NOVA: NULLSPACE DYNAMICS IN NEURAL BASIS SOLVERS FOR MOVING BOUNDARY PDES, TRANSPORT, RANK TRANSITIONS, AND REPROJECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Multiphysics coupling with moving boundaries poses a major challenge to neural solvers, as evolving interfaces induce time dependent constraints. Existing boundary treatments face an inherent trade-off between constraint accuracy and update cost. To address this challenge, we propose NOVA (**N**ullspace Dynamics in Neural Basis S**O**lvers for Mo**V**ing Bound**A**ry PDEs, Transport, Rank Transitions, and Reprojection). The method explicitly decomposes neural basis coefficients into a boundary lift and a component parameterized by instantaneous nullspace coordinates, then reprojects the previous physical state onto the updated nullspace. This construction avoids boundary parameter sweeps by eliminating penalty weights and stabilization parameters. To avoid costly singular value decompositions (SVDs) at every time step, we derive horizontal transport equations for the nullspace basis on the Stiefel manifold and develop an accompanying algorithm. Furthermore, spectral analysis establishes a rank saturation criterion for boundary collocation. Across eight multiphysics benchmarks covering phase change and flow, multiphase flow, and free boundary problems, NOVA achieves best or competitive field accuracy with substantially lower boundary errors than baselines, providing a new paradigm for physics-informed neural networks based on features and mesh-free representations to efficiently handle complex moving geometries. Our code is available here: `https://anonymous.4open.science/r/NOVA-3C44`.

1 INTRODUCTION

Multiphysics problems with moving boundaries arise widely in complex natural and engineering processes. Their simulation requires accurate coupling of interior partial differential equation (PDE) solutions, interface evolution, and conservative flux exchange. Body-fitted arbitrary Lagrangian-Eulerian (ALE) methods (Donea et al., 2004) require frequent reconstruction under severe deformation, whereas fixed-background-grid approaches, such as the immersed boundary method (IBM) (Peskin, 2002), face substantial integration errors and conservation challenges. Efficiently reconciling geometric compatibility with physical coupling at evolving interfaces remains a bottleneck in computational mechanics (Griffith & Patankar, 2020, Scheufler & Roenby, 2023).

Physics-informed neural networks (PINNs) have enabled mesh-free solutions of PDEs (Raissi et al., 2017). However, gradient pathologies can hinder optimization and cause training failure (Wang et al., 2021, Krishnapriyan et al., 2021). The Fourier neural operator (FNO) (Li et al., 2020) relies on costly precomputed solutions and faces generalization challenges in unfamiliar dynamics. Frozen-PINN (Datar et al., 2026) freezes spatial features and reduces PDEs to ordinary differential equations (ODEs) in coefficient space, preserving temporal causality without training. Nevertheless, extending these methods beyond simple, static geometries is challenging because moving boundaries and interfaces induce time-dependent constraints in coefficient space.

Under evolving boundary constraints, existing approaches trade off accuracy against update cost. Soft constraints cannot generally enforce interface conditions exactly (Berrone et al., 2023, Chiba & Saito, 2023), while larger penalty weights worsen conditioning (Juntunen & Stenberg, 2009). Baumgarte stabilization (Baumgarte, 1972) uses negative feedback. Weak gains permit leakage, whereas large gains introduce stiffness and restrict time steps. Nitsche's method (Nitsche, 1971)

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.

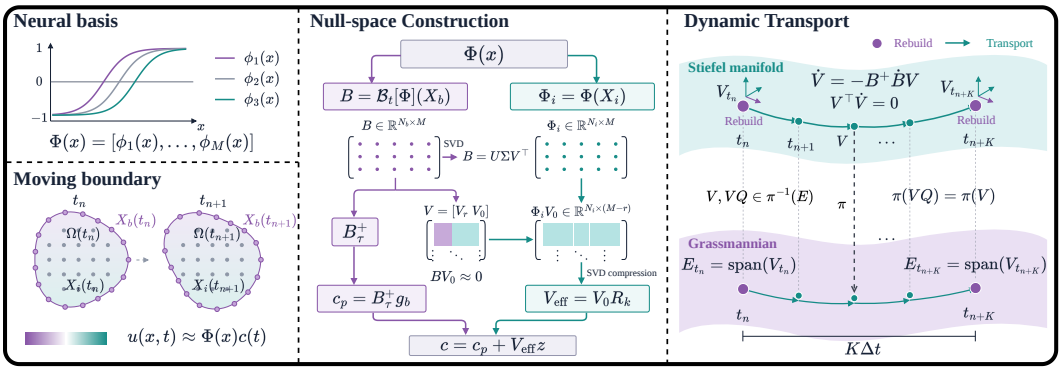


Figure 1: Overview of NOVA. **Left.** Frozen neural features represent the physical field on a moving domain. **Middle.** The current boundary matrix defines a boundary lift and numerical nullspace, with optional interior compression, yielding $c = c_p + V_{\text{eff}} z$. **Right.** Horizontal transport evolves the nullspace frame between full rebuilds on the Stiefel manifold, while its projection onto the Grassmann manifold represents the evolving admissible subspace.

requires costly boundary quadrature and stabilization updates on moving geometries. These treatments depend on hyperparameter tuning, including penalty weights, Baumgarte feedback gains, and Nitsche stabilization parameters, potentially requiring extensive parameter sweeps and repeated solves. Meanwhile, geometric evolution shifts the constraint spectrum, and accumulated interface errors can compromise interface advancement.

Nullspace methods such as LSE-ELM (De Falco et al., 2025) and NP-PIELM (Mishra et al., 2026) separate coefficients into a boundary lift and a nullspace component to enforce discrete boundary conditions without penalties. Moving boundaries, however, require maintaining an evolving subspace. Recomputing its SVD at every step is costly, and reusing old coordinates after basis updates can introduce physical-state jumps.

Temporal causality poses a further challenge. Global space-time methods jointly solve an entire time window, requiring future inputs that are unavailable in online multiphysics interactions. They cannot directly accommodate real-time feedback from an external agent whose next action depends on the current state. Global assembly also imposes a substantial memory burden over long time horizons.

To address these challenges, we propose NOVA, a nullspace dynamics method for neural-basis solvers of moving-boundary PDEs. Figure 1 illustrates the overall framework. At each time step, NOVA computes only the current state and marches to the next, so it can respond to instantaneous interactions. Within each step, a boundary lift and a nullspace component enforce the current boundary constraints. This separates boundary enforcement from interior evolution, so boundary accuracy no longer has to be traded against interior accuracy through a tunable weight. Field-value-based least-squares reprojection then transfers the existing state to the updated representation with minimal transfer error, supporting continued evolution across basis refreshes and numerical-rank changes.

To reduce the cost of repeated full SVDs, we update the nullspace basis between occasional SVDs through a transport equation driven by the boundary motion. Because rotations within the nullspace leave the subspace unchanged, we select the horizontal transport given by the principal-bundle structure relating the Stiefel and Grassmann manifolds (Edelman et al., 1998), which excludes these rotations and moves the basis the least. Spectral analysis further yields a rank-saturation criterion to guide boundary collocation.

Our contributions are summarized as follows:

- **A mesh-free solution framework.** We introduce a causal, time-marching, mesh-free framework with separated spatial features for moving-boundary multiphysics. A boundary lift and nullspace decomposition restrict coefficients to the current boundary-compatible affine space, without mesh reconstruction or backpropagation-based training. Enforcing the constraints by construction avoids tuning penalty weights or other parameters.

- **Nullspace transport and state reprojection.** Horizontal transport follows the Stiefel and Grassmann bundle geometry, reducing rebuild cost. Reprojection transfers physical states across basis refreshes and numerical-rank transitions.
- **Boundary spectral analysis and a collocation criterion.** We derive a rank-saturation criterion that indicates when boundary collocation points capture the space represented by the chosen neural basis, guiding their selection before the solve.
- **Extensive multiphysics evaluation.** Eight benchmarks spanning phase-change melting, incompressible multiphase flow, and freely moving boundaries demonstrate the method's advantages in eliminating interface leakage, improving long-time coupling stability, and reducing computational cost.

2 METHOD

2.1 AN INITIAL BOUNDARY VALUE PROBLEM ON A MOVING DOMAIN

We consider a physical field $u(x,t)$ on a moving domain $\Omega(t) \subset \mathbb{R}^d$, with boundary $\Gamma(t) = \partial\Omega(t)$. Over the time interval $[0,T]$, the initial boundary value problem is

$$\begin{aligned} \partial_t u(x,t) + \mathcal{L}_t[u](x,t) + \mathcal{N}_t[u](x,t) &= f(x,t), \qquad x \in \Omega(t), \quad t \in (0,T], \\ \mathcal{B}_t[u](x,t) &= g(x,t), \qquad x \in \Gamma(t), \quad t \in (0,T], \\ u(x,0) &= u_0(x), \qquad x \in \Omega(0). \end{aligned} \tag{1}$$

Here, $\mathcal{L}_t$ and $\mathcal{N}_t$ are the linear and nonlinear spatial operators, $\mathcal{B}_t$ is a linear boundary operator, and f, g, and u_0 are the source term, boundary data, and initial condition. Advancing this problem requires enforcing the boundary conditions on each updated geometry while retaining enough freedom to resolve the interior dynamics.

2.2 FROZEN NEURAL BASES AND A BOUNDARY COMPATIBLE REPRESENTATION

We represent the field as a combination of time-independent spatial basis functions, with time-dependent coefficients controlling their contributions. With the spatial basis fixed, we obtain the time-dependent coefficients by solving the discretized PDE, enabling training-free time marching (Datar et al., 2026). Given M fixed spatial basis functions $\phi_i(x) = \sigma(w_i^\top x + b_i)$, let $\Phi(x) = [\phi_1(x), \ldots, \phi_M(x)]$ and approximate the physical field by $u_h(x,t) = \Phi(x)c(t)$. Here, σ is the activation function, $w_i \in \mathbb{R}^d$ is a weight vector, and $b_i \in \mathbb{R}$ is a bias. The coefficient vector $c(t) \in \mathbb{R}^M$ evolves in time, while w_i, b_i, and σ remain fixed. Sampling the current boundary at N_b collocation points $X_b(t) \subset \Gamma(t)$ and applying $\mathcal{B}_t$ gives the matrix $B(t) = \mathcal{B}_t[\Phi](X_b(t)) \in \mathbb{R}^{n_b \times M}$ and the boundary constraint $B(t)c(t) = g_b(t)$. Here, $g_b(t)$ samples $g(x,t)$ at $X_b(t)$, and n_b counts scalar constraints. In coefficient space, the boundary equations restrict some directions and leave the others free for the interior PDE. We find these free directions from the full SVD $B = U\Sigma V^\top$ with $V \in \mathbb{R}^{M\times M}$. With B^+ the Moore-Penrose pseudoinverse, $\#$ the number of elements in a set, and a relative threshold τ, define

$$r = \#\{j : \sigma_j(B) > \tau\sigma_1(B)\}, \quad B_\tau = U_r\Sigma_r V_r^\top, \quad B_\tau^+ = V_r\Sigma_r^{-1}U_r^\top, \quad V_0 = V_{[:,r+1:M]}. \tag{2}$$

Algorithm 1 NOVA

Input (1), $\Phi, X_b, X_i, \{t_n\}, K$, geometry and tolerances
Output $u_h(x,t_n) = \Phi(x)c_n$

Initialize the frame, row factors and $c_{p,0}$, then fit c_0 by (5)
for $n = 0, \ldots, N-1$ **do**
  repeat
    Assemble the candidate geometry, $B, g_b, \Phi_i, R_i, \dot{B}$
    if scheduled refresh **then**
      Rebuild $V_{\mathrm{eff},n+1}$ and row factors, select r, k and set $P = B_\tau^+$
    else ▷ Appendix B.5
      Update row factors with fixed r and obtain P
      $\widetilde{V} \leftarrow V_{\mathrm{eff},n} - \Delta t P\dot{B}V_{\mathrm{eff},n}$
      for $j = 1, 2$ **do**
        $\widetilde{V} \leftarrow \widetilde{V} - PB\widetilde{V}$
      end for
      $V_{\mathrm{eff},n+1} \leftarrow \mathrm{NS}_1(\widetilde{V})$ ▷ one Newton-Schulz step
    end if
    $c_{p,n+1} \leftarrow Pg_b$ ▷ $P = B_\tau^+$ on rebuilds
    Solve z_{n+1} by (6) ▷ startup if history is unavailable
    $c_{n+1} \leftarrow c_{p,n+1} + V_{\mathrm{eff},n+1}z_{n+1}$
  until Γ prescribed or coupling converged
  Store c_{n+1}, Γ_{n+1} and advance history
end for

We then decompose c into a boundary lift c_p and a nullspace component $V_0 z$,

$$c = c_p + V_0 z, \qquad c_p = B_\tau^+ g_b. \tag{3}$$

Here, $Bc_p \approx g_b$, V_0 spans the numerical nullspace of B with $BV_0 \approx 0$, and $z \in \mathbb{R}^{M-r}$ contains the unknown reduced coefficients. Substituting this representation into the interior discretization leaves only z to solve for. Every z satisfies the sampled boundary conditions, so the solver needs no penalty weight to balance boundary and interior residuals.

2.3 Interior compression and coefficient dynamics

Let $X_i(t)$ contain N_i interior collocation points, with $\Phi_i(t) = \Phi(X_i(t)) \in \mathbb{R}^{N_i \times M}$.

Although V_0 satisfies the boundary constraints, some directions have little effect on the interior field or contribute nearly redundant information. Retaining them increases the number of unknowns while allowing large coefficient changes with little effect on the sampled field. We remove these directions to obtain a smaller, better-conditioned interior representation. Specifically, we compute $\Phi_i V_0 = \widehat{U}\widehat{\Sigma}R^\top$ and select the truncated rank $k = \#\{j : \widehat{\sigma}_j > \tau_M \widehat{\sigma}_1\}$. Here, $\widehat{\sigma}_j$ are the singular values of $\Phi_i V_0$, and τ_M is the relative compression threshold. The leading k right singular vectors form $R_k \in \mathbb{R}^{(M-r)\times k}$, giving $V_{\text{eff}} = V_0 R_k \in \mathbb{R}^{M\times k}$. While V_0 defines the boundary-compatible space, V_{eff} selects k working directions within it for the interior solve.

Because the boundary moves, the lift c_p and the frame V_0 also change in time. In fixed laboratory coordinates, write the interior right-hand side as $R_i(c,t) = f_i(t) - \mathcal{L}_{i,t}[\Phi c] - \mathcal{N}_{i,t}[\Phi c]$, where i denotes evaluation at $X_i(t)$. Differentiating the uncompressed representation $c = c_p + V_0 z$ exposes this motion as $\dot{c} = \dot{c}_p + \dot{V}_0 z + V_0 \dot{z}$. The interior collocation relation is

$$\Phi_i V_0 \dot{z} \approx R_i(c,t) - \Phi_i \dot{c}_p - \Phi_i \dot{V}_0 z. \tag{4}$$

The $\dot{c}_p$ and $\dot{V}_0$ terms account for changes in the representation itself. For this reason, Section 2.4 discretizes the full coefficients c instead of integrating this relation for z.

2.4 State reprojection and time integration

As the boundary moves, we update B and its nullspace basis. Since the coordinates z depend on the basis, reusing them after a basis update can introduce an artificial jump in the physical field. Reprojection finds coordinates z_{new} in the new basis that preserve the old field u_{old} as closely as possible. Here, $u_{\text{old}} = \Phi c$ is the field in the previous representation. We compare the old and new field values at sample points X_{tr}, with weights specified by the diagonal matrix W_{tr}, and solve

$$z_{\text{new}} = \arg\min_z \left\| W_{\text{tr}}^{1/2} \left[\Phi(X_{\text{tr}})(c_{p,\text{new}} + V_{\text{eff,new}} z) - u_{\text{old}}(X_{\text{tr}}) \right] \right\|_2^2. \tag{5}$$

We store previous states as full coefficient vectors in the fixed spatial basis. This keeps the history in the same coordinates when the reduced basis or its dimension changes. We apply the second-order backward differentiation formula (BDF2) (Wanner & Hairer, 1996) to this history, then solve for the new coordinates. On the time grid t_n, a subscript n denotes evaluation at t_n, and $\Phi_{i,n+1} = \Phi(X_i(t_{n+1}))$. Let $R_{i,n+1}$ denote the interior right-hand side at the current time and geometry. We write the new coefficient vector as $c_{n+1}(z) = c_{p,n+1} + V_{\text{eff},n+1} z$, where $z \in \mathbb{R}^k$ is unknown. Let W_i be the diagonal matrix of interior residual weights. For constant step Δt,

$$z_{n+1} = \arg\min_z \left\| W_i^{1/2} \left\{ \Phi_{i,n+1} \left[3c_{n+1}(z) - 4c_n + c_{n-1} \right] - 2\Delta t \, R_{i,n+1}\big(c_{n+1}(z)\big) \right\} \right\|_2^2. \tag{6}$$

Because the history c_n, c_{n-1} and the new lift and frame are known, z is the only unknown, and each step reduces to one least-squares problem, which is linearized for nonlinear problems. We solve it by truncated SVD with tolerance τ_{LS}, then reconstruct $c_{n+1} = c_{p,n+1} + V_{\text{eff},n+1} z_{n+1}$. We use backward Euler (BE) for startup and restart. Appendix B.5 gives the derivation.

2.5 HORIZONTAL TRANSPORT OF THE NULLSPACE AND BOUNDARY LIFT

To reduce the number of full SVDs, we update the existing nullspace frame through a transport equation as the boundary moves and perform a full SVD only every K steps. On a smooth interval, differentiating $BV_0 = 0$ relates the frame motion to the boundary change $\dot{B}$. This relation does not determine the update uniquely, because V_0 and V_0Q span the same subspace for any orthogonal matrix Q. We therefore choose the horizontal update, which excludes these rotations. The theorem below justifies this choice. Among all frame velocities that keep V_0 in the kernel and orthonormal, the horizontal one is unique and moves the frame the least.

Theorem 2.1 (Horizontal transport of the nullspace and boundary lift)**.** *On a local transport interval $I = [t_0, t_0 + h]$ with $h > 0$, assume $B \in C^1(I, \mathbb{R}^{n_b \times M})$ and $\operatorname{rank} B(t) = r$ for all $t \in I$. Set $m_0 = M - r$ and define the Stiefel manifold $\mathrm{St}(M, m_0) = \{V \in \mathbb{R}^{M \times m_0} \mid V^\top V = I_{m_0}\}$. Let $V_0(t) \in \mathrm{St}(M, m_0)$ span $\ker B(t)$, with $V_0(t_0)$ prescribed. Among velocities X satisfying $BX = -\dot{B}V_0$ and $V_0^\top X + X^\top V_0 = 0$, the horizontal condition $V_0^\top X = 0$ uniquely gives*

$$\dot{V}_0 = -B^+ \dot{B} V_0, \qquad V_0^\top \dot{V}_0 = 0. \tag{7}$$

It also uniquely minimizes the instantaneous Frobenius norm among these velocities.

If $g_b \in C^1(I, \mathbb{R}^{n_b})$, $g_b(t) \in \operatorname{range} B(t)$ on I, and $Bc_p = g_b$, then among velocities y satisfying $By = \dot{g}_b - \dot{B}c_p$, the condition $V_0^\top y = 0$ uniquely gives

$$\dot{c}_p = B^+(\dot{g}_b - \dot{B}c_p), \qquad V_0^\top \dot{c}_p = 0. \tag{8}$$

Proof sketch. Differentiating $BV_0 = 0$ gives $BX = -\dot{B}V_0$. By Lemma B.3, the right-hand side belongs to $\operatorname{range} B$, so $-B^+\dot{B}V_0$ is a particular solution. Proposition B.4 gives the general solution $X = -B^+\dot{B}V_0 + V_0 S$. Combining $V_0^\top B^+ = 0$ with $V_0^\top X + X^\top V_0 = 0$ yields $S + S^\top = 0$.

The term $V_0 S$ rotates the basis within the same subspace, and under the ambient Frobenius metric the horizontal condition $V_0^\top X = 0$ removes it by setting $S = 0$ (Proposition B.2). Since $\|X\|_F^2 = \|B^+\dot{B}V_0\|_F^2 + \|S\|_F^2$, this horizontal velocity also has the minimum Frobenius norm (Corollary B.5).

The lift follows in the same way. Every lift velocity has the form $B^+(\dot{g}_b - \dot{B}c_p) + V_0 a$ with $a = V_0^\top \dot{c}_p$, so horizontality sets $a = 0$ (Proposition B.6). This keeps the lift out of the nullspace, where z already carries the interior freedom. Full proofs are given in Appendices B.2 and B.3. □

For fixed R_k, the transport law also applies to the compressed frame $V_{\text{eff}} = V_0 R_k$, giving $\dot{V}_{\text{eff}} = -B^+\dot{B}V_{\text{eff}}$. Algorithm 1 approximates this law using a first-order transport step followed by boundary correction and near-orthogonalization. Appendix B.5 details the discrete update, including the single near-orthogonal Newton-Schulz step NS_1 in equation 21.

Between full SVD refreshes, r and k remain fixed. At each step, we evaluate B and g_b at t_{n+1} on the candidate geometry, which is the prescribed geometry for moving boundaries and the current coupling iterate for free boundaries. The solver then recomputes the lift $c_p = Pg_b$. Here, $P = B_\tau^+$ at a refresh, and between refreshes P is updated from the current B with the rank r held fixed. Appendix B.3 explains how this choice differs from integrating the horizontal lift.

2.6 RANK SATURATION

Insufficient boundary sampling can miss constraints, while redundant points add decomposition cost. Rank saturation relates the required sampling to the geometry, operator, and basis, guiding point selection before the PDE solve. At fixed t, the trace operator $\mathcal{T}_t c = \mathcal{B}_t[\Phi c]|_{\Gamma(t)}$ gives a function on the entire boundary before sampling. Its range $\mathcal{Y}_t = \operatorname{range} \mathcal{T}_t$ has dimension $r_{\text{tr}} \le M$, the number of independent trace directions. The sampling operator S_{X_b} evaluates boundary functions at X_b, giving $B = S_{X_b}\mathcal{T}_t$ with n_b scalar constraints.

Theorem 2.2 (Rank saturation of the boundary trace space)**.** *The sampling map S_{X_b} is injective on $\mathcal{Y}_t$ if and only if $\operatorname{rank} B = r_{\text{tr}}$. Equivalently,*

$$\operatorname{rank} B = r_{\text{tr}} \quad \Longleftrightarrow \quad \ker B = \ker \mathcal{T}_t. \tag{9}$$

These conditions require $n_b \ge r_{\text{tr}}$. When they hold, for any $g \in \mathcal{Y}_t$ and $g_b = S_{X_b} g$, satisfying $Bc = g_b$ implies $\mathcal{T}_t c = g$ on the entire boundary, including unsampled points.

If $\operatorname{rank} B < r_{\mathrm{tr}}$, then $\ker \mathcal{T}_t \subsetneq \ker B$. Hence there is a coefficient perturbation δc with $B\delta c = 0$ and $\mathcal{T}_t \delta c \not\equiv 0$. Adding it to coefficients satisfying the full boundary condition preserves the sampled constraints but violates the condition at some unsampled boundary points.

Once $\operatorname{rank} B = r_{\mathrm{tr}}$, any enlarged set $\widetilde{X}_b \supseteq X_b$, with $\widetilde{B} = S_{\widetilde{X}_b} \mathcal{T}_t$, satisfies $r_{\mathrm{tr}} = \operatorname{rank} B \leq \operatorname{rank} \widetilde{B} \leq r_{\mathrm{tr}}$. Further points cannot increase the exact rank because they add no independent boundary constraints. Appendix B.4 gives the proof. Section 3.6 tests the numerical counterparts of rank saturation in Figure 3 and boundary compatibility at unsampled points in Table 3.

3 Experiments

We evaluate NOVA on eight physical benchmarks covering time-dependent boundary values, prescribed moving boundaries, changes in domain topology, and physically coupled free boundaries. For mesh-free methods that require no gradient-based training, we compare three variants of Frozen-PINN with different boundary treatments, namely Baumgarte, Nitsche, and the penalty method. Baumgarte corresponds exactly to the augmented ODE of Frozen-PINN. For neural methods that require training, we optimize PINNs using Adam (Kingma & Ba, 2014) and limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) (Liu & Nocedal, 1989) separately. Table 1 reports E_{field}, E_B, runtime, and peak resident set size (RSS). Appendix C provides the metric and problem definitions, method configurations, and evaluation protocols.

Across all eight cases, we scan 23 penalty weights λ from 10 to 10^{12} and all combinations of 16 Baumgarte damping values κ from 10 to 10^8 with 15 boundary row weights w from 10^{-3} to 10^5. Nitsche uses 15 safety factors s_N from 1.01 to 10^5, jointly scanned with the same row weights for Moving cylinder and Navier-Stokes. Appendix C.1 lists the exact grids, and Tables 5 to 12 report representative scanned configurations for each case, showing how strongly the baseline errors depend on these parameters. NOVA has no counterpart to these boundary parameters, and its remaining settings are fixed across all eight cases.

The main comparison in Table 1 uses a single NOVA configuration ($K = 20$, $\tau_M = 10^{-13}$) for all eight cases, whereas each Frozen baseline is reported at the configuration with the lowest E_{field} over its completed scan candidates, a selection that favors the baselines. The refresh interval K only sets how often the nullspace is fully rebuilt, and E_{field} changes little across $K \in \{1, 5, 10, 20\}$ in all eight benchmarks (Tables 2 and 13). We mark the best and second best eligible errors in **bold** and cyan, respectively. NOVA rows use Alice Blue shading. Under this protocol, NOVA structurally reduces E_B to the lowest value in all eight cases while keeping E_{field} low or competitive, with the lowest E_{field} in four.

3.1 Fixed geometries with time-dependent boundary values

Burgers. We solve the one-dimensional viscous Burgers equation on a fixed interval with time-dependent endpoint values, testing whether the boundary conditions remain satisfied as the solution evolves. NOVA attains the lowest E_B, roughly 4×10^5 times lower than the penalty method, while keeping E_{field} competitive with Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.3.

3.2 Moving boundaries with prescribed trajectories

Dynamic disk. We solve the two-dimensional heat equation on a disk that moves and changes size along a prescribed trajectory. NOVA attains the lowest E_B, roughly 708 times lower than the penalty method, and the lowest E_{field}, roughly 4 times lower than Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.4.

Flow around a moving cylinder. We solve the incompressible Navier-Stokes equations around a cylinder that moves along a prescribed path through a periodic domain. All methods are evaluated against the same manufactured analytic solution for velocity and drag, with boundary error measured against the prescribed no-slip condition. NOVA attains the lowest E_B, roughly 9×10^5 times lower than Baumgarte, and a slightly lower E_{field} than Baumgarte (Table 1). Nitsche has no completed scan candidate. Setup and further sweeps appear in Appendix C.5.

Table 1: **Main comparison across eight benchmarks.** NOVA uses $K = 20$ in all eight cases without a sweep. Each Frozen baseline is reported at its lowest E_{field} over completed candidates scanned on $\lambda \in [10, 10^{12}]$ (penalty), $\kappa \in [10, 10^8]$ and $w \in [10^{-3}, 10^5]$ (Baumgarte), and $s_N \in [1.01, 10^5]$ (Nitsche). Selected settings and exact grids appear in Table 4 and Appendix C.1, with representative sweeps in Tables 5 to 12. ↓ indicates lower is better. Black **bold** and cyan mark the best and second-best eligible errors before rounding. Alice Blue marks NOVA. Costs come from different batches and are not ranked. + marks GPU training times. † denotes PINN results. Modal Scriven PINN values use different initial data and reference fields and are shown as separate references.

Burgers	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)	**Dynamic disk**	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)
PINN (Adam)†	1.66e-03	3.44e-04	230.15+	1218.23	PINN (Adam)†	3.50e-02	2.57e-02	445.38+	1209.45
PINN (L-BFGS)†	2.02e-04	6.98e-05	182.50+	1131.77	PINN (L-BFGS)†	6.51e-03	8.19e-03	754.33+	1135.09
Penalty	2.25e-05	7.01e-11	3.13	106.93	Penalty	3.63e-02	1.48e-07	33.21	318.25
Nitsche	2.25e-05	3.85e-10	3.34	109.07	Nitsche	3.72e-02	3.67e-06	43.72	295.61
Baumgarte	**2.82e-06**	1.36e-06	2.35	101.62	Baumgarte	3.52e-06	1.19e-06	19.44	201.43
NOVA (Ours)	3.02e-06	**1.94e-16**	1.50	101.15	NOVA (Ours)	**8.26e-07**	**2.09e-10**	11.78	225.79
Moving cylinder	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)	**Navier-Stokes**	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)
PINN (Adam)†	9.73e-02	8.52e-03	2486.72+	1229.39	PINN (Adam)†	9.45e-03	2.91e-02	1805.44+	1217.88
PINN (L-BFGS)†	8.24e-02	4.09e-03	3417.45+	1155.04	PINN (L-BFGS)†	4.21e-03	1.15e-02	2749.81+	1143.50
Penalty	3.84e-01	2.10e-08	32.10	288.26	Penalty	6.33e-03	1.21e-07	9.01	262.80
Nitsche	failed	failed	failed	failed	Nitsche	7.28e-03	8.84e-03	7.34	274.22
Baumgarte	8.19e-02	1.07e-09	46.44	263.96	Baumgarte	2.77e-05	4.13e-05	8.22	256.56
NOVA (Ours)	**7.93e-02**	**1.13e-15**	28.04	295.62	NOVA (Ours)	**7.94e-06**	**3.57e-08**	7.83	268.83
Topology change	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)	**Stefan**	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)
PINN (Adam)†	1.43e-02	1.28e-02	423.22+	1216.11	PINN (Adam)†	3.26e-03	3.38e-03	873.94+	1193.88
PINN (L-BFGS)†	1.73e-03	1.75e-03	746.67+	1142.45	PINN (L-BFGS)†	3.44e-04	3.44e-04	686.07+	1117.76
Penalty	1.98e-02	4.07e-06	3.47	178.61	Penalty	3.54e-02	2.01e-02	1.98	95.79
Nitsche	2.09e-02	4.52e-06	4.39	178.59	Nitsche	1.06e-02	1.39e-03	2.28	96.27
Baumgarte	3.36e-06	1.64e-06	2.68	180.11	Baumgarte	**5.83e-06**	5.75e-07	1.05	98.47
NOVA (Ours)	**2.36e-06**	**1.85e-08**	2.29	183.00	NOVA (Ours)	3.22e-05	**3.05e-16**	0.76	94.08
Scriven bubble	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)	**Scriven modal**	E_{field} ↓	E_B ↓	Time (s)	RSS (MiB)
PINN (Adam)†	5.11e-02	9.63e-02	573.51+	1205.39	PINN (Adam)†	1.62e-02	1.30e-02	1989.36+	1232.20
PINN (L-BFGS)†	5.41e-02	1.10e-01	858.06+	1170.28	PINN (L-BFGS)†	2.13e-03	8.37e-04	2672.26+	1149.45
Penalty	3.29e-05	2.71e-11	15.67	136.68	Penalty	9.60e-01	2.25e+00	74.91	560.38
Nitsche	3.29e-05	9.41e-11	16.64	129.98	Nitsche	5.94e-06	9.05e-13	61.72	223.43
Baumgarte	**4.77e-06**	1.87e-07	14.47	116.25	Baumgarte	**1.54e-07**	3.70e-07	45.66	204.21
NOVA (Ours)	6.37e-06	**6.33e-13**	7.31	116.14	NOVA (Ours)	2.54e-07	**3.51e-15**	35.26	210.64

Navier-Stokes. We solve the incompressible Navier-Stokes equations around a petal-shaped obstacle with prescribed translation and rotation inside a fixed outer boundary and prescribe analytic velocity values on the boundary. The first row of Figure 2 shows the geometry and flow. NOVA attains the lowest E_B, roughly 3 times lower than the penalty method, and the lowest E_{field}, roughly 3 times lower than Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.6.

3.3 Changes in domain topology

Topology change. We solve the heat equation on a Cassini domain that undergoes a prescribed, regularized splitting into two components. As in the Navier-Stokes case, this prescribed moving domain uses manufactured Dirichlet data to test nullspace transport. The second row of Figure 2 shows the split. NOVA attains the lowest E_B, roughly 89 times lower than Baumgarte, and a lower E_{field} than Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.7.

3.4 Physically coupled free boundaries

Stefan. We solve the two-phase Stefan problem, whose interface moves according to the heat fluxes on both sides (Wang & Perdikaris, 2021). Since this Stefan condition couples the interface to the temperature gradients, the temperature fields and interface must be solved together. NOVA attains the lowest E_B, roughly 2×10^9 times lower than Baumgarte, while keeping E_{field} competitive with Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.8.

Scriven bubble. We solve the spherically symmetric Scriven problem of bubble growth driven by heat transfer from the surrounding liquid. The liquid temperature evolves through advection and diffusion,

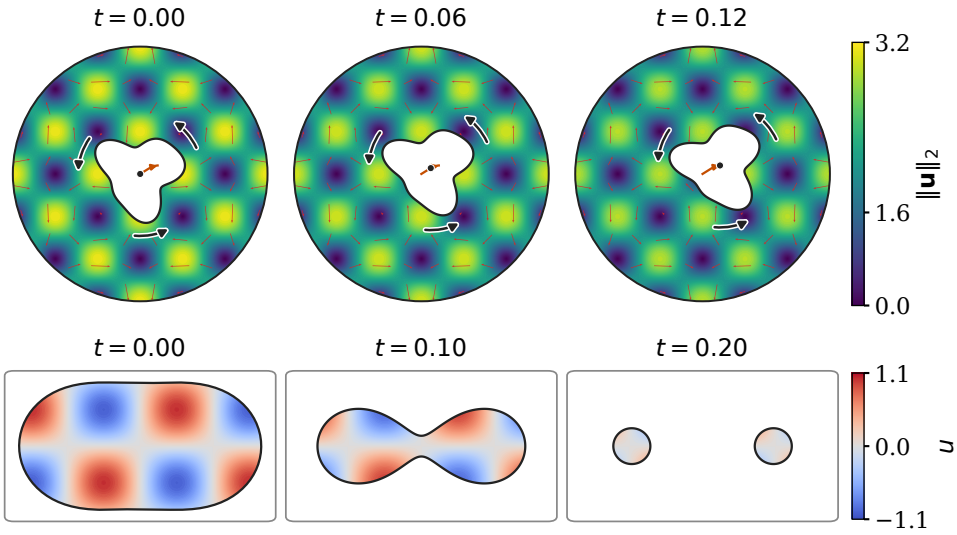


Figure 2: NOVA with $K = 20$ on two prescribed moving domains at three times. The first row (Navier-Stokes) shows a petal-shaped obstacle translating and rotating inside a fixed outer boundary. Colors show velocity magnitude and red arrows show velocity, both on a common scale across times. The black curved arrow indicates rotation, and the orange path and black dot mark the center trajectory and current center. The second row (Topology change) shows a Cassini domain splitting into two components, with the scalar field on a common color scale across times.

Table 2: **NOVA refresh interval ablation.** Full results appear in Appendix C.11. **Bold** and cyan mark best and second-best errors before rounding.

Dynamic disk	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	8.35e-07	**1.66e-10**	23.76	251.19
NOVA ($K = 5$)	8.31e-07	1.73e-10	14.14	246.34
NOVA ($K = 10$)	8.27e-07	1.89e-10	12.59	230.59
NOVA ($K = 20$)	**8.26e-07**	2.09e-10	11.78	225.79
Scriven bubble	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	**6.37e-06**	9.87e-13	16.47	119.30
NOVA ($K = 5$)	6.37e-06	**5.67e-13**	8.62	116.87
NOVA ($K = 10$)	6.37e-06	7.36e-13	7.83	118.38
NOVA ($K = 20$)	6.37e-06	6.33e-13	7.31	116.14

Table 3: **NOVA boundary-point ablation.** Navier-Stokes with $K = 1$ (Appendix C.12). The dashed line marks saturation at all saved times.

N_b	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
17	8.43e-02	5.80e-01	15.54	284.71
34	3.30e-02	1.54e-01	14.73	276.86
51	3.33e-04	1.49e-03	**11.78**	281.50
102	7.99e-06	4.24e-08	15.52	**274.07**
136	**7.92e-06**	**3.90e-08**	14.13	283.89
204	7.92e-06	3.91e-08	14.59	285.60
306	7.92e-06	3.93e-08	15.99	286.15
510	7.92e-06	3.91e-08	13.82	294.67
1020	7.92e-06	3.91e-08	14.14	300.17

while its gradient at the interface determines the bubble growth rate through the Stefan condition. NOVA attains the lowest E_B, roughly 43 times lower than the penalty method, while keeping E_{field} competitive with Baumgarte (Table 1). Setup and further sweeps appear in Appendix C.9.

Scriven modal. We extend the Scriven bubble to axisymmetric shape changes, jointly evolving modes through order two for the temperature, flow, and interface deformation from compatible initial data. The base temperature T_0 and mean radius $a_0(t) = R(t)$ are evaluated against the analytic Scriven reference, while higher-order modes remain in the coupled solve. NOVA attains the lowest E_B, roughly 258 times lower than Nitsche, while keeping E_{field} competitive with Baumgarte (Table 1). The penalty method completes with $E_{\text{field}} \approx 0.960$. Setup and further sweeps appear in Appendix C.10.

3.5 Ablation of the nullspace refresh interval

To assess the benefit of horizontal transport, we vary the full nullspace refresh interval over $K = 1, 5, 10, 20$, holding all other settings fixed within each benchmark. We use the same metrics as in Table 1. Table 2 presents the Dynamic disk and Scriven bubble results, with all eight benchmarks reported in Appendix C.11.

For both cases in Table 2, $K = 5, 10, 20$ all reduce runtime relative to $K = 1$, by roughly 40 to 50% for Dynamic disk and 48 to 56% for Scriven bubble. With $K = 1$, every step repeats the boundary SVD and the interior compression (Figure 1, Middle). For $K > 1$, intermediate steps transport the existing frame instead (Figure 1, Right, and Algorithm 1), so both SVDs are skipped between refreshes. Since the field solve is still required at every step, the runtime benefit depends on the share of work spent on these rebuilds.

These savings come with similar E_{field}, E_B, and peak RSS across all four settings in the two displayed cases, indicating that transport preserves accuracy without adding substantial memory overhead. Together, these results favor $K > 1$, with $K = 10$ or 20 as practical starting choices that save roughly half the runtime in both cases. A larger K would save more rebuild cost, but the frame would then rely on first-order transport for longer, which motivates the adaptive refresh discussed in Appendix D.

3.6 Rank saturation

We evaluate the sampling criterion in Section 2.6 through numerical rank saturation and errors at independent boundary points. On the Navier-Stokes case, we vary only N_b, with $K = 1$, $\tau = \tau_M = 10^{-12}$ and two velocity constraints per point, so $n_b = 2N_b$. At each of 33 saved times, we use the numerical rank from independent dense sampling as an approximation to r_{tr} at this threshold. Appendix C.12 gives the configuration.

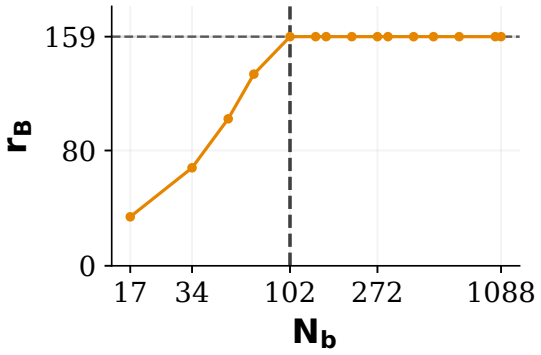


Figure 3: Navier-Stokes boundary rank $r_B = r$ at $t = 0.00$ and $\tau = 10^{-12}$. Dashed lines mark dense reference rank 159 and $N_b = 102$.

Figure 3 shows the numerical counterpart of the rank argument in Section 2.6. At $t = 0$, adding points increases r to the dense reference value 159 at $N_b = 102$, after which r remains constant. Across all saved times, the tested counts $N_b \geq 136$ match reference ranks of 159 to 160, giving the conservative sampling cutoff in Table 3.

Theorem 2.2 explains why complete sampling captures all represented boundary constraints, after which additional points add no independent constraints. Table 3 provides numerical support through errors at independent test points. Increasing N_b from 51 to 136 reduces E_B by 3.8×10^4 times and E_{field} by 42 times. From 136 to 1020 points, E_B stays near 3.9×10^{-8} and E_{field} near 7.92×10^{-6}. Thus, beyond the observed sampling cutoff, further points bring little accuracy gain, consistent with the theoretical saturation criterion. In practice, N_b can therefore be chosen before the PDE solve by increasing it until the numerical rank of B stops growing, which only requires SVDs of the boundary matrix and no field solves.

4 Conclusion

We propose NOVA, a neural basis method for multiphysics coupling with moving boundaries. The method separates boundary constraints from interior evolution through a boundary lift and a nullspace representation. Building the constraints into the representation removes the boundary parameters that the baselines must sweep. We derive the horizontal transport equation for the nullspace basis and develop a corresponding numerical algorithm to reduce the frequency of full SVDs, thereby reducing computational cost. To connect these evolving representations across time steps, we use physical state reprojection to transfer the previous state to the updated admissible space when geometry or rank changes. The rank saturation criterion provides a theoretical basis for selecting boundary collocation points before solving the PDE, supported by numerical validation. Across eight benchmarks, NOVA combines accurate fields with small boundary residuals, while the refresh-interval ablations demonstrate reduced computational cost. This framework extends neural basis methods to multiphysics coupling with moving boundaries and opens new directions for coupled simulation and control of complex dynamical systems. Limitations and future work are discussed in Appendix D.

AI USE STATEMENT

In this work, we used generative AI tools for translation, code implementation, and experimental design. We have not used generative AI tools for synthetic data generation or dataset cleaning, and qualitative analyses are not applicable to this work. Additionally, we used generative AI tools for literature searches, manuscript drafting and editing, LaTeX formatting, code debugging, and visualization. We have reviewed all AI-assisted work. We checked and revised the text, mathematics, code, and experimental analyses. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Joachim Baumgarte. Stabilization of constraints and integrals of motion in dynamical systems. *Computer methods in applied mechanics and engineering*, 1(1):1–16, 1972.

Stefano Berrone, Claudio Canuto, Moreno Pintore, and Natarajan Sukumar. Enforcing dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks. *Heliyon*, 9(8), 2023.

Yuki Chiba and Norikazu Saito. Nitsche's method for a robin boundary value problem in a smooth domain. *Numerical Methods for Partial Differential Equations*, 39(6):4126–4144, 2023.

Chinmay Datar, Taniya Kapoor, Abhishek Chandra, Qing Sun, Erik L Bolager, Iryna Burak, Anna Veselovska, Massimo Fornasier, and Felix Dietrich. Fast training of accurate physics-informed neural networks without gradient descent. In *International Conference on Learning Representations*, volume 2026, pp. 104770–104829, 2026.

Davide Elia De Falco, Enrico Schiassi, and Francesco Calabrò. Least squares with equality constraints extreme learning machines for the resolution of pdes. *Journal of Computational Physics*, pp. 114553, 2025.

Jean Donea, Antonio Huerta, J-Ph Ponthot, and Antonio Rodríguez-Ferran. Arbitrary l agrangian–e ulerian methods. *Encyclopedia of computational mechanics*, 2004.

Vikas Dwivedi and Balaji Srinivasan. Physics informed extreme learning machine (pielm)–a rapid method for the numerical solution of partial differential equations. *Neurocomputing*, 391:96–118, 2020.

Gerhard Dziuk and Charles M Elliott. A fully discrete evolving surface finite element method. *SIAM Journal on Numerical Analysis*, 50(5):2677–2694, 2012.

Alan Edelman, Tomás A Arias, and Steven T Smith. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.

Gene H Golub and Victor Pereyra. The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM Journal on numerical analysis*, 10(2):413–432, 1973.

Boyce E Griffith and Neelesh A Patankar. Immersed methods for fluid–structure interaction. *Annual review of fluid mechanics*, 52(1):421–448, 2020.

Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: theory and applications. *Neurocomputing*, 70(1-3):489–501, 2006.

Mika Juntunen and Rolf Stenberg. Nitsche's method for general boundary conditions. *Mathematics of computation*, 78(267):1353–1374, 2009.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in neural information processing systems*, 34:26548–26560, 2021.

Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023a.

Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36: 35836–35854, 2023b.

Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 1(3):1–27, 2024.

Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.

Zongjia Long, Qi Zhou, Aiqing Zhu, Dong Dai, and Yiqun Liu. Deep neural operator for free boundary problems. *Nature Machine Intelligence*, pp. 1–12, 2026.

Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.

Rishi Mishra, Balaji Srinivasan, Sundararajan Natarajan, Ganapathy Krishnamurthi, et al. Exact constraint enforcement in physics-informed extreme learning machines using null-space projection framework. *arXiv preprint arXiv:2601.10999*, 2026.

Joachim Nitsche. Über ein variationsprinzip zur lösung von dirichlet-problemen bei verwendung von teilräumen, die keinen randbedingungen unterworfen sind. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 36, pp. 9–15. Springer, 1971.

Charles S Peskin. The immersed boundary method. *Acta numerica*, 11:479–517, 2002.

Peter Petersen. *Riemannian geometry*. Springer, 2006.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.

Fei Ren, Pei-Zhi Zhuang, Xiaohui Chen, Hai-Sui Yu, and He Yang. Physics-informed extreme learning machine (pielm) for stefan problems. *Computer Methods in Applied Mechanics and Engineering*, 441:118015, 2025.

Henning Scheufler and Johan Roenby. Twophaseflow: A framework for developing two phase flow solvers in openfoam. *OpenFOAM® Journal*, 3:200–224, 2023.

LE Scriven. On the dynamics of phase growth. *Chemical engineering science*, 50(24):3907–3917, 1995.

Sifan Wang and Paris Perdikaris. Deep learning of free boundary and stefan problems. *Journal of Computational Physics*, 428:109914, 2021.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality is all you need for training physics-informed neural networks. *arXiv preprint arXiv:2203.07404*, 2022.

Gerhard Wanner and Ernst Hairer. *Solving ordinary differential equations II*, volume 375. Springer Berlin Heidelberg New York, 1996.

APPENDIX CONTENTS

A Related work **14**

A.1 Physics-informed neural networks 14

A.2 Neural operators on varying domains 14

A.3 Frozen features and exact boundary constraints 14

A.4 Moving domains and coupled interfaces 14

B Theory and proofs **15**

B.1 Compatible coefficients and frame geometry 15

B.2 Local horizontal transport 17

B.3 Boundary lift and transport trajectories 18

B.4 Rank saturation 20

B.5 Discrete solver 21

C Experimental details **24**

C.1 Common protocol 24

C.2 Analytic reference fields 25

C.3 Burgers with a time-dependent boundary trace 25

C.4 Dynamic disk with prescribed motion 26

C.5 Flow around a moving cylinder 26

C.6 Navier-Stokes flow on a prescribed petal 28

C.7 Topology change through prescribed Cassini splitting 28

C.8 Stefan with a two-phase free boundary 30

C.9 Spherical Scriven bubble growth 30

C.10 Axisymmetric modal Scriven bubble 31

C.11 Complete nullspace refresh interval ablation 33

C.12 Boundary-count ablation configuration 33

D Limitations and Future Work **34**

A Related work

A.1 Physics-informed neural networks

Physics-informed neural networks (PINNs) approximate PDE solutions by minimizing equation residuals together with initial and boundary errors (Raissi et al., 2017). Their mesh-free representation supports irregular domains and inverse problems. Imbalanced gradients across loss terms and difficulties in optimizing the PDE residual can hinder training (Wang et al., 2021, Krishnapriyan et al., 2021). For time-dependent problems, causal training weights encourage earlier times to be resolved before later ones (Wang et al., 2022). PINNs have also been extended to free boundary and Stefan problems by learning the physical fields and the unknown interface jointly (Wang & Perdikaris, 2021). These methods demonstrate joint approximation of physical fields and unknown interfaces. NOVA focuses on advancing such problems without gradient-based training while maintaining the evolving boundary constraints in coefficient space.

A.2 Neural operators on varying domains

Neural operators learn mappings between input functions and PDE solutions. DeepONet represents these mappings through branch and trunk networks (Lu et al., 2019), while the Fourier neural operator (FNO) parameterizes integral kernels in Fourier space (Li et al., 2020). Physics-informed neural operators combine solution data with PDE constraints during training (Li et al., 2024). Once trained, these models can amortize the cost of solving many related problem instances.

Geometry-aware extensions broaden the domains on which these operators can be applied. Geo-FNO learns a deformation to a regular latent domain (Li et al., 2023a), and the geometry-informed neural operator (GINO) combines graph and Fourier operators to process irregular geometries (Li et al., 2023b). The free boundary neural operator learns a transformed evolution together with a mapping that supports prediction on evolving domains (Long et al., 2026). NOVA solves each problem through a fixed spatial feature map and coefficient updates, without learning a solution operator across instances. Its central concern is maintaining boundary compatibility as the geometry and physical state evolve together.

A.3 Frozen features and exact boundary constraints

Freezing the hidden parameters turns a neural representation into a finite trial space. Extreme learning machines determine the output weights through linear algebra (Huang et al., 2006), and PIELM applies this construction to differential equations (Dwivedi & Srinivasan, 2020). Frozen-PINN separates space and time, combining frozen spatial features with coefficient evolution to obtain a causal PDE solver without gradient descent (Datar et al., 2026). NOVA shares this frozen spatial representation and imposes moving boundary constraints through compatible coefficient coordinates. It discretizes the field equations in time before solving for the new compatible coefficients.

Boundary conditions can be imposed through the trial function, the variational formulation, or algebraic constraints. Early neural trial functions embedded prescribed boundary data directly (Lagaris et al., 1998), and subsequent work compares exact constructions with penalty methods and other weak enforcement techniques (Berrone et al., 2023). Nitsche's method imposes boundary conditions through the weak form (Nitsche, 1971), while Baumgarte stabilization damps drift in differentiated constraints through feedback (Baumgarte, 1972). In a frozen trial space, LSE-ELM enforces boundary equations through equality-constrained least squares (De Falco et al., 2025). The NP-PIELM preprint likewise uses a particular solution and a nullspace basis to enforce discrete constraints (Mishra et al., 2026).

As boundaries move, the boundary lift and nullspace basis must be updated. NOVA transports the compatible nullspace frame between scheduled rebuilds to reduce repeated decompositions.

A.4 Moving domains and coupled interfaces

Classical moving-domain methods organize the spatial discretization around the evolving geometry. Arbitrary Lagrangian-Eulerian (ALE) methods use a moving coordinate description (Donea et al., 2004), and evolving surface finite elements discretize PDEs on moving manifolds (Dziuk & Elliott,

2012). Immersed boundary methods couple a moving interface to a fixed background grid (Peskin, 2002, Griffith & Patankar, 2020). NOVA uses fixed mesh-free features and updates the boundary equations in coefficient space, so geometry changes do not require a body-fitted mesh.

For free boundary problems, the interface motion must also be determined from the physical solution. PIELM has been applied to Stefan problems through a coupled nonlinear least-squares formulation for the temperature and interface (Ren et al., 2025). NOVA uses a marching formulation in which the current fields supply the fluxes needed for interface advancement. This connects boundary-compatible field solves with multiphysics coupling at successive time steps, while nullspace transport updates the compatible frame between rebuilds.

B Theory and proofs

We prove the compatible representation, horizontal transport and rank saturation, then describe the discrete frame and field updates. The boundary matrix is $B \in \mathbb{R}^{n_b \times M}$, and $g_b \in \mathbb{R}^{n_b}$ contains the boundary data. We write B^+ for the exact Moore-Penrose inverse and B_τ^+ for the truncated inverse in equation 2. The notation I_ℓ denotes the $\ell \times \ell$ identity, and $\|\cdot\|_2$ denotes the Euclidean vector norm or spectral matrix norm. The Frobenius inner product and norm are

$$\langle X, Y\rangle_F = \mathrm{tr}(X^\top Y), \qquad \|X\|_F^2 = \langle X, X\rangle_F,$$

where tr is the sum of diagonal entries.

B.1 Compatible coefficients and frame geometry

Definition B.1 (Orthonormal frames and their subspaces)**.** For integers $0 \le m_0 \le M$, define

$$\begin{aligned} \mathrm{St}(M, m_0) &= \{V \in \mathbb{R}^{M\times m_0} \mid V^\top V = I_{m_0}\}, \\ \mathrm{Gr}(M, m_0) &= \{E \subseteq \mathbb{R}^M \mid E \text{ is a linear subspace, } \dim E = m_0\}, \\ O(m_0) &= \{Q \in \mathbb{R}^{m_0\times m_0} \mid Q^\top Q = I_{m_0}\}. \end{aligned} \tag{10}$$

The Stiefel manifold stores an ordered orthonormal frame. The Grassmann manifold stores its span. The projection $\pi(V) = \mathrm{range}\, V$ maps $\mathrm{St}(M, m_0)$ to $\mathrm{Gr}(M, m_0)$. Its fibre over E is $\pi^{-1}(E)$, the set of frames spanning E. The right action of $O(m_0)$ is $V \mapsto VQ$.

The tangent space $T_V\mathrm{St}(M, m_0)$ consists of velocities of differentiable frame curves through V (Petersen, 2006). Vertical velocities stay within a fibre. Horizontal velocities are their Frobenius orthogonal complement in the tangent space. We use the standard matrix quotient geometry of Edelman et al. (1998).

Proposition B.2 (Frame geometry under the Frobenius metric)**.** *For $V \in \mathrm{St}(M, m_0)$, the fibre is $\pi^{-1}(\pi(V)) = \{VQ \mid Q \in O(m_0)\}$. The tangent, vertical and horizontal spaces are*

$$\begin{aligned} T_V\mathrm{St}(M, m_0) &= \{X \in \mathbb{R}^{M\times m_0} \mid V^\top X + X^\top V = 0\}, \\ \mathcal{V}_V &= \{VS \mid S \in \mathbb{R}^{m_0\times m_0},\ S^\top = -S\}, \\ \mathcal{H}_V &= \{X \in T_V\mathrm{St}(M, m_0) \mid V^\top X = 0\}. \end{aligned} \tag{11}$$

Every tangent vector splits orthogonally into horizontal and vertical components,

$$X = (I_M - VV^\top)X + V(V^\top X).$$

Proof. If W is another orthonormal frame with the same span, then $VV^\top W = W$. Setting $Q = V^\top W$ gives

$$Q^\top Q = W^\top VV^\top W = W^\top W = I_{m_0}, \qquad VQ = W.$$

Thus the action is transitive on each fibre. It is free because $VQ = V$ implies $Q = I_{m_0}$.

The constraint map $G(V) = V^\top V$ takes values in symmetric matrices. Its derivative satisfies

$$dG_V[X] = V^\top X + X^\top V, \qquad dG_V[VC/2] = C \quad \text{for every } C = C^\top.$$

The derivative is therefore onto, so $\mathrm{St}(M, m_0)$ is a smooth level set. The free action of the compact group $O(m_0)$ gives the smooth quotient $\mathrm{Gr}(M, m_0)$ and the corresponding principal bundle.

Differentiating $V^\top V = I_{m_0}$ along a frame curve gives the tangent condition. Conversely, suppose $V^\top X + X^\top V = 0$ and define

$$\gamma(s) = (V + sX)(I_{m_0} + s^2 X^\top X)^{-1/2}.$$

The inverse square root is that of a positive definite matrix. The tangent condition yields

$$\begin{aligned}(V + sX)^\top (V + sX) &= I_{m_0} + s^2 X^\top X,\\ \gamma(s)^\top \gamma(s) = I_{m_0}, \qquad \gamma(0) &= V, \qquad \dot\gamma(0) = X.\end{aligned}$$

The derivative at zero follows because the normalizing factor is $I_{m_0} + O(s^2)$. Hence every vector satisfying the tangent condition is realized by a frame curve.

A curve in the fibre has the form $VQ(s)$. At $Q(0) = I_{m_0}$, write $S = \dot Q(0)$. Differentiating $Q(s)^\top Q(s) = I_{m_0}$ gives

$$S^\top + S = 0, \qquad \frac{d}{ds}\big(VQ(s)\big)\Big|_{s=0} = VS.$$

Every such direction is realized by $Q(s) = \exp(sS)$, where exp is the matrix exponential. This proves the formula for the vertical space.

For tangent X, the matrix $A = V^\top X$ is skew symmetric. If $V^\top X = 0$, then

$$\langle X, VS\rangle_F = \mathrm{tr}\big((V^\top X)^\top S\big) = 0 \quad \text{for every } S^\top = -S.$$

Conversely, orthogonality to the vertical vector VA gives

$$0 = \langle X, VA\rangle_F = \mathrm{tr}(A^\top A) = \|A\|_F^2.$$

Hence $A = 0$, which proves the formula for the horizontal space. In the stated splitting, $(I_M - VV^\top)X$ is horizontal and $V(V^\top X)$ is vertical, so the components are orthogonal. □

Boundary-compatible coefficients. At a fixed time, let $r = \operatorname{rank} B$, $m_0 = M - r$, and let $V_0 \in \mathrm{St}(M, m_0)$ span $\ker B$. If $g_b \in \operatorname{range} B$, then $c_p = B^+ g_b$ satisfies $Bc_p = g_b$. Consequently,

$$Bc = g_b \quad \Longleftrightarrow \quad B(c - c_p) = 0 \quad \Longleftrightarrow \quad c - c_p \in \operatorname{range} V_0.$$

Every compatible coefficient vector therefore has the unique representation

$$c = c_p + V_0 z, \qquad z = V_0^\top (c - c_p) \in \mathbb{R}^{m_0}.$$

For incompatible data, projection onto the orthogonal complement of $\operatorname{range} B$ gives

$$(I_{n_b} - BB^+)(Bc - g_b) = -(I_{n_b} - BB^+)g_b.$$

This residual component cannot be removed by any coefficient vector.

For the truncated representation, use the retained rank $r = \operatorname{rank} B_\tau$ and $m_0 = M - r$. Take $V_0 \in \mathrm{St}(M, m_0)$ to span $\ker B_\tau$ and set $c = B_\tau^+ g_b + V_0 z$ with $z \in \mathbb{R}^{m_0}$. Since $B_\tau V_0 = 0$,

$$\begin{aligned}Bc - g_b &= (B - B_\tau)c + B_\tau c - g_b\\ &= (B - B_\tau)c - (I_{n_b} - B_\tau B_\tau^+)g_b.\end{aligned}$$

The triangle inequality gives

$$\|Bc - g_b\|_2 \le \|B - B_\tau\|_2 \|c\|_2 + \|(I_{n_b} - B_\tau B_\tau^+)g_b\|_2. \tag{12}$$

Exact compatibility with B_τ requires $g_b \in \operatorname{range} B_\tau$ and does not imply compatibility with B.

B.2 LOCAL HORIZONTAL TRANSPORT

Fix a local interval $I = [t_0, t_0 + h]$ with $h > 0$. Assume independently that $B \in C^1(I, \mathbb{R}^{n_b \times M})$ and $\operatorname{rank} B(t) = r$ for every $t \in I$. Set $m_0 = M - r$. Shortening the interval alone does not establish constant rank. Different intervals may have different ranks, and no continuous transport claim is made across rank changes. A numerical rank retained between rebuilds does not establish this exact-rank hypothesis. Applying the theorem to a C^1 constant-rank truncated curve $B_\tau(t)$ requires its own derivative $\dot{B}_\tau$.

Lemma B.3 (Local frames and differentiated compatibility)**.** *Under these assumptions, $P_0 = I_M - B^+B$ is a C^1 orthogonal projector onto $\ker B$, and local C^1 orthonormal kernel frames exist. At every complete kernel frame V_0, one has*

$$BB^+\dot{B}V_0 = \dot{B}V_0. \tag{13}$$

If $g_b \in C^1(I, \mathbb{R}^{n_b})$ and $g_b(t) \in \operatorname{range} B(t)$ on I, then every compatible value $Bc_p = g_b$ satisfies

$$BB^+(\dot{g}_b - \dot{B}c_p) = \dot{g}_b - \dot{B}c_p. \tag{14}$$

Proof. The Moore-Penrose inverse is C^1 on a fixed-rank matrix curve (Golub & Pereyra, 1973), so $P_0 = I_M - B^+B$ is C^1. The Moore-Penrose identities make B^+B symmetric and idempotent. Hence

$$\begin{aligned} P_0^\top &= I_M - (B^+B)^\top = P_0, \\ P_0^2 &= I_M - 2B^+B + (B^+B)^2 = I_M - B^+B = P_0, \\ BP_0 &= B - BB^+B = 0. \end{aligned}$$

The last identity gives $\operatorname{range} P_0 \subseteq \ker B$. Conversely, $P_0 v = v$ for every $v \in \ker B$, so $\operatorname{range} P_0 = \ker B$.

Fix $t_* \in I$ and a complete orthonormal kernel frame F_* at t_*. Define $H(t) = F_*^\top P_0(t) F_*$. Since $H(t_*) = I_{m_0}$, continuity makes $H(t)$ positive definite near t_*. The inverse square root varies smoothly on positive definite matrices, so

$$F(t) = P_0(t)F_*\big(F_*^\top P_0(t)F_*\big)^{-1/2} \tag{15}$$

is C^1. The projector identities verify its frame properties directly,

$$\begin{aligned} B(t)F(t) &= 0, \\ F(t)^\top F(t) &= H(t)^{-1/2}F_*^\top P_0(t)^2 F_* H(t)^{-1/2} = I_{m_0}, \\ F(t_*) &= F_*. \end{aligned}$$

Since $\dim \ker B(t) = m_0$, these m_0 orthonormal columns form a complete kernel frame. Differentiating $BF = 0$ now gives

$$\dot{B}F = -B\dot{F}, \qquad BB^+\dot{B}F = -BB^+B\dot{F} = -B\dot{F} = \dot{B}F.$$

Taking $F_* = V_0$ at t_* proves equation 13.

For a compatible value c_p^* at t_*, choose

$$\begin{aligned} a &= F_*^\top\big(c_p^* - B(t_*)^+ g_b(t_*)\big), \\ \widehat{c}_p(t) &= B(t)^+ g_b(t) + F(t)a. \end{aligned}$$

The difference defining a lies in $\ker B(t_*)$, so $\widehat{c}_p(t_*) = c_p^*$. This local lift is C^1 and satisfies $B\widehat{c}_p = g_b$. Differentiation and projection give

$$\begin{aligned} \dot{g}_b - \dot{B}\widehat{c}_p &= B\dot{\widehat{c}}_p, \\ BB^+(\dot{g}_b - \dot{B}\widehat{c}_p) &= BB^+B\dot{\widehat{c}}_p = \dot{g}_b - \dot{B}\widehat{c}_p. \end{aligned}$$

Evaluating at t_* proves equation 14 for the prescribed value. □

Proposition B.4 (Velocities following the kernel)**.** *At a fixed time in I, let $V_0 \in \operatorname{St}(M, m_0)$ span $\ker B$ and set $W = -B^+\dot{B}V_0$. Every solution of $BX = -\dot{B}V_0$ has the unique form*

$$X = W + V_0 S, \qquad S \in \mathbb{R}^{m_0 \times m_0}.$$

Proof. Lemma B.3 gives

$$BW = -BB^+\dot{B}V_0 = -\dot{B}V_0.$$

Thus any solution X satisfies $B(X - W) = 0$. Each column of $X - W$ lies in the span of the complete frame V_0, so

$$X - W = V_0 S, \qquad S = V_0^\top (X - W).$$

This expression also proves uniqueness of S. Conversely, every such matrix satisfies $B(W + V_0 S) = BW = -\dot{B}V_0$. □

Corollary B.5 (Horizontal projection and minimum norm)**.** *Under the hypotheses of Proposition B.4, the admissible orthonormal frame velocities are exactly $X = W + V_0 S$ with $S^\top = -S$. Their horizontal and vertical components are*

$$(I_M - V_0 V_0^\top)X = W, \qquad V_0 V_0^\top X = V_0 S.$$

The velocity W is the unique horizontal choice and the unique minimizer of $\|X\|_F$.

Proof. Since $\operatorname{range} B^+ = (\ker B)^\perp$, one has $V_0^\top W = 0$. Therefore

$$V_0^\top X = S, \qquad V_0^\top X + X^\top V_0 = S + S^\top.$$

The tangent condition is equivalent to $S^\top = -S$. Proposition B.2 gives the stated projections and selects $S = 0$ for horizontality. Moreover,

$$\langle W, V_0 S\rangle_F = 0, \qquad \|V_0 S\|_F^2 = \operatorname{tr}(S^\top V_0^\top V_0 S) = \|S\|_F^2.$$

Expanding the squared norm gives

$$\|W + V_0 S\|_F^2 = \|W\|_F^2 + \|S\|_F^2. \tag{16}$$

The minimum is attained uniquely at $S = 0$. □

B.3 Boundary lift and transport trajectories

Proposition B.6 (Horizontal velocity of a boundary lift)**.** *Retain the local assumptions on B and suppose $g_b \in C^1(I, \mathbb{R}^{n_b})$ with $g_b(t) \in \operatorname{range} B(t)$ on I. For a compatible $Bc_p = g_b$ and a complete orthonormal kernel frame V_0, every velocity satisfying $By = \dot{g}_b - \dot{B}c_p$ has the unique form*

$$y = B^+(\dot{g}_b - \dot{B}c_p) + V_0 a, \qquad a \in \mathbb{R}^{m_0}.$$

The condition $V_0^\top y = 0$ uniquely selects $a = 0$.

Proof. Set $y_p = B^+(\dot{g}_b - \dot{B}c_p)$. The affine part of Lemma B.3 gives

$$By_p = BB^+(\dot{g}_b - \dot{B}c_p) = \dot{g}_b - \dot{B}c_p.$$

For any other solution y, the difference satisfies $B(y - y_p) = 0$. Completeness and orthonormality of V_0 therefore give

$$y - y_p = V_0 a, \qquad a = V_0^\top (y - y_p).$$

Since $V_0^\top B^+ = 0$, projection onto the kernel frame yields

$$V_0^\top y = V_0^\top B^+(\dot{g}_b - \dot{B}c_p) + V_0^\top V_0 a = a.$$

Thus horizontality is equivalent to $a = 0$, which uniquely determines y. □

Proof of Theorem 2.1. Proposition B.4 and Corollary B.5 give

$$\begin{aligned} X &= -B^+\dot{B}V_0 + V_0 S, \\ V_0^\top X + X^\top V_0 = 0 \quad &\Longleftrightarrow \quad S + S^\top = 0, \\ V_0^\top X = 0 \quad &\Longleftrightarrow \quad S = 0. \end{aligned}$$

Thus the unique horizontal velocity is equation 7. Equation equation 16 proves its minimum Frobenius norm. Proposition B.6 gives the lift velocity equation 8. It remains to show that these velocities generate compatible trajectories from the prescribed initial values.

Take a local frame F from Lemma B.3 near t_0 and define $A(t) = F(t)^\top \dot{F}(t)$. Differentiating $F^\top F = I_{m_0}$ shows that $A^\top = -A$. Solve

$$\dot{Q} = -AQ, \qquad Q(t_0) = F(t_0)^\top V_0(t_0).$$

The initial matrix is orthogonal because $F(t_0)$ and $V_0(t_0)$ are complete orthonormal frames of the same kernel. Along the solution,

$$\frac{d}{dt}(Q^\top Q) = \dot{Q}^\top Q + Q^\top \dot{Q} = -Q^\top (A^\top + A)Q = 0.$$

Hence $Q^\top Q = I_{m_0}$. Define $V_0 = FQ$. This frame has the prescribed initial value and satisfies

$$\begin{aligned} BV_0 &= BFQ = 0, \\ V_0^\top V_0 &= Q^\top F^\top FQ = I_{m_0}, \\ V_0^\top \dot{V}_0 &= Q^\top F^\top (\dot{F}Q + F\dot{Q}) = Q^\top (AQ + \dot{Q}) = 0. \end{aligned}$$

Differentiating the first identity gives $B\dot{V}_0 = -\dot{B}V_0$. The unique horizontal velocity already proved therefore yields

$$\dot{V}_0 = -B^+ \dot{B} V_0.$$

The coefficient $-B^+\dot{B}$ is continuous on the compact interval I, so its linear matrix ODE has a unique solution on I. The frame and kernel identities are preserved under limits. At any endpoint reached by the local construction, they therefore provide admissible initial data for the next local frame. Uniqueness makes the local solutions agree on overlaps, giving the claimed horizontal trajectory on all of I.

For the lift, define $\widehat{c}_p(t) = B(t)^+ g_b(t)$ and choose

$$\begin{aligned} \dot{a}(t) &= -V_0(t)^\top \dot{\widehat{c}}_p(t), \\ a(t_0) &= V_0(t_0)^\top \big(c_p(t_0) - \widehat{c}_p(t_0)\big), \\ c_p(t) &= \widehat{c}_p(t) + V_0(t)a(t). \end{aligned}$$

The initial difference lies in $\ker B(t_0)$, so this construction has the prescribed initial value. At every time,

$$Bc_p = BB^+ g_b + BV_0 a = g_b.$$

Differentiating the representation and using the horizontal frame gives

$$\begin{aligned} \dot{c}_p &= \dot{\widehat{c}}_p + \dot{V}_0 a + V_0 \dot{a}, \\ V_0^\top \dot{c}_p &= V_0^\top \dot{\widehat{c}}_p + V_0^\top \dot{V}_0 a + \dot{a} = V_0^\top \dot{\widehat{c}}_p + \dot{a} = 0. \end{aligned}$$

Differentiating $Bc_p = g_b$ and applying Proposition B.6 now give

$$\dot{c}_p = B^+(\dot{g}_b - \dot{B}c_p) = -B^+\dot{B}c_p + B^+\dot{g}_b.$$

This is a linear ODE with continuous coefficients and forcing, so the prescribed initial lift determines a unique solution on I. This completes both parts of the theorem. □

The horizontal lift depends on its initial value and generally differs from the pointwise minimum-norm lift $B^+ g_b$. Algorithm 1 recomputes $c_p = Pg_b$ at each candidate geometry. It does not integrate the lift ODE.

For the horizontal frame V_0 and fixed $R_k \in \mathbb{R}^{m_0 \times k}$ with $0 \le k \le m_0$ and $R_k^\top R_k = I_k$, define $V_{\text{eff}} = V_0 R_k$. Then

$$\begin{aligned} V_{\text{eff}}^\top V_{\text{eff}} &= R_k^\top V_0^\top V_0 R_k = I_k, \\ \dot{V}_{\text{eff}} &= \dot{V}_0 R_k = -B^+ \dot{B} V_0 R_k = -B^+ \dot{B} V_{\text{eff}}, \\ V_{\text{eff}}^\top \dot{V}_{\text{eff}} &= R_k^\top V_0^\top \dot{V}_0 R_k = 0. \end{aligned}$$

This construction holds R_k fixed. When $0 < k < m_0$, the compressed horizontal condition alone does not determine a unique motion within the full kernel. Indeed, one may add $V_0 Z$ to the displayed velocity for any nonzero $Z \in \mathbb{R}^{m_0 \times k}$ satisfying $R_k^\top Z = 0$, because

$$B(V_0 Z) = 0, \qquad V_{\mathrm{eff}}^\top V_0 Z = R_k^\top Z = 0.$$

The transported frame also need not remain the instantaneous leading interior singular subspace.

B.4 Rank saturation

At fixed t, the trace operator acts on the entire boundary before sampling. Define

$$\mathcal{T}_t c = \mathcal{B}_t[\Phi c]|_{\Gamma(t)}, \qquad \mathcal{Y}_t = \operatorname{range} \mathcal{T}_t, \qquad r_{\mathrm{tr}} = \dim \mathcal{Y}_t.$$

By linearity, $\mathcal{Y}_t$ is spanned by the M traces $\mathcal{B}_t[\phi_j]|_{\Gamma(t)}$, so $r_{\mathrm{tr}} \le M$. The sampling map S_{X_b} evaluates these traces at the selected points and gives

$$B = S_{X_b} \mathcal{T}_t \in \mathbb{R}^{n_b \times M}.$$

Proof of Theorem 2.2. If $\mathcal{T}_t c = 0$, then sampling also gives zero. Thus

$$\mathcal{T}_t c = 0 \quad \Longrightarrow \quad Bc = S_{X_b} \mathcal{T}_t c = 0, \qquad \ker \mathcal{T}_t \subseteq \ker B.$$

Suppose S_{X_b} is injective on $\mathcal{Y}_t$. For any $c \in \ker B$, the trace $w = \mathcal{T}_t c$ belongs to $\mathcal{Y}_t$ and satisfies $S_{X_b} w = Bc = 0$. Injectivity implies $w = 0$, so $c \in \ker \mathcal{T}_t$. This proves equality of the kernels.

Conversely, suppose the kernels are equal and take $w \in \mathcal{Y}_t$ with $S_{X_b} w = 0$. By the definition of $\mathcal{Y}_t$, there is a coefficient vector c such that $w = \mathcal{T}_t c$. Then

$$S_{X_b} w = 0 \quad \Longrightarrow \quad Bc = 0 \quad \Longrightarrow \quad c \in \ker \mathcal{T}_t \quad \Longrightarrow \quad w = 0.$$

Thus sampling is injective on $\mathcal{Y}_t$.

Both B and $\mathcal{T}_t$ have domain $\mathbb{R}^M$. Rank-nullity gives

$$\dim \ker B = M - \operatorname{rank} B, \qquad \dim \ker \mathcal{T}_t = M - r_{\mathrm{tr}}.$$

Together with the kernel inclusion, these identities show

$$\ker B = \ker \mathcal{T}_t \quad \Longleftrightarrow \quad \operatorname{rank} B = r_{\mathrm{tr}}.$$

Since B has n_b rows, the rank condition requires $n_b \ge r_{\mathrm{tr}}$. This dimension bound alone is insufficient because the sample locations must distinguish all traces in $\mathcal{Y}_t$.

For inhomogeneous data $g \in \mathcal{Y}_t$, choose c_g with $\mathcal{T}_t c_g = g$. If $Bc = S_{X_b} g$, then

$$\begin{aligned} B(c - c_g) &= Bc - S_{X_b} \mathcal{T}_t c_g = S_{X_b} g - S_{X_b} g = 0, \\ \mathcal{T}_t(c - c_g) &= 0, \\ \mathcal{T}_t c &= \mathcal{T}_t c_g = g. \end{aligned}$$

The second line uses kernel equality. The last equality holds as a boundary function, so it includes unsampled points. The homogeneous case follows by taking $g = 0$. □

If $\operatorname{rank} B < r_{\mathrm{tr}}$, rank-nullity instead gives

$$\dim \ker B - \dim \ker \mathcal{T}_t = r_{\mathrm{tr}} - \operatorname{rank} B > 0.$$

Thus there is a coefficient perturbation $\delta c \in \ker B \setminus \ker \mathcal{T}_t$. For any c_g representing $g \in \mathcal{Y}_t$,

$$\begin{aligned} B(c_g + \delta c) &= Bc_g = S_{X_b} g, \\ \mathcal{T}_t(c_g + \delta c) - g &= \mathcal{T}_t \delta c \not\equiv 0. \end{aligned}$$

Its boundary trace vanishes at the sampled constraints but is nonzero somewhere else on the boundary. This identifies the information that insufficient sampling leaves unconstrained.

To obtain rank saturation, suppose $\operatorname{rank} B = r_{\mathrm{tr}}$ and add samples while retaining the existing points. Write $\widetilde{X}_b \supseteq X_b$ and $\widetilde{B} = S_{\widetilde{X}_b}\mathcal{T}_t$. The old rows remain among the new constraints, so $\operatorname{rank}\widetilde{B} \geq \operatorname{rank} B$. On the other hand,

$$\operatorname{rank}\widetilde{B} = \dim S_{\widetilde{X}_b}(\mathcal{Y}_t) \leq \dim \mathcal{Y}_t = r_{\mathrm{tr}}.$$

Combining the bounds gives

$$r_{\mathrm{tr}} = \operatorname{rank} B \leq \operatorname{rank}\widetilde{B} \leq r_{\mathrm{tr}}.$$

Thus the exact rank remains r_{tr}. This conclusion fixes t, the geometry, the features and the boundary operator. If the points are replaced rather than retained, injectivity must be checked for the new configuration. A temporary rank plateau below r_{tr} does not establish sufficient sampling.

Numerically, r in equation 2 counts singular values above $\tau\sigma_1(B)$. Dense reference ranks estimate the trace dimension resolved at this threshold and need not equal the exact r_{tr}. Changing the sampling can move singular values across the relative cutoff, so numerical rank need not increase monotonically with point count. Matching reference ranks and stable errors on independent test points provide numerical evidence for a sampling choice. They do not establish exact injectivity, and the theorem does not guarantee an error plateau or exact matching of boundary data outside $\mathcal{Y}_t$.

B.5 Discrete solver

B.5.1 Frame updates

Boundary derivatives. Let $x_i^b(t)$ be C^1 labeled boundary trajectories and assume $F_j(t,x) = \mathcal{B}_t[\phi_j](x)$ has a C^1 extension near them. Then

$$\dot{B}_{ij}(t) = \partial_t F_j(t, x_i^b(t)) + \nabla_x F_j(t, x_i^b(t))^\top \dot{x}_i^b(t). \tag{17}$$

For Dirichlet data with $\phi_j \in C^1$, the features have no explicit time dependence, so

$$B_{ij}(t) = \phi_j(x_i^b(t)), \qquad \dot{B}_{ij}(t) = \nabla\phi_j(x_i^b(t))^\top \dot{x}_i^b(t).$$

For a Neumann operator, assume $\phi_j \in C^2$ and a C^1 unit normal $\nu_i(t)$. Differentiating $B_{ij} = \nu_i^\top \nabla\phi_j(x_i^b)$ by the product and chain rules gives

$$\dot{B}_{ij} = \dot{\nu}_i^\top \nabla\phi_j(x_i^b) + \nu_i^\top \nabla^2\phi_j(x_i^b)\dot{x}_i^b. \tag{18}$$

A Robin operator $a(t,x)u + b(t,x)\partial_\nu u$ has C^1 scalar coefficients a, b. Write $a_i(t) = a(t, x_i^b(t))$ and $b_i(t) = b(t, x_i^b(t))$, so that

$$\dot{a}_i = \partial_t a(t, x_i^b) + \nabla_x a(t, x_i^b)^\top \dot{x}_i^b, \qquad \dot{b}_i = \partial_t b(t, x_i^b) + \nabla_x b(t, x_i^b)^\top \dot{x}_i^b.$$

Under the same feature and normal regularity as for Neumann data, the full derivative is

$$\begin{aligned} B_{ij} &= a_i\phi_j(x_i^b) + b_i\nu_i^\top\nabla\phi_j(x_i^b), \\ \dot{B}_{ij} &= \dot{a}_i\phi_j(x_i^b) + a_i\nabla\phi_j(x_i^b)^\top\dot{x}_i^b + \dot{b}_i\nu_i^\top\nabla\phi_j(x_i^b) \\ &\quad + b_i\dot{\nu}_i^\top\nabla\phi_j(x_i^b) + b_i\nu_i^\top\nabla^2\phi_j(x_i^b)\dot{x}_i^b. \end{aligned}$$

A secant approximation to $\dot{B}$ requires matching boundary row labels. If the rows cannot be matched, rebuild the frame. Unchanged spatial features still permit using the stored full-coefficient histories.

Fixed-rank row update. Here r is the retained numerical rank. At a full rebuild, store $V_r \in \mathbb{R}^{M\times r}$ from equation 2. Let $\mathrm{qf}(Y)$ denote the orthonormal factor of a thin QR factorization of a full-column-rank matrix Y. Assuming $\operatorname{rank}(BV_r) = r$ at the new geometry, form

$$\begin{aligned} Q_L &= \mathrm{qf}(BV_r), \qquad Q_R = \mathrm{qf}(B^\top Q_L), \\ C &= Q_L^\top B Q_R = U_C\Sigma_C V_C^\top. \end{aligned} \tag{19}$$

Here $Q_L \in \mathbb{R}^{n_b\times r}$ and $Q_R \in \mathbb{R}^{M\times r}$ have orthonormal columns, and $C \in \mathbb{R}^{r\times r}$. To see why the second QR factorization is well defined, take $y \in \operatorname{range} B$. Then

$$B^\top y = 0 \quad\Longrightarrow\quad y \in \operatorname{range} B \cap \ker B^\top = \operatorname{range} B \cap (\operatorname{range} B)^\perp = \{0\}.$$

Thus $B^\top$ is injective on $\operatorname{range} B$, which contains the columns of Q_L. Hence $B^\top Q_L$ has rank r. If its thin QR factorization is $B^\top Q_L = Q_R T_R$, then T_R is invertible and

$$C = (B^\top Q_L)^\top Q_R = T_R^\top Q_R^\top Q_R = T_R^\top.$$

Thus C is nonsingular under the stated rank assumption. Retain all r singular values of C and set

$$P = (Q_R V_C)\Sigma_C^{-1}(Q_L U_C)^\top, \qquad V_r \leftarrow Q_R V_C. \tag{20}$$

To verify this inverse, write the rank-r approximation as

$$Q_L C Q_R^\top = (Q_L U_C)\Sigma_C (Q_R V_C)^\top.$$

Both outer factors have orthonormal columns, so this is a thin SVD and equation 20 is its Moore-Penrose inverse. It need not equal a fresh truncated inverse of B. Constant rank of B alone does not ensure the assumed overlap with V_r over a finite step. A singular core causes an error. At full rebuilds, $P = B_\tau^+$.

Transport and correction. The implementation requires $k > 0$ and stops if no resolved kernel direction remains. With the current compressed frame V_{eff} and step size Δt, form the Euler transport prediction and apply two corrections,

$$\begin{aligned} \widetilde{V}^{(0)} &= V_{\text{eff}} - \Delta t\, P\dot{B}V_{\text{eff}}, \\ \widetilde{V}^{(j+1)} &= (I_M - PB)\widetilde{V}^{(j)}, \qquad j = 0, 1, \\ \widetilde{V} &= \widetilde{V}^{(2)}. \end{aligned}$$

Then take one Newton-Schulz step

$$V_{\text{new}} = \mathrm{NS}_1(\widetilde{V}) = \tfrac{1}{2}\widetilde{V}(3I_k - \widetilde{V}^\top \widetilde{V}). \tag{21}$$

Let $E = \widetilde{V}^\top \widetilde{V} - I_k$. This matrix is symmetric, and equation 21 becomes $V_{\text{new}} = \widetilde{V}(I_k - E/2)$. Since all factors below are polynomials in E, they commute, giving

$$\begin{aligned} V_{\text{new}}^\top V_{\text{new}} - I_k &= (I_k - E/2)(I_k + E)(I_k - E/2) - I_k \\ &= (I_k - E + E^2/4)(I_k + E) - I_k \\ &= -\tfrac{3}{4}E^2 + \tfrac{1}{4}E^3. \end{aligned}$$

In particular,

$$\|V_{\text{new}}^\top V_{\text{new}} - I_k\|_2 \le \tfrac{3}{4}\|E\|_2^2 + \tfrac{1}{4}\|E\|_2^3.$$

Thus small orthogonality errors become second-order quantities, but the update is not exact orthonormalization. Likewise, correction with the approximate P does not assert exact membership in $\ker B$. For $c = Pg_b + V_{\text{new}} z$, substitution gives the boundary residual

$$\begin{aligned} Bc - g_b &= BPg_b + BV_{\text{new}} z - g_b \\ &= (BP - I_{n_b})g_b + BV_{\text{new}} z. \end{aligned}$$

If $B^+\dot{B}$ is C^1 on I, the exact transported frame is C^2. Taylor expansion at t gives

$$\begin{aligned} V_{\text{eff}}(t + \Delta t) &= V_{\text{eff}}(t) + \Delta t\, \dot{V}_{\text{eff}}(t) + O(\Delta t^2) \\ &= V_{\text{eff}}(t) - \Delta t\, B(t)^+ \dot{B}(t) V_{\text{eff}}(t) + O(\Delta t^2). \end{aligned}$$

Thus Euler transport with the exact coefficient has local frame error $O(\Delta t^2)$. This rate uses stronger regularity than the C^1 assumption on B in Theorem 2.1. The scale $\|B^+\dot{B}\|_2$ explains sensitivity to small nonzero singular values. The implemented step additionally incurs errors from the retained-rank approximation and finite corrections.

Refresh events. With $K = 1$, every step rebuilds the frame. With $K > 1$, at most $K - 1$ transport steps occur between scheduled rebuilds, and the current boundary matrix is still evaluated at every step. A frame change alone does not invalidate full-coefficient histories when the spatial features are unchanged.

B.5.2 TIME DISCRETIZATION IN COMPATIBLE COORDINATES

The fixed features give $u_h(x,t) = \Phi(x)c(t)$. For a complete frame, differentiating $c = c_p + V_0 z$ gives

$$\dot{c} = \dot{c}_p + \dot{V}_0 z + V_0 \dot{z}.$$

Substituting into the interior collocation relation $\Phi_i \dot{c} \approx R_i(c,t)$ and moving the known lift and frame terms to the right gives

$$\begin{aligned}\Phi_i \dot{c}_p + \Phi_i \dot{V}_0 z + \Phi_i V_0 \dot{z} &\approx R_i(c,t),\\ \Phi_i V_0 \dot{z} &\approx R_i(c,t) - \Phi_i \dot{c}_p - \Phi_i \dot{V}_0 z.\end{aligned}$$

This is equation 4. The implementation instead discretizes the full coefficient history and then solves for the new compatible coordinates.

At a fixed candidate geometry, let $\Phi_{i,n+1} = \Phi(X_i(t_{n+1}))$ and let $R_{i,n+1}(c)$ be the interior right-hand side defined in Section 2.3. The values $c_n, c_{n-1}, c_{p,n+1}, V_{\mathrm{eff},n+1}$ are known. With constant step Δt, BE and BDF2 give

$$\Phi_{i,n+1}(c_{n+1} - c_n) \approx \Delta t\, R_{i,n+1}(c_{n+1}), \tag{22}$$

$$\Phi_{i,n+1}(3c_{n+1} - 4c_n + c_{n-1}) \approx 2\Delta t\, R_{i,n+1}(c_{n+1}). \tag{23}$$

BE is used for startup and declared history restarts. Each history field $\Phi_{i,n+1}c_j$ is evaluated at the same current physical points. This approximates the laboratory-frame partial time derivative. Substituting $c_{n+1} = c_{p,n+1} + V_{\mathrm{eff},n+1}z$ gives equation 6 without a preliminary full-state solve. Differencing reduced coordinates in different frames would omit lift and frame changes.

Linear or linearized solve. Write the right-hand side of the chosen linear or linearized subproblem as $q_{n+1} + J_{n+1}c$, with constant term q_{n+1} and coefficient matrix J_{n+1}. For Taylor linearization about $\bar{c}$,

$$\begin{aligned}J_{n+1} &= \partial_c R_{i,n+1}(\bar{c}),\\ q_{n+1} &= R_{i,n+1}(\bar{c}) - J_{n+1}\bar{c},\\ R_{i,n+1}(c) &\approx R_{i,n+1}(\bar{c}) + J_{n+1}(c - \bar{c}) = q_{n+1} + J_{n+1}c.\end{aligned}$$

The time coefficients are

$$(a_0, a_1, a_2, \beta) = \begin{cases}(3,4,-1,2), & \text{for BDF2,}\\ (1,1,0,1), & \text{for BE.}\end{cases}$$

The linearized residual is therefore

$$\Phi_{i,n+1}(a_0 c_{n+1} - a_1 c_n - a_2 c_{n-1}) - \beta\Delta t\,(q_{n+1} + J_{n+1}c_{n+1}).$$

Dividing by a_0 and grouping the new coefficient gives the normalized residual $A_{n+1}c_{n+1} - b_{n+1}$, where

$$A_{n+1} = \Phi_{i,n+1} - \frac{\beta\Delta t}{a_0} J_{n+1}, \tag{24}$$

$$b_{n+1} = \Phi_{i,n+1}\left(\frac{a_1}{a_0}c_n + \frac{a_2}{a_0}c_{n-1}\right) + \frac{\beta\Delta t}{a_0} q_{n+1}. \tag{25}$$

The older-history term is omitted for BE. Let W_i be the diagonal matrix of nonnegative interior residual weights. Substituting the new compatible representation into the weighted residual gives

$$\begin{aligned}W_i^{1/2}(A_{n+1}c_{n+1} - b_{n+1}) &= W_i^{1/2}\big[A_{n+1}(c_{p,n+1} + V_{\mathrm{eff},n+1}z) - b_{n+1}\big]\\ &= Dz - d,\end{aligned}$$

where

$$D = W_i^{1/2} A_{n+1} V_{\mathrm{eff},n+1}, \qquad d = W_i^{1/2}(b_{n+1} - A_{n+1}c_{p,n+1}).$$

Take an SVD $D = U_D \Sigma_D Q_D^\top$. Let u_j, v_j be the corresponding left and right singular vectors, so $Dv_j = \sigma_j(D)u_j$. Retain the index set

$$\mathcal{J} = \{j \mid \sigma_j(D) > \tau_{\mathrm{LS}}\sigma_1(D)\}.$$

For $z = \sum_{j\in\mathcal{J}} \alpha_j v_j$, orthogonality of the left singular vectors gives

$$\|Dz - d\|_2^2 = \sum_{j\in\mathcal{J}} \left(\sigma_j(D)\alpha_j - u_j^\top d\right)^2 + \left\| \left(I_{N_i} - \sum_{j\in\mathcal{J}} u_j u_j^\top \right) d \right\|_2^2 .$$

The last term is independent of the coefficients. Each retained singular value is positive, so minimizing gives $\alpha_j = u_j^\top d/\sigma_j(D)$ and hence

$$z_{n+1} = \sum_{j\in\mathcal{J}} v_j \frac{u_j^\top d}{\sigma_j(D)}, \qquad c_{n+1} = c_{p,n+1} + V_{\text{eff},n+1} z_{n+1}. \tag{26}$$

Here τ_{LS} is the relative least-squares threshold. This is the minimum-norm least-squares solution restricted to retained right singular directions. Spatial weights enter once, and reconstruction requires no second solve. Algebraic variables such as pressure use their declared coupled blocks. The displayed time differences apply to differential field blocks.

C Experimental details

This appendix records problem definitions and evaluation protocols for the eight benchmarks in Section 3.

C.1 Common protocol

This subsection describes the shared settings and how we measure errors, runtime, and memory use across the eight cases.

Within each case, NOVA and the penalty, Nitsche, and Baumgarte baselines use the same fixed neural basis functions, problem data, collocation points, time grid, and independent test points. We test NOVA with refresh intervals $K \in \{1, 5, 10, 20\}$.

The same parameter grids apply across all eight cases. For the penalty method, we scan 23 boundary weights λ given by 10, 50, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, 10^5, 5×10^5, 10^6, 5×10^6, 10^7, 5×10^7, 10^8, 5×10^8, 10^9, 5×10^9, 10^{10}, 5×10^{10}, 10^{11}, 5×10^{11}, and 10^{12}.

For Baumgarte, the 16 damping values κ are 10, 50, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, 10^5, 2×10^5, 5×10^5, 10^6, 5×10^6, 10^7, 5×10^7, and 10^8. The 15 boundary row weights w are 10^{-3}, 10^{-2}, 0.1, 0.25, 1, 5, 10, 50, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, and 10^5. We scan all 240 combinations per case.

For Nitsche, the 15 safety factors s_N are 1.01, 1.1, 1.5, 2, 3, 5, 10, 50, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, and 10^5. Moving cylinder and Navier-Stokes additionally use the same 15 boundary row weights w, with all 225 combinations scanned in each case. Each of the other six cases uses the 15 safety factors alone.

Across all eight cases, these grids give 184 penalty, 1920 Baumgarte, and 540 Nitsche configurations, totaling 2644 trials, of which 2213 complete. Each case has 23 penalty and 240 Baumgarte trials. Nitsche has 225 trials each for Moving cylinder and Navier-Stokes and 15 for each remaining case.

Table 4 lists the boundary settings with the lowest E_{field} among completed scan candidates for each case and baseline. Ties are resolved by E_B and then run identifier. All metrics in a row come from that selected configuration. All 225 Moving cylinder Nitsche candidates failed.

The case tables show five representative parameter configurations per method, including those attaining the lowest E_{field} and E_B among completed runs. The full sweep results will be submitted as supplementary material.

Black **bold** and cyan mark the best and second-best displayed errors before rounding, ranked within each method for parameter sweeps and within each benchmark for the NOVA ablation. Alice Blue highlights NOVA at $K = 20$. $\star$ marks the Table 1 configuration.

Tables 1, 2, and 13 use boundary rank threshold $\tau = 10^{-12}$ and least-squares threshold $\tau_{\text{LS}} = 10^{-10}$. NOVA fixes $K = 20$ and $\tau_M = 10^{-13}$ across all eight cases in Table 1. For the Frozen baselines, shared initialization routines use compression tolerance 10^{-12} where applicable.

Table 4: Boundary parameters selected by minimum field error in each baseline scan for Table 1. λ is the penalty weight, s_N is Nitsche safety, κ is Baumgarte damping, and w is the boundary row weight. Moving cylinder Nitsche has no completed candidate.

Case	λ	s_N	Nitsche w	Baumgarte (κ, w)
Burgers	5×10^8	10^5	n/a	$(100, 50)$
Dynamic disk	10^{12}	10^5	n/a	$(10^4, 1)$
Moving cylinder	10^{12}	failed	n/a	$(10^8, 10^5)$
Navier-Stokes	10^{12}	5	50	$(500, 1)$
Topology change	10^{12}	10^5	n/a	$(10^5, 0.25)$
Stefan	100	2	n/a	$(5000, 500)$
Scriven bubble	10^{12}	10^5	n/a	$(10^8, 0.1)$
Scriven modal	1000	3	n/a	$(5\times10^6, 0.01)$

We report two error metrics, E_{field} and E_B. Let u_h and u^* denote the numerical and analytic fields, $x_{n,i}$ independent spatial test points, $w_{n,i}$ physical quadrature weights, and ω_n trapezoidal time weights. Define

$$A_n = \sum_i w_{n,i} \|u_h(x_{n,i}, t_n) - u^*(x_{n,i}, t_n)\|_2^2, \qquad C_n = \sum_i w_{n,i} \|u^*(x_{n,i}, t_n)\|_2^2. \tag{27}$$

The relative space time field error is

$$E_{\text{field}} = \left(\frac{\sum_n \omega_n A_n}{\sum_n \omega_n C_n} \right)^{1/2}. \tag{28}$$

For Moving cylinder and Navier-Stokes, the field is velocity. For modal Scriven, the field error compares the zero order temperature T_0 with its spherical analytic reference.

The boundary error is the root mean square residual on boundary test points $x^b_{n,j}$ with physical boundary weights $w^\Gamma_{n,j}$,

$$E_B = \left(\frac{\sum_{n,j} \omega_n w^\Gamma_{n,j} \|\mathcal{B}_{t_n}[u_h](x^b_{n,j}) - g(x^b_{n,j}, t_n)\|_2^2}{\sum_{n,j} \omega_n w^\Gamma_{n,j}} \right)^{1/2}. \tag{29}$$

Field evaluation uses independent test points. Boundary evaluation uses independent samples on curves and surfaces and the endpoints of one dimensional domains. Report times coincide with saved solution times. Spatial quadrature uses physical measures on every moving domain.

For free boundary problems, numerical and reference fields are compared on a common fixed region using their respective interfaces. Scriven temperatures use spherical volume weights and are extended by zero inside each bubble.

The PINN runs use seed zero and the final checkpoint, with five hidden layers of width 30, $\tanh$ activation, and float64 arithmetic. Adam uses 5000 epochs and L-BFGS uses 2500 outer iterations. Their runtime is measured on an NVIDIA A100-SXM4-80GB.

C.2 Analytic reference fields

Figures 4 and 5 show the analytic reference fields for the eight benchmarks described below.

C.3 Burgers with a time-dependent boundary trace

Figure 5(d) shows the traveling-wave reference on the fixed interval.

Reference construction The fixed spatial domain is $\Omega = (-1, 1)$. The scalar field satisfies

$$u_t + uu_x - \nu u_{xx} = 0, \qquad \nu = 0.2. \tag{30}$$

Let

$$\xi = x - ct - x_0, \qquad k = \frac{a}{2\nu}, \qquad u^*(x,t) = c - a\tanh(k\xi), \tag{31}$$

with $c = 0.5$, $a = 0.7$, and $x_0 = -0.3$. Direct differentiation gives

$$u_t^* + u^* u_x^* - \nu u_{xx}^* = 0. \tag{32}$$

The initial value and both endpoint values are traces of u^*. The geometry remains fixed while $g_L(t) = u^*(-1,t)$ and $g_R(t) = u^*(1,t)$ change with time.

Table 5: Representative baseline sweeps for Burgers. ⋆ marks the Table 1 configuration.

Penalty			**Nitsche**			**Baumgarte**		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	s_N	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	3.47e-03	3.37e-03	1.01	4.83e-05	1.24e-04	$(10, 10^{-3})$	9.66e-03	1.21e-02
5×10^6	2.26e-05	6.95e-09	3	2.59e-05	1.67e-05	$(100, 50)$ ⋆	**2.82e-06**	1.36e-06
5×10^8 ⋆	**2.25e-05**	7.01e-11	50	2.27e-05	7.81e-07	$(10^5, 50)$	3.02e-06	1.37e-09
10^{11}	2.34e-05	**3.78e-12**	5×10^3	2.26e-05	7.70e-09	$(10^8, 10^4)$	3.02e-06	**1.37e-12**
10^{12}	5.19e-05	3.85e-12	10^5 ⋆	**2.25e-05**	**3.85e-10**	$(10^8, 10^5)$	3.10e-06	1.37e-12

C.4 Dynamic disk with prescribed motion

Figure 4(a) shows the manufactured temperature field on a translating and pulsating disk.

Reference construction The moving domain is

$$\Omega(t) = \left\{ x \in \mathbb{R}^2 \;\middle|\; \|x - c_\Gamma(t)\|_2 < R(t) \right\}, \tag{33}$$

where

$$c_\Gamma(t) = \big(0.3\sin(2\pi t), 0\big)^\top, \qquad R(t) = 1 + 0.15\sin(2\pi t). \tag{34}$$

The heat equation and manufactured field are

$$u_t - \Delta u = f, \qquad u^*(x,y,t) = e^{-t/2}\sin(2\pi x)\sin(2\pi y). \tag{35}$$

Consequently

$$f^* = \left(8\pi^2 - \frac{1}{2}\right) u^*, \qquad g = u^*|_{\partial\Omega(t)}. \tag{36}$$

The prescribed disk acts as a moving computational window through the global reference field.

Table 6: Representative baseline sweeps for Dynamic disk. ⋆ marks the Table 1 configuration.

Penalty			**Nitsche**			**Baumgarte**		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	s_N	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	2.57e-01	2.45e-01	1.01	5.63e-02	2.06e-02	$(10, 10^{-3})$	9.65e+00	9.57e+00
10^4	3.79e-02	3.90e-04	3	3.95e-02	2.43e-03	$(10, 10^5)$	3.79e-03	6.83e-04
5×10^6	3.74e-02	8.41e-06	50	3.78e-02	1.81e-04	$(10^4, 1)$ ⋆	**3.52e-06**	1.19e-06
5×10^9	3.68e-02	3.31e-07	5×10^3	3.75e-02	1.41e-05	$(10^5, 50)$	2.44e-05	6.60e-08
10^{12} ⋆	**3.63e-02**	**1.48e-07**	10^5 ⋆	**3.72e-02**	**3.67e-06**	$(10^8, 10^5)$	2.40e-02	**2.80e-09**

C.5 Flow around a moving cylinder

Figure 4(b) shows the manufactured velocity around the prescribed moving cylinder.

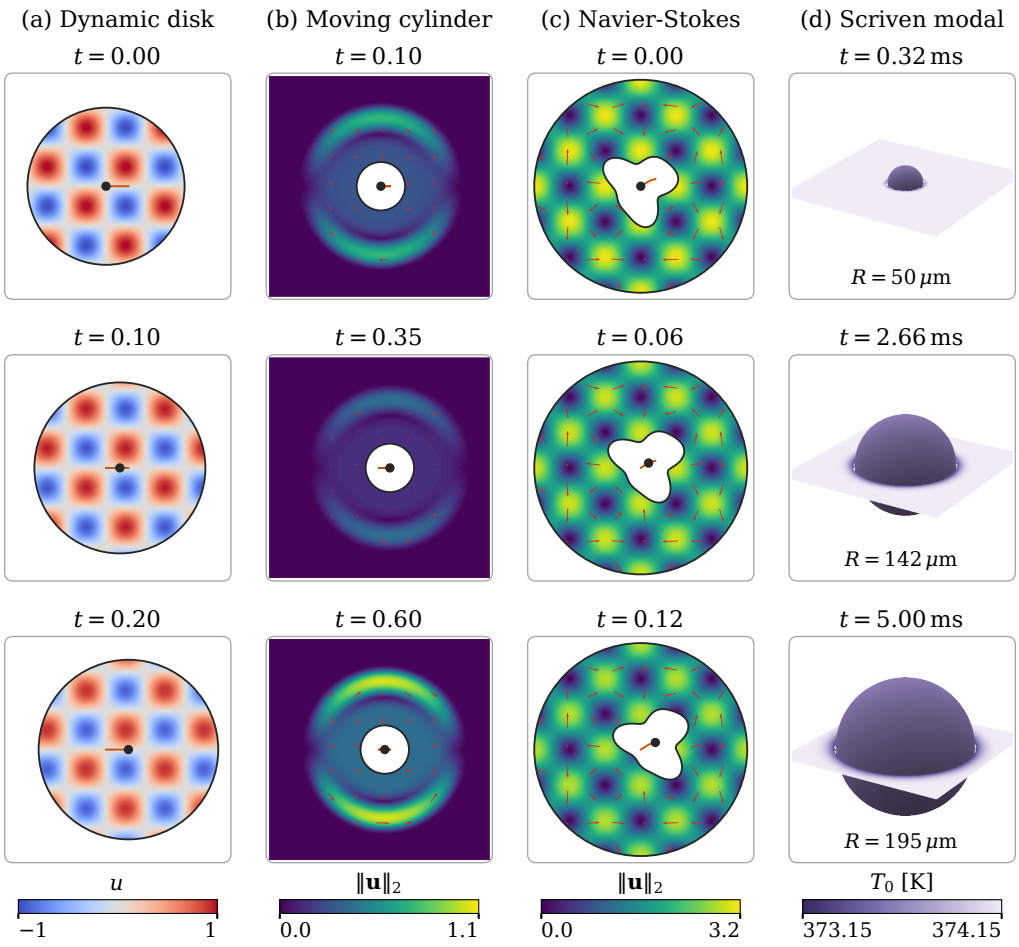


Figure 4: Analytic reference fields for (a) Dynamic disk, (b) Moving cylinder, (c) Navier-Stokes, and (d) Scriven modal. Each column shows one case at three times increasing from top to bottom. Spatial limits and color scales are fixed within each column. Red arrows show velocity with a common scale across times. Dark dots and orange paths in (a) to (c) mark the current centers and their prescribed trajectories. Column (d) shows the analytic base-mode temperature T_0 on a complete equatorial liquid plane and the spherical radius $a_0 = R$. Sphere shading represents geometry and lighting, while its color bar describes liquid temperature.

Reference construction The periodic box is $[0, 2]^2$. A cylinder of radius $a = 0.22$ follows

$$c(t) = \big(1 + 0.05(1 - \cos \omega t), 1\big)^\top, \qquad \omega = \frac{2\pi}{0.8}. \tag{37}$$

Let $r = x - c(t)$ and let $V(t) = \dot{c}(t)$. A compact C^3 radial cutoff q_u equals one near the cylinder. The streamfunction and velocity are

$$\Psi_{\mathrm{rb}} = V_x r_y - V_y r_x, \qquad \Psi^* = q_u \Psi_{\mathrm{rb}}, \qquad u^* = \nabla^\perp \Psi^*. \tag{38}$$

This construction gives $\nabla \cdot u^* = 0$ in the box and $u^* = V$ on the cylinder. The target drag is $D^*(t) = \sin(\omega t)$. A second cutoff defines the exterior pressure

$$p^*_{\mathrm{out}} = g(t) r_x q_p, \qquad g(t) = -\frac{D^*(t)}{\pi a^2}. \tag{39}$$

The pressure inside the cylinder is zero. The regular body force follows by substitution into the incompressible Navier-Stokes equation. The pressure jump supplies the prescribed interface traction in the distributional formulation. The cylinder trajectory remains externally prescribed.

Table 7: Representative baseline sweeps for Moving cylinder. ⋆ marks the Table 1 configuration.

Penalty			Nitsche			Baumgarte		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(s_N, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	8.47e-01	3.83e-01	$(1.01, 10^{-3})$	failed	failed	$(10, 10^{-3})$	1.46e+00	1.05e+00
10^4	4.03e-01	6.42e-04	$(1.01, 10^5)$	failed	failed	$(10, 10^5)$	8.26e-02	8.56e-03
5×10^6	3.97e-01	1.16e-05	$(50, 50)$	failed	failed	$(10^5, 50)$	8.43e-02	1.08e-06
5×10^9	3.89e-01	2.89e-07	$(10^5, 10^{-3})$	failed	failed	$(10^8, 10^{-3})$	8.62e-02	1.73e-04
10^{12} ⋆	**3.84e-01**	**2.10e-08**	$(10^5, 10^5)$	failed	failed	$(10^8, 10^5)$ ⋆	**8.19e-02**	**1.07e-09**

C.6 Navier-Stokes flow on a prescribed petal

Figure 4(c) shows the manufactured velocity around the prescribed moving petal.

Reference construction The fluid region is a disk of radius 1.5 with a prescribed moving petal removed. Its center and angle are

$$c_b(t) = \big(0.30\sin(2\pi t), 0.10\sin(4\pi t)\big)^\top, \qquad \alpha(t) = \pi t. \tag{40}$$

The petal radius in body coordinates is

$$r(\vartheta) = 0.45\big(1 + 0.25\sin(3\vartheta) + 0.1\cos(5\vartheta)\big). \tag{41}$$

Define

$$\psi^*(x, y, t) = e^{-t/2}\sin(\pi x)\sin(\pi y), \qquad u^* = \nabla^\perp \psi^*, \tag{42}$$

and

$$p^*(x, y, t) = e^{-t/2}\cos(\pi x)\cos(\pi y). \tag{43}$$

The velocity is divergence free. The manufactured forcing is

$$f^* = \partial_t u^* + (u^* \cdot \nabla)u^* - \nu\Delta u^* + \nabla p^*, \qquad \nu = 0.5. \tag{44}$$

The analytic fields are defined independently of the prescribed petal motion. Both boundary components use the trace of u^*.

Table 8: Representative baseline sweeps for Navier-Stokes. ⋆ marks the Table 1 configuration.

Penalty			Nitsche			Baumgarte		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(s_N, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	6.37e-02	2.72e-01	$(1.01, 10^{-3})$	8.70e-03	2.29e-02	$(10, 10^{-3})$	6.79e-01	3.17e+00
10^4	1.73e-02	1.05e-03	$(5, 50)$ ⋆	**7.28e-03**	8.84e-03	$(500, 1)$ ⋆	**2.77e-05**	4.13e-05
5×10^6	1.51e-02	3.66e-05	$(10, 5\times10^4)$	7.29e-03	**7.46e-03**	$(10^5, 50)$	4.24e-04	1.08e-06
5×10^9	1.15e-02	1.21e-06	$(50, 50)$	1.58e+01	2.78e+01	$(10^8, 10^3)$	4.56e-03	**3.17e-07**
10^{12} ⋆	**6.33e-03**	**1.21e-07**	$(10^5, 10^5)$	failed	failed	$(10^8, 10^5)$	4.06e-03	6.10e-07

C.7 Topology change through prescribed Cassini splitting

Figure 5(a) shows the manufactured heat field before and after the prescribed Cassini domain split.

Reference construction The moving region is

$$\Omega(t) = \big\{(x, y) \;\big|\; \big((x-a)^2 + y^2\big)\big((x+a)^2 + y^2\big) < b(t)^4\big\}, \tag{45}$$

with $a = 0.5$ and $b(t) = 0.66 - 1.55t$. The region has one component for $b > a$ and two components for $b < a$. The exact event time is

$$t^* = \frac{0.66 - 0.5}{1.55}. \tag{46}$$

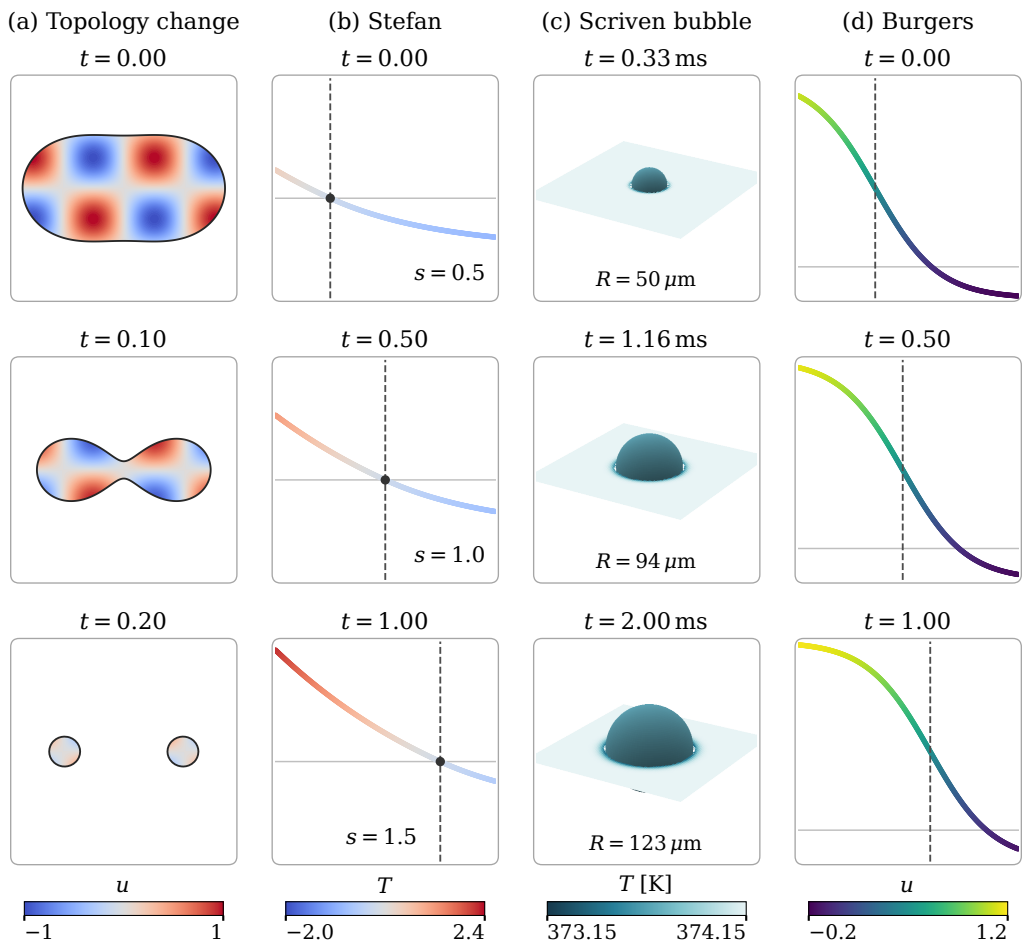


Figure 5: Analytic references for (a) Topology change, (b) Stefan, (c) Scriven bubble, and (d) Burgers. Time increases from top to bottom, with fixed spatial limits and color scales within each column. In (b), the dashed line and dot mark the analytic interface. The oblique view in (c) combines the analytic sphere with a complete equatorial liquid temperature plane. Sphere shading represents geometry and lighting, while its color bar describes liquid temperature. The dashed line in (d) marks the traveling-wave center.

The geometry routine keeps $|b - a| \geq 0.002$ and switches between the two regular sides of the singular lemniscate. The heat reference is

$$u^*(x, y, t) = e^{-t/2} \sin(2\pi x) \sin(2\pi y), \qquad f^* = \left(8\pi^2 - \frac{1}{2}\right) u^*. \tag{47}$$

The boundary trace and initial value come from this field. The topology schedule is prescribed.

Table 9: Representative baseline sweeps for Topology change. ⋆ marks the Table 1 configuration.

Penalty			Nitsche			Baumgarte		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	s_N	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	1.60e-01	1.29e-01	1.01	3.36e-02	1.10e-02	$(10, 10^{-3})$	6.86e+02	5.91e+02
10^4	2.24e-02	2.96e-04	3	2.37e-02	1.18e-03	$(10, 10^5)$	3.75e-02	1.85e-02
5×10^6	2.13e-02	8.71e-06	50	2.21e-02	1.10e-04	$(10^5, 0.25)$ ⋆	**3.36e-06**	1.64e-06
5×10^9	2.04e-02	4.07e-06	5×10^3	2.13e-02	9.48e-06	$(10^5, 50)$	2.09e-05	4.84e-07
10^{12} ⋆	**1.98e-02**	**4.07e-06**	10^5 ⋆	**2.09e-02**	**4.52e-06**	$(10^8, 10^5)$	2.57e-02	**6.32e-09**

C.8 Stefan with a two-phase free boundary

Figure 5(b) shows the analytic temperatures and interface for the two phase Stefan problem.

Reference construction An unknown interface $s(t)$ divides $[0, 2]$ into two phases. The temperature fields satisfy

$$u_{1,t} = 2u_{1,xx} \qquad 0 < x < s(t), \tag{48}$$

$$u_{2,t} = u_{2,xx} \qquad s(t) < x < 2. \tag{49}$$

The interfacial temperature and Stefan law are

$$u_1(s(t), t) = u_2(s(t), t) = 0, \qquad \dot{s}(t) = -2u_{1,x}(s(t)^-, t) + u_{2,x}(s(t)^+, t). \tag{50}$$

The exact coupled solution is

$$u_1^*(x, t) = 2\exp\left(\frac{2t - 2x + 1}{4}\right) - 2, \tag{51}$$

$$u_2^*(x, t) = \exp\left(\frac{2t - 2x + 1}{2}\right) - 1, \tag{52}$$

$$s^*(t) = t + \frac{1}{2}. \tag{53}$$

Both temperatures vanish at $s^*(t)$. Their interface gradients give $\dot{s}^* = 1$. The solver receives $s(0) = 0.5$ and computes subsequent positions from the numerical temperature gradients. This reference follows Stefan Case 3 used by PIELM (Ren et al., 2025).

Table 10: Representative baseline sweeps for Stefan. ⋆ marks the Table 1 configuration.

Penalty			Nitsche			Baumgarte		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	s_N	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	1.94e-01	1.69e-01	1.01	1.45e-01	5.87e-03	$(10, 10^{-3})$	2.20e-03	2.36e-03
100 ⋆	**3.54e-02**	2.01e-02	2 ⋆	**1.06e-02**	1.39e-03	$(5\times10^3, 500)$ ⋆	**5.83e-06**	5.75e-07
5×10^6	3.90e-02	4.15e-07	3	2.32e-02	7.87e-04	$(10^5, 50)$	6.10e-06	2.87e-08
5×10^{11}	3.90e-02	**2.30e-11**	50	3.83e-02	3.68e-05	$(10^8, 5\times10^4)$	2.62e-05	**2.87e-11**
10^{12}	3.90e-02	2.58e-11	10^5	3.90e-02	**1.82e-08**	$(10^8, 10^5)$	1.05e-04	2.87e-11

C.9 Spherical Scriven bubble growth

Figure 5(c) shows the analytic temperature and radius for spherical Scriven growth.

Reference construction The liquid occupies $R(t) < r < R_\infty$. Define the normalized superheat

$$\theta = \frac{T - T_{\text{sat}}}{T_\infty - T_{\text{sat}}}. \tag{54}$$

The spherical transport and Stefan equations are

$$\theta_t + u_r\theta_r = \alpha_l\left(\theta_{rr} + \frac{2}{r}\theta_r\right), \tag{55}$$

$$u_r(r,t) = q\dot{R}\frac{R^2}{r^2}, \tag{56}$$

$$\rho_v L_{\text{eff}}\dot{R} = \lambda_l(T_\infty - T_{\text{sat}})\theta_r(R^+,t), \tag{57}$$

where $q = 1 - \rho_v/\rho_l$. The Scriven radius is

$$R^*(t) = 2\beta_S\sqrt{\alpha_l t}. \tag{58}$$

The constant β_S is the positive root of

$$2\beta_S^2\int_0^1 \exp\left\{-\beta_S^2\left[(1-\eta)^{-2} - 2q\eta - 1\right]\right\}\,\mathrm{d}\eta = \frac{\rho_l c_{p,l}(T_\infty - T_{\text{sat}})}{\rho_v L_{\text{eff}}}. \tag{59}$$

The current properties give $\beta_S = 3.3820039708919154$. Define

$$A = \frac{2\beta_S^2\rho_v L_{\text{eff}}}{\rho_l c_{p,l}(T_\infty - T_{\text{sat}})}, \qquad \eta = 1 - \frac{R^*(t)}{r}. \tag{60}$$

The reference temperature is

$$\theta^*(r,t) = 1 - A\int_\eta^1 \exp\left\{-\beta_S^2\left[(1-\zeta)^{-2} - 2q\zeta - 1\right]\right\}\,\mathrm{d}\zeta. \tag{61}$$

The integral equation gives $\theta^*(R^+,t) = 0$ and makes the temperature gradient consistent with $\dot{R}^*$. This is the spherical Scriven solution (Scriven, 1995). The finite numerical shell uses $R_\infty = 500$ micrometers. The analytic far field gradient is negligible at this radius for the reported time interval.

The liquid properties are $\lambda_l = 0.67135$, $\rho_l = 958.4$, and $c_{p,l} = 4216$. The vapor density is $\rho_v = 0.585$ and $c_{p,v} = 2030$. The latent heat is 2.26×10^6. The temperatures are $T_{\text{sat}} = 373.15$ K and $T_\infty = 374.15$ K.

Table 11: Representative baseline sweeps for Scriven bubble. ⋆ marks the Table 1 configuration.

Penalty			**Nitsche**			**Baumgarte**		
λ	$E_{\text{field}}\downarrow$	$E_B\downarrow$	s_N	$E_{\text{field}}\downarrow$	$E_B\downarrow$	(κ, w)	$E_{\text{field}}\downarrow$	$E_B\downarrow$
10	7.72e-02	6.40e-01	1.01	7.77e-04	1.70e-04	$(10, 10^{-3})$	failed	failed
10^4	8.53e-04	2.64e-03	3	4.72e-05	4.60e-06	$(10^5, 50)$	3.41e-05	9.96e-05
5×10^6	3.45e-05	5.29e-06	50	3.35e-05	1.92e-07	$(10^8, 0.1)^\star$	**4.77e-06**	1.87e-07
5×10^9	3.29e-05	5.29e-09	5×10^3	3.29e-05	1.88e-09	$(10^8, 10^3)$	1.83e-04	**9.12e-08**
$10^{12}{}^\star$	**3.29e-05**	**2.71e-11**	$10^5{}^\star$	**3.29e-05**	**9.41e-11**	$(10^8, 10^5)$	3.61e-03	2.05e-07

C.10 Axisymmetric modal Scriven bubble

The axisymmetric model retains Legendre modes through order two. Figure 4(d) shows its analytic base temperature T_0 and spherical radius $a_0 = R$.

Physical model Let $y = r - R(t)$ be the radial coordinate measured from the mean interface and let $\mu = \cos\vartheta$. The interface and normalized liquid temperature are

$$r_s(\vartheta,t) = R(t) + \sum_{\ell=1}^{2} a_\ell(t)P_\ell(\mu), \tag{62}$$

$$\Theta(y,\vartheta,t) = T_0(y,t) + \sum_{\ell=1}^{2} T_\ell(y,t)P_\ell(\mu). \tag{63}$$

The initial perturbation excites only the second mode with $a_2(t_0)/R_0 = 0.005$ and $a_1(t_0) = 0$. The pole and equator radii are $R + a_2$ and $R - a_2/2$.

The zero order temperature satisfies

$$\partial_t T_0|_y = \alpha_l \left(T_0'' + \frac{2}{r} T_0' \right) - (U - \dot{R}) T_0', \tag{64}$$

$$U(r,t) = \left(1 - \frac{\rho_v}{\rho_l} \right) \frac{R^2 \dot{R}}{r^2}, \tag{65}$$

$$T_0(0,t) = 0, \qquad T_0(\infty,t) = 1, \tag{66}$$

$$\dot{R} = k_{\mathrm{st}} T_0'(0,t), \qquad k_{\mathrm{st}} = \frac{\lambda_l (T_\infty - T_{\mathrm{sat}})}{\rho_v L_{\mathrm{eff}}}. \tag{67}$$

This block determines the mean free boundary. The solver receives $R(t_0) = R_0$ and computes every later radius from the numerical interface gradient.

For $q_\ell = \ell(\ell+1)$, the retained high order thermal equation is

$$\partial_t T_\ell|_y = \alpha_l \left(T_\ell'' + \frac{2}{r} T_\ell' - \frac{q_\ell}{r^2} T_\ell \right) - (U - \dot{R}) T_\ell' - \frac{q_\ell}{r^2} \psi_\ell T_0'. \tag{68}$$

The linearized isothermal and Stefan conditions are

$$T_\ell(0,t) + a_\ell(t) T_0'(0,t) = 0, \qquad T_\ell(\infty,t) = 0, \qquad \dot{a}_\ell = k_{\mathrm{st}} \left[T_\ell'(0,t) + a_\ell T_0''(0,t) \right]. \tag{69}$$

The unknown coefficient vector for T_ℓ^{n+1} and the unknown a_ℓ^{n+1} enter one linear system. The next modal interface is therefore obtained from the numerical thermal field.

The streamfunction and vorticity variables satisfy

$$g_\ell = \mathcal{D}_\ell \psi_\ell, \qquad \mathcal{D}_\ell = \partial_{yy} - \frac{q_\ell}{r^2}, \tag{70}$$

$$\partial_t g_\ell = \nu_l \mathcal{D}_\ell g_\ell + (\dot{R} - U) g_\ell' + \frac{2U}{r} g_\ell. \tag{71}$$

The interface rows impose tangential stress, modal normal velocity, and the relation between g_ℓ and ψ_ℓ. The far field rows impose $\psi_\ell = g_\ell = 0$. Equation 68 receives the resulting flow through its last term.

Analytic reference and coupling scope The zero order reference uses the Scriven solution from Appendix C.9. The case time is mapped to the compatible similarity time by

$$t_{\mathrm{eff}} = t_{R_0} + t - t_0, \qquad t_{R_0} = \frac{R_0^2}{4\beta_S^2 \alpha_l}. \tag{72}$$

Hence

$$R^*(t) = 2\beta_S \sqrt{\alpha_l t_{\mathrm{eff}}} \tag{73}$$

passes exactly through the prescribed initial radius. The temperature T_0^* follows the integral expression in Appendix C.9 at t_{eff}. This reference provides the initial T_0 state and the evaluation values R^* and T_0^*. Subsequent states follow the coupled numerical equations.

The first order perturbation model has one way modal coupling, with (R, T_0) driving $(a_\ell, T_\ell, \psi_\ell, g_\ell)$. The spherical Scriven solution provides the analytic reference for R and T_0. Higher-order shape, temperature, and flow modes are computed from Equations 68–71.

Compatible initial state The high order temperature starts from the first order displacement of the base profile

$$T_\ell(y, t_0) = -a_\ell(t_0) T_0'(y, t_0). \tag{74}$$

The initial projection enforces the interface trace, far field decay, and

$$T_\ell'(0, t_0) = -a_\ell(t_0) T_0''(0, t_0). \tag{75}$$

Equation 69 then gives $\dot{a}_\ell(t_0) = 0$. The initial flow is the minimum norm frozen feature state that satisfies Equation 71 and all five instantaneous flow traces. All four NOVA settings and all three Frozen-PINN variants use this same initial state.

The infinite liquid region is compactified by

$$\xi = \frac{y}{L+y}, \qquad y = \frac{L\xi}{1-\xi}, \qquad L = 25 \text{ micrometers}. \tag{76}$$

Table 12: Representative baseline sweeps for Scriven modal. ⋆ marks the Table 1 configuration.

Penalty			**Nitsche**			**Baumgarte**		
λ	$E_{\text{field}} \downarrow$	$E_B \downarrow$	s_N	$E_{\text{field}} \downarrow$	$E_B \downarrow$	(κ, w)	$E_{\text{field}} \downarrow$	$E_B \downarrow$
10	failed	failed	1.01	5.94e-06	8.80e-13	$(10, 10^{-3})$	1.30e-02	5.98e-02
10^3 ⋆	**9.60e-01**	2.25e+00	3 ⋆	**5.94e-06**	9.05e-13	$(10, 10^5)$	1.00e-03	1.14e-05
10^6	9.88e-01	**2.45e-03**	50	1.01e-04	2.12e-13	$(10^5, 50)$	1.58e-06	3.31e-08
5×10^6	9.88e-01	1.17e+00	10^4	1.48e-04	**9.76e-14**	$(5\times10^6, 0.01)$ ⋆	**1.54e-07**	3.70e-07
10^{12}	failed	failed	10^5	1.51e-04	9.84e-14	$(10^8, 10^5)$	1.00e-03	**3.13e-11**

C.11 Complete nullspace refresh interval ablation

Table 13 reports all eight benchmarks with $K = 1, 5, 10, 20$. Within each benchmark, features, collocation points, time integration settings, spectral thresholds, independent test points, and the random seed remain fixed. The two panels in the main-text Table 2 reproduce the corresponding entries from this complete table.

Table 13: **NOVA refresh interval ablation across eight benchmarks.** We compare $K \in \{1, 5, 10, 20\}$ using the same metrics as Table 1.

Burgers	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)	**Dynamic disk**	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	**3.02e-06**	1.98e-16	2.40	101.14	NOVA ($K = 1$)	8.35e-07	**1.66e-10**	23.76	251.19
NOVA ($K = 5$)	3.02e-06	**1.78e-16**	1.64	101.36	NOVA ($K = 5$)	8.31e-07	1.73e-10	14.14	246.34
NOVA ($K = 10$)	3.02e-06	1.85e-16	1.55	101.17	NOVA ($K = 10$)	8.27e-07	1.89e-10	12.59	230.59
NOVA ($K = 20$)	3.02e-06	1.94e-16	1.50	101.15	NOVA ($K = 20$)	**8.26e-07**	2.09e-10	11.78	225.79
Moving cylinder	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)	**Navier-Stokes**	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	7.93e-02	1.14e-15	32.62	300.60	NOVA ($K = 1$)	**7.92e-06**	3.91e-08	8.12	279.68
NOVA ($K = 5$)	7.93e-02	1.13e-15	28.97	313.23	NOVA ($K = 5$)	7.97e-06	**3.52e-08**	7.92	273.71
NOVA ($K = 10$)	**7.93e-02**	1.15e-15	28.65	284.32	NOVA ($K = 10$)	7.95e-06	3.59e-08	7.42	263.81
NOVA ($K = 20$)	7.93e-02	**1.13e-15**	28.04	295.62	NOVA ($K = 20$)	7.94e-06	3.57e-08	7.83	268.83
Topology change	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)	**Stefan**	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	2.27e-06	1.18e-10	2.86	176.71	NOVA ($K = 1$)	3.22e-05	5.74e-16	1.26	94.50
NOVA ($K = 5$)	**2.24e-06**	1.01e-10	2.31	174.46	NOVA ($K = 5$)	3.22e-05	3.39e-16	0.84	94.14
NOVA ($K = 10$)	2.33e-06	**9.09e-11**	2.28	174.74	NOVA ($K = 10$)	3.22e-05	3.39e-16	0.79	94.36
NOVA ($K = 20$)	2.36e-06	1.85e-08	2.29	183.00	NOVA ($K = 20$)	**3.22e-05**	**3.05e-16**	0.76	94.08
Scriven bubble	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)	**Scriven modal**	$E_{\text{field}} \downarrow$	$E_B \downarrow$	Time (s)	RSS (MiB)
NOVA ($K = 1$)	**6.37e-06**	9.87e-13	16.47	119.30	NOVA ($K = 1$)	**2.54e-07**	5.11e-15	53.41	209.54
NOVA ($K = 5$)	6.37e-06	**5.67e-13**	8.62	116.87	NOVA ($K = 5$)	2.54e-07	3.88e-15	38.09	212.95
NOVA ($K = 10$)	6.37e-06	7.36e-13	7.83	118.38	NOVA ($K = 10$)	2.54e-07	3.71e-15	36.07	210.87
NOVA ($K = 20$)	6.37e-06	6.33e-13	7.31	116.14	NOVA ($K = 20$)	2.54e-07	**3.51e-15**	35.26	210.64

C.12 Boundary-count ablation configuration

The Navier-Stokes boundary-count ablation in Table 3 varies only the boundary point count N_b, distinct from the scalar constraint count n_b. Each point contributes two velocity constraints, giving $n_b = 2N_b$. All runs use NOVA with $K = 1$, $\tau = \tau_M = 10^{-12}$, $\tau_{\text{LS}} = 10^{-10}$, and seed zero. The fixed inputs include 240 velocity features, 80 pressure features, feature scale 0.8, the same 700 interior points at each step, and time step 0.00375. The rank scan uses the 240 velocity columns. The independent field and boundary evaluations use 4096 and 1024 points, respectively, with identical point sets, weights, and reference values across configurations.

The rank scan evaluates 15 boundary counts at all 33 saved times with $\tau = 10^{-12}$. Independent boundary sets of 2176 and 4352 points provide dense reference ranks of 159 to 160, estimating the trace dimension resolved at this threshold. At $t = 0$, the rank matches the dense reference value 159 from $N_b = 102$ onward. At this count, three other saved times retain one more direction than the dense references. For all tested counts $N_b \geq 136$, the sampled ranks match the dense references at every saved time. Independent field and boundary errors also reach a plateau. For these tested times and settings, the scan supports $N_b = 136$ as a conservative sampling choice. The main experiment uses 272 boundary points.

D Limitations and Future Work

The current implementation performs full nullspace rebuilds at a fixed interval of K time steps. Future work will explore an adaptive refresh strategy that uses boundary compatibility residuals to adjust the rebuild frequency as the boundary evolves.