermite bridge	Yes	**0.069417**	0.917	2
	KMF calibrated fusion	Yes	0.081420	1.218	2
	RK2, 1 substep	No	85.026716	0.909	2
	RK2, 2 substeps	No	529611.47	1.445	4
	RK2, 4 substeps	No	3.4124×10^{14}	2.553	8
	RK2, 8 substeps	No	4.6613×10^{30}	4.761	16
	RK2, 16 substeps	No	6.9847×10^{60}	9.024	32
	RK4, 1 substep	No	1641249.4	1.251	4
	RK4, 2 substeps	No	2.8936×10^{15}	2.200	8
	RK4, 4 substeps	No	6.1763×10^{32}	4.086	16
	RK4, 8 substeps	No	1.1680×10^{66}	7.756	32
	RK4, 16 substeps	No	7.8260×10^{123}	15.095	64

Table 16: FM-endpoint temporal-refinement matrix. Every row is an evaluated configuration; unavailable dataset or contract pairs are not imputed.

PDE	Model	req. Δt	span Δt	Linear	Raw FM	Hermite	Calibrated	Gain/Lin.	Gain/FM
ACE	Polymathic-FNO	0.05	0.10	0.21557	N/A	0.21256	0.21256	+1.40%	N/A
ACE	Polymathic-FNO	0.10	0.10	0.21557	N/A	0.21256	0.21256	+1.40%	N/A
ACE	Polymathic-FNO	0.20	0.20	0.24856	N/A	0.24251	0.24251	+2.43%	N/A
ACE	Polymathic-FNO	0.30	0.30	0.28259	N/A	0.27350	0.27350	+3.22%	N/A
ACE	Polymathic-FNO	0.40	0.40	0.31723	N/A	0.30510	0.30510	+3.82%	N/A
ACE	Polymathic-TFNO	0.05	0.10	0.21431	N/A	0.21123	0.21123	+1.44%	N/A
ACE	Polymathic-TFNO	0.10	0.10	0.21431	N/A	0.21123	0.21123	+1.44%	N/A
ACE	Polymathic-TFNO	0.20	0.20	0.24686	N/A	0.24071	0.24071	+2.49%	N/A
ACE	Polymathic-TFNO	0.30	0.30	0.28056	N/A	0.27135	0.27135	+3.28%	N/A
ACE	Polymathic-TFNO	0.40	0.40	0.31495	N/A	0.30272	0.30272	+3.88%	N/A
ACE	Polymathic-UNetConvNext	0.05	0.10	0.04039	N/A	0.04084	0.04084	-1.12%	N/A
ACE	Polymathic-UNetConvNext	0.10	0.10	0.04039	N/A	0.04084	0.04084	-1.12%	N/A
ACE	Polymathic-UNetConvNext	0.20	0.20	0.07388	N/A	0.07443	0.07443	-0.75%	N/A
ACE	Polymathic-UNetConvNext	0.30	0.30	0.10952	N/A	0.11012	0.11012	-0.55%	N/A
ACE	Polymathic-UNetConvNext	0.40	0.40	0.14612	N/A	0.14676	0.14676	-0.44%	N/A
ACE	Poseidon-B	0.05	0.10	0.12361	0.16839	0.12230	0.12230	+1.06%	+27.37%
ACE	Poseidon-B	0.10	0.10	0.12361	0.16839	0.12230	0.12230	+1.06%	+27.37%
ACE	Poseidon-B	0.20	0.20	0.19017	0.24800	0.18621	0.18621	+2.08%	+24.91%
ACE	Poseidon-B	0.30	0.30	0.22637	0.32269	0.22018	0.22018	+2.73%	+31.77%
ACE	Poseidon-B	0.40	0.40	0.23925	0.38204	0.23304	0.23304	+2.60%	+39.00%
ACE	Poseidon-L	0.05	0.10	0.11382	0.13448	0.11259	0.11259	+1.09%	+16.28%
ACE	Poseidon-L	0.10	0.10	0.11382	0.13448	0.11259	0.11259	+1.09%	+16.28%
ACE	Poseidon-L	0.20	0.20	0.17910	0.22836	0.17545	0.17545	+2.04%	+23.17%
ACE	Poseidon-L	0.30	0.30	0.21763	0.30379	0.21177	0.21177	+2.69%	+30.29%
ACE	Poseidon-L	0.40	0.40	0.23408	0.35981	0.22783	0.22783	+2.67%	+36.68%
ACE	Poseidon-T	0.05	0.10	0.11793	0.19875	0.11664	0.11664	+1.09%	+41.31%
ACE	Poseidon-T	0.10	0.10	0.11793	0.19875	0.11664	0.11664	+1.09%	+41.31%
ACE	Poseidon-T	0.20	0.20	0.18499	0.23659	0.18115	0.18115	+2.08%	+23.43%
ACE	Poseidon-T	0.30	0.30	0.22464	0.31215	0.21836	0.21836	+2.79%	+30.05%
ACE	Poseidon-T	0.40	0.40	0.23845	0.37162	0.23187	0.23187	+2.76%	+37.61%
FNS-KF	CNO-FM	0.05	0.10	0.03110	0.09023	0.03059	0.03059	+1.67%	+66.10%
FNS-KF	CNO-FM	0.10	0.10	0.03110	0.09023	0.03059	0.03059	+1.67%	+66.10%
FNS-KF	CNO-FM	0.20	0.20	0.03680	0.05515	0.03479	0.03479	+5.47%	+36.92%
FNS-KF	CNO-FM	0.30	0.30	0.05063	0.04555	0.04723	0.04270	+15.66%	+6.25%
FNS-KF	CNO-FM	0.40	0.40	0.06441	0.04502	0.06006	0.04437	+31.11%	+1.44%
FNS-KF	Dpot-S	0.10	0.10	0.02380	N/A	0.02371	0.02371	+0.36%	N/A
FNS-KF	Dpot-Ti	0.10	0.10	0.02981	N/A	0.02991	0.02991	-0.33%	N/A
FNS-KF	Morph-S	0.10	0.10	0.02641	N/A	0.02633	0.02633	+0.33%	N/A
FNS-KF	Morph-Ti	0.10	0.10	0.02611	N/A	0.02600	0.02600	+0.41%	N/A
FNS-KF	Polymathic-FNO	0.05	0.10	0.12173	N/A	0.12172	0.12172	+0.01%	N/A
FNS-KF	Polymathic-FNO	0.10	0.10	0.12173	N/A	0.12172	0.12172	+0.01%	N/A
FNS-KF	Polymathic-FNO	0.20	0.20	0.12635	N/A	0.12598	0.12598	+0.29%	N/A
FNS-KF	Polymathic-FNO	0.30	0.30	0.13229	N/A	0.13148	0.13148	+0.62%	N/A
FNS-KF	Polymathic-FNO	0.40	0.40	0.13863	N/A	0.13742	0.13742	+0.87%	N/A
FNS-KF	Polymathic-TFNO	0.05	0.10	0.12815	N/A	0.12813	0.12813	+0.02%	N/A
FNS-KF	Polymathic-TFNO	0.10	0.10	0.12815	N/A	0.12813	0.12813	+0.02%	N/A
FNS-KF	Polymathic-TFNO	0.20	0.20	0.13251	N/A	0.13214	0.13214	+0.28%	N/A
FNS-KF	Polymathic-TFNO	0.30	0.30	0.13816	N/A	0.13735	0.13735	+0.59%	N/A
FNS-KF	Polymathic-TFNO	0.40	0.40	0.14421	N/A	0.14302	0.14302	+0.83%	N/A
FNS-KF	Polymathic-UNetConvNext	0.05	0.10	0.03077	N/A	0.03080	0.03080	-0.13%	N/A
FNS-KF	Polymathic-UNetConvNext	0.10	0.10	0.03077	N/A	0.03080	0.03080	-0.13%	N/A
FNS-KF	Polymathic-UNetConvNext	0.20	0.20	0.05554	N/A	0.05559	0.05559	-0.10%	N/A
FNS-KF	Polymathic-UNetConvNext	0.30	0.30	0.07750	N/A	0.07759	0.07759	-0.11%	N/A
FNS-KF	Polymathic-UNetConvNext	0.40	0.40	0.09641	N/A	0.09654	0.09654	-0.14%	N/A
FNS-KF	Poseidon-B	0.05	0.10	0.01018	0.01248	0.00908	0.00873	+14.33%	+30.08%
FNS-KF	Poseidon-B	0.10	0.10	0.01018	0.01248	0.00908	0.00873	+14.33%	+30.08%
FNS-KF	Poseidon-B	0.20	0.20	0.02594	0.00580	0.02317	0.00580	+77.63%	+0.00%
FNS-KF	Poseidon-B	0.30	0.30	0.04339	0.00812	0.03924	0.00812	+81.29%	+0.00%
FNS-KF	Poseidon-B	0.40	0.40	0.05972	0.01076	0.05475	0.01076	+81.99%	+0.00%
FNS-KF	Poseidon-L	0.05	0.10	0.01030	0.01150	0.00918	0.00878	+14.71%	+23.63%
FNS-KF	Poseidon-L	0.10	0.10	0.01030	0.01150	0.00918	0.00878	+14.71%	+23.63%
FNS-KF	Poseidon-L	0.20	0.20	0.02602	0.00588	0.02325	0.00588	+77.41%	+0.00%
FNS-KF	Poseidon-L	0.30	0.30	0.04342	0.00813	0.03929	0.00813	+81.28%	+0.00%
FNS-KF	Poseidon-L	0.40	0.40	0.05967	0.01079	0.05470	0.01079	+81.92%	+0.00%
FNS-KF	Poseidon-T	0.05	0.10	0.01199	0.07945	0.01103	0.01103	+7.98%	+86.11%
FNS-KF	Poseidon-T	0.10	0.10	0.01199	0.07945	0.01103	0.01103	+7.98%	+86.11%
FNS-KF	Poseidon-T	0.20	0.20	0.02664	0.01187	0.02396	0.01179	+55.75%	+0.67%
FNS-KF	Poseidon-T	0.30	0.30	0.04398	0.01344	0.03993	0.01320	+70.00%	+1.85%
FNS-KF	Poseidon-T	0.40	0.40	0.06035	0.01488	0.05549	0.01487	+75.36%	+0.06%
NS-Gauss	CNO-FM	0.05	0.10	0.06106	0.03681	0.03296	0.02684	+56.05%	+27.09%
NS-Gauss	CNO-FM	0.10	0.10	0.06106	0.03681	0.03296	0.02684	+56.05%	+27.09%
NS-Gauss	CNO-FM	0.20	0.20	0.12657	0.02974	0.09114	0.02974	+76.50%	+0.00%
NS-Gauss	CNO-FM	0.30	0.30	0.17864	0.03730	0.14708	0.03730	+79.12%	+0.00%
NS-Gauss	CNO-FM	0.40	0.40	0.22013	0.04075	0.20168	0.04075	+81.49%	+0.00%
NS-Gauss	Dpot-S	0.10	0.10	0.07581	N/A	0.07799	0.07799	-2.88%	N/A
NS-Gauss	Dpot-Ti	0.10	0.10	0.07670	N/A	0.08113	0.08113	-5.77%	N/A
NS-Gauss	Morph-S	0.10	0.10	0.09871	N/A	0.10018	0.10018	-1.49%	N/A
NS-Gauss	Morph-Ti	0.10	0.10	0.09954	N/A	0.09893	0.09893	+0.62%	N/A
NS-Gauss	Polymathic-FNO	0.05	0.10	0.31744	N/A	0.30985	0.30985	+2.39%	N/A
NS-Gauss	Polymathic-FNO	0.10	0.10	0.31744	N/A	0.30985	0.30985	+2.39%	N/A
NS-Gauss	Polymathic-FNO	0.20	0.20	0.34106	N/A	0.32671	0.32671	+4.21%	N/A
NS-Gauss	Polymathic-FNO	0.30	0.30	0.36452	N/A	0.35126	0.35126	+3.64%	N/A
NS-Gauss	Polymathic-FNO	0.40	0.40	0.38540	N/A	0.38036	0.38036	+1.31%	N/A
NS-Gauss	Polymathic-TFNO	0.05	0.10	0.31642	N/A	0.30862	0.30862	+2.46%	N/A

Continued on next page

Table 16 continued

PDE	Model	req. Δt	span Δt	Linear	Raw FM	Hermite	Calibrated	Gain/Lin.	Gain/FM
NS-Gauss	Polymathic-TFNO	0.10	0.10	0.31642	N/A	0.30862	0.30862	+2.46%	N/A
NS-Gauss	Polymathic-TFNO	0.20	0.20	0.34030	N/A	0.32540	0.32540	+4.38%	N/A
NS-Gauss	Polymathic-TFNO	0.30	0.30	0.36407	N/A	0.34984	0.34984	+3.91%	N/A
NS-Gauss	Polymathic-TFNO	0.40	0.40	0.38541	N/A	0.37889	0.37889	+1.69%	N/A
NS-Gauss	Polymathic-UNetConvNext	0.05	0.10	0.12538	N/A	0.13686	0.13686	-9.16%	N/A
NS-Gauss	Polymathic-UNetConvNext	0.10	0.10	0.12538	N/A	0.13686	0.13686	-9.16%	N/A
NS-Gauss	Polymathic-UNetConvNext	0.20	0.20	0.21397	N/A	0.24096	0.24096	-12.62%	N/A
NS-Gauss	Polymathic-UNetConvNext	0.30	0.30	0.28006	N/A	0.32551	0.32551	-16.23%	N/A
NS-Gauss	Polymathic-UNetConvNext	0.40	0.40	0.33113	N/A	0.39784	0.39784	-20.15%	N/A
NS-Gauss	Poseidon-B	0.05	0.10	0.06165	0.03004	0.03200	0.02509	+59.30%	+16.47%
NS-Gauss	Poseidon-B	0.10	0.10	0.06165	0.03004	0.03200	0.02509	+59.30%	+16.47%
NS-Gauss	Poseidon-B	0.20	0.20	0.13079	0.00298	0.09596	0.00298	+97.72%	-0.00%
NS-Gauss	Poseidon-B	0.30	0.30	0.18458	0.00376	0.15594	0.00376	+97.96%	-0.00%
NS-Gauss	Poseidon-B	0.40	0.40	0.22715	0.00425	0.21346	0.00425	+98.13%	-0.00%
NS-Gauss	Poseidon-L	0.05	0.10	0.06166	0.03693	0.03208	0.02740	+55.57%	+25.81%
NS-Gauss	Poseidon-L	0.10	0.10	0.06166	0.03693	0.03208	0.02740	+55.57%	+25.81%
NS-Gauss	Poseidon-L	0.20	0.20	0.13079	0.00243	0.09596	0.00243	+98.14%	+0.00%
NS-Gauss	Poseidon-L	0.30	0.30	0.18455	0.00314	0.15607	0.00314	+98.30%	+0.02%
NS-Gauss	Poseidon-L	0.40	0.40	0.22718	0.00357	0.21364	0.00357	+98.43%	+0.00%
NS-Gauss	Poseidon-T	0.05	0.10	0.06139	0.05336	0.03103	0.02873	+53.20%	+46.15%
NS-Gauss	Poseidon-T	0.10	0.10	0.06139	0.05336	0.03103	0.02873	+53.20%	+46.15%
NS-Gauss	Poseidon-T	0.20	0.20	0.13071	0.00635	0.09538	0.00635	+95.14%	+0.00%
NS-Gauss	Poseidon-T	0.30	0.30	0.18442	0.00796	0.15534	0.00796	+95.68%	-0.00%
NS-Gauss	Poseidon-T	0.40	0.40	0.22699	0.00858	0.21234	0.00858	+96.22%	+0.00%
Wave-Gauss	Polymathic-FNO	0.05	0.10	0.08871	N/A	0.09480	0.09480	-6.86%	N/A
Wave-Gauss	Polymathic-FNO	0.10	0.10	0.08871	N/A	0.09480	0.09480	-6.86%	N/A
Wave-Gauss	Polymathic-FNO	0.30	0.30	0.13343	N/A	0.40157	0.40157	-200.96%	N/A
Wave-Gauss	Polymathic-TFNO	0.05	0.10	0.13143	N/A	0.13043	0.13043	+0.76%	N/A
Wave-Gauss	Polymathic-TFNO	0.10	0.10	0.13143	N/A	0.13043	0.13043	+0.76%	N/A
Wave-Gauss	Polymathic-TFNO	0.30	0.30	0.16109	N/A	0.23867	0.23867	-48.15%	N/A
Wave-Gauss	Polymathic-UNetConvNext	0.05	0.10	0.07829	N/A	0.10846	0.10846	-38.53%	N/A
Wave-Gauss	Polymathic-UNetConvNext	0.10	0.10	0.07829	N/A	0.10846	0.10846	-38.53%	N/A
Wave-Gauss	Polymathic-UNetConvNext	0.30	0.30	0.17185	N/A	0.86046	0.86046	-400.69%	N/A
Wave-Gauss	Poseidon-B	0.05	0.10	0.08150	0.09202	0.07425	0.08537	-4.75%	+7.22%
Wave-Gauss	Poseidon-B	0.10	0.10	0.08150	0.09202	0.07425	0.08537	-4.75%	+7.22%
Wave-Gauss	Poseidon-B	0.30	0.30	0.14899	0.14727	0.22094	0.13335	+10.50%	+9.45%
Wave-Gauss	Poseidon-L	0.05	0.10	0.07320	0.06548	0.06584	0.07753	-5.92%	-18.41%
Wave-Gauss	Poseidon-L	0.10	0.10	0.07320	0.06548	0.06584	0.07753	-5.92%	-18.41%
Wave-Gauss	Poseidon-L	0.30	0.30	0.14399	0.12408	0.20361	0.11533	+19.91%	+7.05%
Wave-Gauss	Poseidon-T	0.05	0.10	0.07305	0.08673	0.07168	0.08218	-12.49%	+5.24%
Wave-Gauss	Poseidon-T	0.10	0.10	0.07305	0.08673	0.07168	0.08218	-12.49%	+5.24%
Wave-Gauss	Poseidon-T	0.30	0.30	0.14227	0.12330	0.23311	0.11525	+18.99%	+6.53%

Table 17: Causal rollout and dense-bridge matrix. The rollout columns use the calibration-selected correction and the dense columns are reported only where an intermediate truth state was available for scoring.

PDE	Model	Δt	Steps	Raw	Projected	KMF	Gain	α^*	Dense raw-lin.	Dense KMF
ACE	Polymathic-FNO	0.05	4	0.43837	0.43837	0.41625	+5.05%	0.200	N/A	N/A
ACE	Polymathic-FNO	0.10	4	0.52159	0.52159	0.49908	+4.32%	0.200	0.48867	0.42121
ACE	Polymathic-FNO	0.20	4	0.66727	0.66727	0.64402	+3.48%	0.200	0.61043	0.54165
ACE	Polymathic-FNO	0.30	3	0.71782	0.71782	0.68602	+4.43%	0.200	0.63813	0.54533
ACE	Polymathic-FNO	0.40	2	0.72555	0.72555	0.67793	+6.56%	0.200	0.61517	0.47827
ACE	Polymathic-TFNO	0.05	4	0.44140	0.44140	0.41927	+5.01%	0.200	N/A	N/A
ACE	Polymathic-TFNO	0.10	4	0.52349	0.52349	0.50107	+4.28%	0.200	0.49102	0.42279
ACE	Polymathic-TFNO	0.20	4	0.66790	0.66790	0.64487	+3.45%	0.200	0.61159	0.54220
ACE	Polymathic-TFNO	0.30	3	0.71747	0.71747	0.68593	+4.40%	0.200	0.63842	0.54487
ACE	Polymathic-TFNO	0.40	2	0.72388	0.72388	0.67658	+6.53%	0.200	0.61433	0.47650
ACE	Polymathic-UNetConvNext	0.05	4	0.11867	0.11852	0.10549	+11.10%	0.200	N/A	N/A
ACE	Polymathic-UNetConvNext	0.10	4	0.20085	0.20076	0.18434	+8.22%	0.200	0.16790	0.14747
ACE	Polymathic-UNetConvNext	0.20	4	0.35923	0.35916	0.33471	+6.82%	0.200	0.29716	0.27252
ACE	Polymathic-UNetConvNext	0.30	3	0.41211	0.41210	0.38515	+6.54%	0.200	0.32310	0.29845
ACE	Polymathic-UNetConvNext	0.40	2	0.41552	0.41551	0.38905	+6.37%	0.200	0.29140	0.26944
ACE	Poseidon-B	0.05	4	0.28498	0.28498	0.25460	+10.66%	0.200	N/A	N/A
ACE	Poseidon-B	0.10	4	0.39916	0.39916	0.36609	+8.28%	0.200	0.36665	0.30070
ACE	Poseidon-B	0.20	4	0.58108	0.58108	0.54553	+6.12%	0.200	0.52484	0.45244
ACE	Poseidon-B	0.30	3	0.62838	0.62837	0.58354	+7.14%	0.200	0.54945	0.45729
ACE	Poseidon-B	0.40	2	0.60250	0.60250	0.55506	+7.87%	0.200	0.49076	0.38680
ACE	Poseidon-L	0.05	4	0.21015	0.21015	0.19174	+8.76%	0.200	N/A	N/A
ACE	Poseidon-L	0.10	4	0.34691	0.34691	0.32116	+7.42%	0.200	0.31583	0.26437
ACE	Poseidon-L	0.20	4	0.53233	0.53233	0.50212	+5.67%	0.200	0.47713	0.41710
ACE	Poseidon-L	0.30	3	0.59157	0.59155	0.55115	+6.83%	0.200	0.51387	0.43311
ACE	Poseidon-L	0.40	2	0.57904	0.57904	0.53685	+7.28%	0.200	0.46927	0.37534
ACE	Poseidon-T	0.05	4	0.31724	0.31724	0.28610	+9.82%	0.200	N/A	N/A
ACE	Poseidon-T	0.10	4	0.35590	0.35590	0.32842	+7.72%	0.200	0.32433	0.27066
ACE	Poseidon-T	0.20	4	0.54401	0.54401	0.51167	+5.94%	0.200	0.48851	0.42523
ACE	Poseidon-T	0.30	3	0.60487	0.60485	0.56200	+7.09%	0.200	0.52701	0.44117
ACE	Poseidon-T	0.40	2	0.58817	0.58817	0.54504	+7.33%	0.200	0.47811	0.38108
FNS-KF	CNO-FM	0.05	4	0.15573	0.15410	0.13710	+11.96%	0.200	N/A	N/A
FNS-KF	CNO-FM	0.10	4	0.12601	0.12091	0.11154	+11.48%	0.200	0.12518	0.09285
FNS-KF	CNO-FM	0.20	4	0.11485	0.11035	0.10960	+4.57%	0.050	0.11819	0.09385
FNS-KF	CNO-FM	0.30	3	0.09904	0.09612	0.09612	+2.96%	0.000	0.10941	0.08279
FNS-KF	CNO-FM	0.40	2	0.08089	0.07901	0.07919	+2.11%	0.020	0.10440	0.07397
FNS-KF	Dpot-S	0.10	4	0.08378	0.08329	0.08303	+0.90%	0.100	0.07305	0.06886
FNS-KF	Dpot-Ti	0.10	4	0.10423	0.10385	0.09962	+4.43%	0.200	0.09620	0.08232
FNS-KF	Morph-S	0.10	4	0.10458	0.09469	0.09445	+9.69%	0.100	0.09267	0.07882
FNS-KF	Morph-Ti	0.10	4	0.09918	0.09342	0.09326	+5.97%	0.100	0.08671	0.07799
FNS-KF	Polymathic-FNO	0.05	4	0.23397	0.23266	0.22149	+5.33%	0.200	N/A	N/A
FNS-KF	Polymathic-FNO	0.10	4	0.23442	0.23311	0.22258	+5.05%	0.200	0.23425	0.19412
FNS-KF	Polymathic-FNO	0.20	4	0.23584	0.23454	0.22599	+4.18%	0.200	0.23528	0.19820
FNS-KF	Polymathic-FNO	0.30	3	0.23722	0.23595	0.22699	+4.31%	0.200	0.23630	0.19150
FNS-KF	Polymathic-FNO	0.40	2	0.23846	0.23724	0.22608	+5.19%	0.200	0.23732	0.17635
FNS-KF	Polymathic-TFNO	0.05	4	0.24836	0.24832	0.23776	+4.27%	0.200	N/A	N/A
FNS-KF	Polymathic-TFNO	0.10	4	0.24878	0.24874	0.23875	+4.03%	0.200	0.24862	0.20818
FNS-KF	Polymathic-TFNO	0.20	4	0.25013	0.25008	0.24185	+3.31%	0.200	0.24961	0.21193
FNS-KF	Polymathic-TFNO	0.30	3	0.25120	0.25114	0.24244	+3.49%	0.200	0.25035	0.20409
FNS-KF	Polymathic-TFNO	0.40	2	0.25192	0.25182	0.24117	+4.26%	0.200	0.25085	0.18728
FNS-KF	Polymathic-UNetConvNext	0.05	4	0.07801	0.07074	0.06830	+12.45%	0.200	N/A	N/A
FNS-KF	Polymathic-UNetConvNext	0.10	4	0.11212	0.10725	0.10627	+5.21%	0.100	0.09773	0.08929
FNS-KF	Polymathic-UNetConvNext	0.20	4	0.15588	0.15243	0.15208	+2.44%	0.050	0.13623	0.13134
FNS-KF	Polymathic-UNetConvNext	0.30	3	0.16926	0.16686	0.16669	+1.52%	0.050	0.14329	0.14005
FNS-KF	Polymathic-UNetConvNext	0.40	2	0.17164	0.17001	0.16978	+1.08%	0.050	0.13742	0.13514
FNS-KF	Poseidon-B	0.05	4	0.02675	0.02627	0.02488	+7.01%	0.100	N/A	N/A
FNS-KF	Poseidon-B	0.10	4	0.01434	0.01431	0.01431	+0.23%	0.000	0.02895	0.01412
FNS-KF	Poseidon-B	0.20	4	0.03249	0.03248	0.03248	+0.03%	0.000	0.05457	0.03334
FNS-KF	Poseidon-B	0.30	3	0.04073	0.04072	0.04072	+0.01%	0.000	0.07299	0.04632
FNS-KF	Poseidon-B	0.40	2	0.04017	0.04016	0.04016	+0.01%	0.000	0.08833	0.05625
FNS-KF	Poseidon-L	0.05	4	0.02626	0.02510	0.02282	+13.11%	0.100	N/A	N/A
FNS-KF	Poseidon-L	0.10	4	0.01501	0.01495	0.01495	+0.38%	0.000	0.02895	0.01477
FNS-KF	Poseidon-L	0.20	4	0.03263	0.03262	0.03262	+0.03%	0.000	0.05432	0.03356
FNS-KF	Poseidon-L	0.30	3	0.04087	0.04086	0.04086	+0.02%	0.000	0.07261	0.04645
FNS-KF	Poseidon-L	0.40	2	0.04031	0.04030	0.04030	+0.02%	0.000	0.08775	0.05625
FNS-KF	Poseidon-T	0.05	4	0.15862	0.14629	0.13089	+17.48%	0.200	N/A	N/A
FNS-KF	Poseidon-T	0.10	4	0.04201	0.04031	0.04031	+4.03%	0.000	0.05040	0.03338
FNS-KF	Poseidon-T	0.20	4	0.05264	0.05187	0.05187	+1.46%	0.000	0.06949	0.04725
FNS-KF	Poseidon-T	0.30	3	0.05869	0.05775	0.05775	+1.60%	0.000	0.08495	0.05684
FNS-KF	Poseidon-T	0.40	2	0.05339	0.05262	0.05262	+1.43%	0.000	0.09454	0.06171
NS-Gauss	CNO-FM	0.05	4	0.07801	0.07592	0.06277	+19.53%	0.200	N/A	N/A
NS-Gauss	CNO-FM	0.10	4	0.07639	0.07535	0.06775	+11.31%	0.200	0.10045	0.06567
NS-Gauss	CNO-FM	0.20	4	0.10296	0.10211	0.09983	+3.04%	0.050	0.14446	0.11401
NS-Gauss	CNO-FM	0.30	3	0.10257	0.10171	0.10100	+1.53%	0.020	0.19111	0.14422
NS-Gauss	CNO-FM	0.40	2	0.09503	0.09383	0.09361	+1.49%	0.010	0.24769	0.18596
NS-Gauss	Dpot-S	0.10	4	0.25495	0.25367	0.22941	+10.02%	0.200	0.21827	0.19637
NS-Gauss	Dpot-Ti	0.10	4	0.26637	0.26538	0.23701	+11.02%	0.200	0.23387	0.20272
NS-Gauss	Morph-S	0.10	4	0.30463	0.29196	0.26066	+14.43%	0.200	0.26477	0.22594
NS-Gauss	Morph-Ti	0.10	4	0.29402	0.28878	0.25836	+12.13%	0.200	0.25480	0.22319
NS-Gauss	Polymathic-FNO	0.05	4	0.59810	0.59709	0.56268	+5.92%	0.200	N/A	N/A
NS-Gauss	Polymathic-FNO	0.10	4	0.59787	0.59686	0.56386	+5.69%	0.200	0.59799	0.49135
NS-Gauss	Polymathic-FNO	0.20	4	0.59731	0.59630	0.56654	+5.15%	0.200	0.59768	0.49967
NS-Gauss	Polymathic-FNO	0.30	3	0.59820	0.59704	0.56214	+6.03%	0.200	0.59892	0.48475

Continued on next page

Table 17 continued

PDE	Model	Δt	Steps	Raw	Projected	KMF	Gain	α^*	Dense raw-lin.	Dense KMF
NS-Gauss	Polymathic-FNO	0.40	2	0.60057	0.59912	0.55624	+7.38%	0.200	0.60147	0.46416
NS-Gauss	Polymathic-TFNO	0.05	4	0.60242	0.60189	0.56676	+5.92%	0.200	N/A	N/A
NS-Gauss	Polymathic-TFNO	0.10	4	0.60222	0.60168	0.56816	+5.66%	0.200	0.60231	0.49487
NS-Gauss	Polymathic-TFNO	0.20	4	0.60184	0.60130	0.57122	+5.09%	0.200	0.60199	0.50346
NS-Gauss	Polymathic-TFNO	0.30	3	0.60239	0.60167	0.56632	+5.99%	0.200	0.60263	0.48724
NS-Gauss	Polymathic-TFNO	0.40	2	0.60359	0.60252	0.55867	+7.44%	0.200	0.60416	0.46369
NS-Gauss	Polymathic-UNetConvNext	0.05	4	0.27585	0.25078	0.22267	+19.28%	0.200	N/A	N/A
NS-Gauss	Polymathic-UNetConvNext	0.10	4	0.37140	0.35310	0.32816	+11.64%	0.200	0.32819	0.30178
NS-Gauss	Polymathic-UNetConvNext	0.20	4	0.48107	0.46701	0.46215	+3.93%	0.050	0.42883	0.45306
NS-Gauss	Polymathic-UNetConvNext	0.30	3	0.51107	0.50255	0.50307	+1.57%	0.050	0.44528	0.52569
NS-Gauss	Polymathic-UNetConvNext	0.40	2	0.51551	0.51057	0.50975	+1.12%	0.020	0.43060	0.50678
NS-Gauss	Poseidon-B	0.05	4	0.05919	0.05890	0.04941	+16.52%	0.200	N/A	N/A
NS-Gauss	Poseidon-B	0.10	4	0.00420	0.00394	0.00394	+6.21%	0.000	0.11326	0.02571
NS-Gauss	Poseidon-B	0.20	4	0.00656	0.00634	0.00634	+3.35%	0.000	0.18216	0.07734
NS-Gauss	Poseidon-B	0.30	3	0.00882	0.00862	0.00862	+2.27%	0.000	0.23653	0.13124
NS-Gauss	Poseidon-B	0.40	2	0.01057	0.01034	0.01034	+2.21%	0.000	0.28787	0.19009
NS-Gauss	Poseidon-L	0.05	4	0.08979	0.08970	0.07362	+18.00%	0.200	N/A	N/A
NS-Gauss	Poseidon-L	0.10	4	0.00335	0.00314	0.00314	+6.27%	0.000	0.11325	0.02572
NS-Gauss	Poseidon-L	0.20	4	0.00541	0.00523	0.00523	+3.34%	0.000	0.18219	0.07737
NS-Gauss	Poseidon-L	0.30	3	0.00789	0.00766	0.00766	+2.88%	0.000	0.23640	0.13140
NS-Gauss	Poseidon-L	0.40	2	0.00996	0.00968	0.00968	+2.76%	0.000	0.28788	0.19017
NS-Gauss	Poseidon-T	0.05	4	0.14967	0.10534	0.08853	+40.85%	0.200	N/A	N/A
NS-Gauss	Poseidon-T	0.10	4	0.00888	0.00849	0.00849	+4.47%	0.000	0.11272	0.02482
NS-Gauss	Poseidon-T	0.20	4	0.01348	0.01309	0.01309	+2.92%	0.000	0.18206	0.07645
NS-Gauss	Poseidon-T	0.30	3	0.01765	0.01725	0.01725	+2.31%	0.000	0.23628	0.13011
NS-Gauss	Poseidon-T	0.40	2	0.02153	0.02107	0.02107	+2.12%	0.000	0.28768	0.18831
Wave-Gauss	Polymathic-FNO	0.05	4	0.12734	0.12471	0.12471	+2.07%	0.000	N/A	N/A
Wave-Gauss	Polymathic-FNO	0.10	4	0.12276	0.12029	0.12029	+2.01%	0.000	0.12834	0.11777
Wave-Gauss	Polymathic-FNO	0.30	2	0.12654	0.12373	0.12373	+2.22%	0.000	0.12243	0.33260
Wave-Gauss	Polymathic-TFNO	0.05	4	0.21844	0.21835	0.21835	+0.04%	0.000	N/A	N/A
Wave-Gauss	Polymathic-TFNO	0.10	4	0.20260	0.20254	0.20254	+0.03%	0.000	0.21050	0.18386
Wave-Gauss	Polymathic-TFNO	0.30	2	0.18135	0.18073	0.18073	+0.34%	0.000	0.19939	0.18953
Wave-Gauss	Polymathic-UNetConvNext	0.05	4	0.16661	0.13591	0.13591	+18.43%	0.000	N/A	N/A
Wave-Gauss	Polymathic-UNetConvNext	0.10	4	0.20624	0.16631	0.16631	+19.36%	0.000	0.18284	0.16067
Wave-Gauss	Polymathic-UNetConvNext	0.30	2	0.21317	0.17931	0.17931	+15.88%	0.000	0.20852	0.66474
Wave-Gauss	Poseidon-B	0.05	4	0.11130	0.10958	0.10958	+1.55%	0.000	N/A	N/A
Wave-Gauss	Poseidon-B	0.10	4	0.13151	0.11462	0.11462	+12.84%	0.000	0.12767	0.10366
Wave-Gauss	Poseidon-B	0.30	2	0.15148	0.13052	0.13052	+13.84%	0.000	0.15322	0.15534
Wave-Gauss	Poseidon-L	0.05	4	0.11213	0.10222	0.10222	+8.84%	0.000	N/A	N/A
Wave-Gauss	Poseidon-L	0.10	4	0.14016	0.11988	0.11988	+14.47%	0.000	0.12877	0.10834
Wave-Gauss	Poseidon-L	0.30	2	0.13842	0.12035	0.12035	+13.06%	0.000	0.14021	0.14097
Wave-Gauss	Poseidon-T	0.05	4	0.10471	0.10649	0.10649	-1.70%	0.000	N/A	N/A
Wave-Gauss	Poseidon-T	0.10	4	0.13948	0.11857	0.11857	+14.99%	0.000	0.12967	0.10904
Wave-Gauss	Poseidon-T	0.30	2	0.14185	0.12054	0.12054	+15.02%	0.000	0.14368	0.16070