

Field-Residual Learning Does Not Reliably Improve a Strong Fixed-Budget Riemann Starter, Even Under Hard-Aware Checkpoint Selection

Ramesh Kolluru

rameshkolluru43@gmail.com

Research Article

Keywords: Riemann problems, approximate Riemann solvers, star-state estimation, physicsinformed residual learning, checkpoint selection, one-dimensional Euler flow.

Posted Date: October 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11148885/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Field-Residual Learning Does Not Reliably Improve a Strong Fixed-Budget Riemann Starter, Even Under Hard-Aware Checkpoint Selection

Ramesh Kolluru (Corresponding author)
Senior CFD Consultant and Independent Researcher, India
rameshkolluru43@gmail.com
ORCID: [0000-0001-5447-8550](https://orcid.org/0000-0001-5447-8550)

24 September 2026

Abstract

Approximate Riemann star-state estimates enter many Godunov-type constructions. This paper tests whether a two-scalar learned correction, $(\Delta \log p_*, \Delta u_*)$, improves a strong fixed-budget estimate of the star pressure and velocity for one-dimensional ideal-gas Euler Riemann problems. The correction is applied around Toro’s branched starter Q_0 ; an analytic three-wave sampler reconstructs the field, and zero correction recovers Q_0 exactly. The same residual improves a weak primitive-variable Riemann solver (PVRS) base on the predefined hard-anchor metric (hard median $0.203 \rightarrow 0.045 \pm 0.007$, five seeds). Under the locked overall-median protocol of the objective-comparison study, none of four tested field- and physics-based objectives yields a dependable hard-anchor improvement over Q_0 . A confirmatory ten-seed study evaluates validation at every epoch and applies two predeclared rules to the same trajectories. The overall-median rule yields no hard-anchor gains. The composite rule reduces average degradation and identifies two practically improved seeds, but eight of ten selected models remain worse than Q_0 and the bootstrap interval for the mean paired difference includes zero. In a diagnostic first-order Godunov-type finite-volume embedding on 100–400 cells, the maximum inter-method difference in the reported median errors is at most 10^{-3} . The conclusions concern this two-scalar residual, the specified one-dimensional Euler family suite, the tested objectives, and the two checkpoint-selection rules; they do not establish a new numerical flux.

Keywords Riemann problems; approximate Riemann solvers; star-state estimation; physics-informed residual learning; checkpoint selection; one-dimensional Euler flow.

1 Introduction

Approximate Riemann solvers remain the local engine of many Godunov-type finite-volume and discontinuous-Galerkin schemes [1–5]. A learned correction to an approximate star state is an attractive hybrid-design target: the learned map may adjust a difficult local estimate while a known wave model retains its physical structure. This paper is not a new finite-volume flux or a general-purpose learned Riemann solver. Instead, it isolates the narrow question of whether field-residual learning can improve a matched, fixed-budget star-state starter under a controlled one-dimensional Euler protocol.

A residual around a weak starter can appear effective while adding no value beyond a stronger physical approximation. We therefore compare every learned correction with Toro’s branched fixed-budget starter Q_0 , using the same input states, wave reconstruction, branch budget, and held-out protocol. The learned map returns only (p_*, u_*) ; the analytic sampler, not a learned field operator, creates the spatial field. The objective is therefore a diagnostic of learned

star-state correction, rather than a general assessment of neural operators or physics-informed learning [6, 7].

The study tests three hypotheses. First, the residual parameterization has capacity: the same architecture can reduce the error of a weak primitive-variable Riemann solver (PVRS) starter. Second, under the locked overall-median selection protocol used for the objective-comparison experiment, none of the four tested objectives yielded a dependable hard-anchor improvement over Q_0 (Sections 6.5 and 6.6). For the field-residual objective, the confirmatory full-epoch composite-selection study reduced degradation but did not produce repeatable improvement (Section 5.3). Third, aggregate soft-case metrics and ordinary checkpoint selection can conceal hard-regime degradation. The field losses follow residual constructions used in PINN and PINO studies, including conservative control-volume and thermodynamically consistent variants [8, 9]. Weak, dual-norm, and locally linearized objectives for discontinuous hyperbolic problems [10–13] are related but are not the objects under test.

The contributions are:

1. A matched-baseline protocol for learned star-state corrections: the same sampler, the same six-pass budget, a residual that is identically Q_0 at initialization, hard-family stratification, and paired multi-seed reporting.
2. Evidence that weak-starter errors are learnable, while under the locked overall-median selection protocol used for the objective-comparison experiment none of the four tested objectives yielded a dependable hard-anchor improvement over Q_0 .
3. A confirmatory checkpoint-selection study (Section 5.3) showing that the main negative result is not caused by sparse validation cadence or selection solely by overall validation median. A predeclared hard-aware composite rule reduces average degradation but does not yield a repeatable practical improvement over Q_0 .
4. A label- and front-aware evaluation showing that the Q_0 -versus-learned ordering is preserved under conserved and plateau metrics.
5. A first-order Godunov-type finite-volume embedding that varies only the star map and shows that, at the tested first-order resolutions, donor-cell truncation masks the differences among the evaluated star maps (Section 5.4).

The conclusions concern this two-scalar residual, the specified one-dimensional ideal-gas Euler family suite, the tested objectives, and the two predeclared checkpoint-selection rules. They do not establish improved finite-volume flux performance, multidimensional interface treatment, non-ideal-EOS performance, or general learned operator performance. A first-order Godunov-type finite-volume embedding of the same stars is reported only as a diagnostic in Section 5.4.

2 Problem

2.1 Ideal-gas Euler IBVP

On a tube $\Omega = [0, L]$ the conserved state $\mathbf{U} = (\rho, \rho u, E)$ satisfies

$$\partial_t \mathbf{U} + \partial_x \mathbf{F}(\mathbf{U}) = 0, \quad \mathbf{F} = (\rho u, \rho u^2 + p, u(E + p)), \quad (1)$$

with caloric equation of state $E = p/(\gamma - 1) + \frac{1}{2}\rho u^2$ [5]. Each training and holdout member is a Riemann IBVP: a diaphragm at x_d separates constant states (ρ_L, u_L, p_L) and (ρ_R, u_R, p_R) . The spatial ends are *fixed Dirichlet* data—the IC endpoints—for all t . The horizon T is clipped so that leading waves stay strictly inside $(0, L)$ (inactive T). Exact labels are Toro’s similarity solution sampled on the same grid; they are consistent with those Dirichlet ends.

The structured surrogate is

$$Q_\theta : (\rho_0, u_0, p_0, \gamma; x, t) \mapsto (\rho, u, p)(x, t). \quad (2)$$

Weights are shared. There is no case-specific residual δ . Hard IC and Dirichlet overwrite only values the IBVP already owns:

$$(\rho, u, p)|_{t=0} = (\rho_0, u_0, p_0), \quad (\rho, u, p)|_{x \in \{0, L\}} = (\rho_0, u_0, p_0)|_{x \in \{0, L\}}. \quad (3)$$

2.2 Nine-regime taxonomy

Classic anchors follow the RICCA / Toro table [5, 14] with $\gamma = 1.4$ (Table 1). They are held out of training. Soft families are $\{1, 2, 6, 7, 8, 9\}$; hard families are the two blasts and the collision, $\{3, 4, 5\}$. Locked validation, in-family, and detuned-family members are random draws from those same nine families, never the named anchors. An in-family (protocol-v2 IID / val) draw jitters each primitive about the family anchor: densities by a relative half-width $\delta_\rho \in [0.05, 0.15]$, velocities by an absolute/fractional half-width $\delta_u \in [0.02, 0.15]$, pressures either relatively or in $\log_{10} p$ ($\delta_{\log p} = 0.15$ on the blasts), plus diaphragm and horizon jitter. A detuned-family draw uses the same boxes with every half-width multiplied by three (“3 $\times$ detuning”): blast $\log_{10} p$ half-width 0.45, diaphragm clipped to $[0.2, 0.8]$, and the stationary-shock Mach box widened from $[1.5, 2.5]$ to $[1.2, 3.5]$. The locked split file stores that bucket under an internal name; in this paper the draw is called detuned-family throughout. It is a limited, family-conditioned perturbation shift, not a qualitatively new distribution or PDE class. Members remain 1D ideal-gas Riemann problems. The locked protocol-v2 file stores one validation member per family ($n = 9$, of which 3 hard), two in-family members per family ($n = 18$, 6 hard), and two detuned-family members per family ($n = 18$, 6 hard). Density and pressure are jittered independently. There is no physical-admissibility rejection and therefore no rejection rate by family. Positivity is a post-jitter clamp ($\rho \geq 10^{-8}$; contact pressure $\geq 10^{-6}$). Blast families force a pressure ratio of at least 100 after the independent $\log_{10} p$ draws. A later 1800-member Monte Carlo audit of the same sampler (Section 6.8 and Supplementary Section 2.3) found no draw with $p < 10^{-3}$ or $\rho < 10^{-2}$; that 1800-member campaign does not cover near-vacuum initial data [3, 5], although a Toro-123 star can still be near vacuum. A separate locked near-vacuum family is studied in Section 6.6. It is a control at the Q_0 accuracy floor, not part of the published hard-error metric. The published hard set remains $\{3, 4, 5\}$; family 10 is never folded into that median.

Table 1 Held-out classic Riemann anchors ($\gamma = 1.4$, $x_d = 0.5$). Inactive T clips the listed Toro time so waves stay off the Dirichlet ends. Case 6 is the Mach-2 stationary normal shock. ID 10 is the locked near-vacuum control of Section 6.6. It is outside the nine paper regimes and is not included in the hard-error metric.

ID	Name	ρ_L	u_L	p_L	ρ_R	u_R	p_R
1	Sod	1	0	1	0.125	0	0.1
2	Toro 123	1	−2	0.4	1	2	0.4
3	Left blast	1	0	1000	1	0	0.01
4	Right blast	1	0	0.01	1	0	100
5	Collision	5.999	19.60	460.9	5.992	−6.196	46.10
6	Stat. shock	1	1	0.179	2.667	0.375	0.804
7	Stat. contact	1.4	0	0.4	1	0	0.4
8	Slow contact	1.4	0.1	1	1	0.1	1
9	Slow shock	3.86	−0.81	10.33	1	−3.44	1
10	Near vacuum	1	−3	0.4	1	3	0.4

2.3 Metrics

Reported field error on a primitive q is

$$\text{Rel-}L^2(q) = \|\hat{q} - q^*\|_2 / (\|q^*\|_2 + \varepsilon), \quad \varepsilon = \max(10^{-6}, 10^{-3} \text{rms}(q^*)\sqrt{N}). \quad (4)$$

The headline score is the mean over $\{\rho, u, p\}$, then the *median* over the nine classic anchors so one ill-conditioned $\text{Rel-}L^2(u)$ ($u^* \equiv 0$) cannot dominate. Hard families are $\{3, 4, 5\}$. Star-state error is $|\log p_* - \log p_*^{\text{ex}}|$ and $|u_* - u_*^{\text{ex}}|$. Front error is $|x_{\text{pred}} - x_{\text{ex}}|$ at $t = T$. The exact contact location is $x_d + u_*^{\text{ex}}T$. The predicted contact is the grid node that maximises $|\partial_x \rho| / (|\partial_x p| + 10^{-6})$ in a window of half-width 0.12 about that location (large density gradient, small pressure gradient). If the Toro structure has a right (else left) shock, the predicted shock is the node of maximum $|\partial_x \rho|$ in the same-width window about $x_d + s_{\text{sh}}^{\text{ex}}T$. Rarefaction heads and tails are not scored as fronts. Locations are grid nodes; there is no subcell interpolation. Wall-clock in Table 4 is the unbatched sequential forward CPU time to map the nine anchors through $\mathcal{S}$. Batched forward and reverse-mode cost, including reconstruction, is Table 17. Large-batch GPU throughput and memory scaling are not measured.

Discretisation floor. The exact Newton star plus $\mathcal{S}$ is an accuracy oracle, not a runtime-equivalent baseline. It does not yield machine-zero field error. Both the label and the oracle sample a similarity solution onto the same 128×33 nodes, but the label is Toro’s float64 primitive sample and $\mathcal{S}$ is a batched float comparison $\xi \leq s$. On eight anchors the two fields agree to $\text{Rel-}L^2 \sim 10^{-8}$. On the collision tube the oracle $\text{Rel-}L^2$ is 0.027. That residual is a one-cell shock-side assignment at one intermediate node, not a Newton miss; Supplementary Appendix C records the cell and shows that a plateau-only and a conserved $\text{Rel-}L^2$ both drive the oracle collision score to the float floor. The Q_0 collision error 0.038 is that floor plus a small Picard star bias ($p_* = 1685$ versus 1692).

2.4 Why an arithmetic midpoint is not enough

The original midpoint-skip map used

$$p_* = p_{\text{mid}} e^{8 \tanh n_p}, \quad u_* = u_{\text{mid}} + (a_L + a_R + |u_L - u_R|) \tanh n_u, \quad (5)$$

with $p_{\text{mid}} = \frac{1}{2}(p_L + p_R)$. The arithmetic-midpoint construction is already three sharp waves with the *wrong* star, so its trained correction must move (n_p, n_u) . Classic median fell to 0.105, but six holdouts stayed odd (Table 2):

Table 2 Odd midpoint-skip holdouts (epoch 20).

ID	Case	Rel- L^2	Failure
2	Toro 123	0.031	Metric good; flat well ($p_* = 0.035$ vs 0.002)
3	Left blast	0.199	p_* stuck at midpoint (499 vs 461)
4	Right blast	0.162	p_* 41.6 vs 46.1; shock/contact merge
6	Stat. shock	0.065	Inconsistent (p_*, u_*); walking contact
7	Stat. contact	—	$u_* \approx -0.045$ (truth 0); $\text{Rel-}L^2(u)$ explodes
8	Slow contact	0.151	$u_* = 0.032$ vs 0.100; fake acoustic well

The arithmetic midpoint is a bad prior for vacuum and blast tubes. On contacts, u_{mid} is already the truth: training *created* the leak. That midpoint construction is retained only as an ablation. The main baseline below is a fixed-budget physical approximation to the star state.

3 Structured surrogate and fixed-budget baseline

The model factorises into a fixed-budget approximate solver, an optional learned residual, and an Euler reconstruction,

$$Q_\theta = \mathcal{S} \circ H_\theta, \quad H_\theta(\text{IC}) = (p_*, u_*), \quad \mathcal{S}(p_*, u_*; \text{IC}, x, t) = (\rho, u, p). \quad (6)$$

The learned residual H_θ is the only trained map. $\mathcal{S}$ contains no Newton iteration on Toro’s pressure function $f(p)$. It is analytic and piecewise differentiable with respect to the predicted star state away from branch-switch surfaces and grid-front crossings (Section 3.5). Crucially, the zero-residual map is a complete, fixed-cost solver in its own right. It is therefore a primary baseline rather than an ablation to exclude from evaluation.

3.1 Fixed-budget approximate star solver Q_0

Toro’s *initial guess* [5], not the Newton root and not a Roe, HLL, or HLLC numerical flux [1, 2, 4], is computed from the IC endpoints and then *detached*. Let $a_k = \sqrt{\gamma p_k / \rho_k}$ and $z = (\gamma - 1)/(2\gamma)$. The primitive-variable Riemann solver (PVRS) is Toro’s linearized acoustic estimate of the star pressure from the left and right primitive states [5]. It is inexpensive, but it is inaccurate on strong shocks and is used here as the weak matched base. Its closed form is

$$p_{\text{PV}} = \frac{1}{2}(p_L + p_R) - \frac{1}{8}(u_R - u_L)(\rho_L + \rho_R)(a_L + a_R). \quad (7)$$

The two-rarefaction (TR) closed form is

$$p_{\text{TR}} = \left(\frac{a_L + a_R - \frac{1}{2}(\gamma - 1)(u_R - u_L)}{a_L p_L^z + a_R p_R^z} \right)^{1/z}. \quad (8)$$

The two-shock Picard recurrence (6 fixed passes, no line search) is

$$A_k = \frac{2}{(\gamma + 1)\rho_k}, \quad B_k = \frac{\gamma - 1}{\gamma + 1} p_k, \quad g_k = \sqrt{\frac{A_k}{p + B_k}}, \quad (9)$$

$$p \leftarrow \frac{g_L p_L + g_R p_R - (u_R - u_L)}{g_L + g_R}. \quad (10)$$

The two-shock initial iterate is the PVRS seed p_{PV} ; densities and pressures in the skip are clamped at 10^{-8} . The branch is

$$p_{\text{guess}} = \begin{cases} p_{\text{PV}}, & p_{\text{max}}/p_{\text{min}} < 2 \text{ and } p_{\text{PV}} \in [p_{\text{min}}, p_{\text{max}}], \\ p_{\text{TR}}, & p_{\text{PV}} < p_{\text{min}}, \\ p_{\text{TS}}, & \text{otherwise.} \end{cases} \quad (11)$$

The matching velocity uses one evaluation of the shock/rarefaction function $f_k(p)$ [5],

$$u_{\text{guess}} = \frac{1}{2}(u_L + u_R) + \frac{1}{2}(f_R(p_{\text{guess}}) - f_L(p_{\text{guess}})). \quad (12)$$

This is the same closed form as Toro’s starter, not the iterated root of $f_L(p) + f_R(p) + u_R - u_L = 0$. With the Euler sampler below, the resulting fixed-cost map is denoted by $Q_0 = \mathcal{S}(p_{\text{guess}}, u_{\text{guess}}; \text{IC})$. It is not a learned model and is evaluated wherever the learned correction is evaluated.

Algorithm 1 Fixed-budget star Q_0 (no Newton, no learned residual)

```

1: compute  $p_{PV}$ ,  $p_{TR}$ , and  $p_{TS}$  after  $N$  Picard passes
2: if  $p_{\max}/p_{\min} < 2$  and  $p_{PV} \in [p_{\min}, p_{\max}]$  then
3:    $p_{\text{guess}} \leftarrow p_{PV}$ 
4: else if  $p_{PV} < p_{\min}$  then
5:    $p_{\text{guess}} \leftarrow p_{TR}$ 
6: else
7:    $p_{\text{guess}} \leftarrow p_{TS}$ 
8: end if
9:  $u_{\text{guess}} \leftarrow \frac{1}{2}(u_L + u_R) + \frac{1}{2}(f_R - f_L)$  at  $p_{\text{guess}}$ 
10: return  $\mathcal{S}(p_{\text{guess}}, u_{\text{guess}}; \text{IC}, x, t)$ 

```

Constituent maps used in the Pareto study freeze a single branch: PVRs only, TR only, or two-shock with $N \in \{0, 1, 3, 6\}$. HLL, HLLC, HLLC, Roe, and entropy-fixed Roe [1–4] are classical numerical-flux constructions, not matched star-state starters for $\mathcal{S}$. They are cited as external references (Table 3) and are not treated as interchangeable baselines for the residual question. This paper does not claim that Q_0 is the best general-purpose Riemann solver, nor that the learned residual competes with that flux literature.

Table 3 Solver roles in this diagnostic study. HLL-family and Roe-type methods are unmatched external references.

Method	Wave structure	Iteration	Role here
PVRs / TR / fixed Picard	Toro-style star approximation	Fixed budget	Matched base family
HLL / HLLC	Two-wave flux	None	Classical external reference
HLLC	Three-wave flux	None	Classical external reference
Roe / entropy-fixed Roe	Linearized characteristic flux	None	Classical external reference
Newton-star oracle	Exact star + $\mathcal{S}$	Iterative	Accuracy reference

3.2 Acoustic gate

A pure contact has $\Delta p = 0$ and $\Delta u = 0$. Acoustic waves need a jump. The gate

$$g = \tanh\left(4\left(\frac{|p_L - p_R|}{p_{\text{mid}}} + \frac{|u_L - u_R|}{\bar{a}}\right)\right) \in [0, 1] \quad (13)$$

vanishes on families 7 and 8 and is ≈ 1 on Sod, the blasts, and Toro 123 ($\Delta u \neq 0$).

3.3 Learned residual H_θ

The IC is encoded as

$$\phi = \left(\frac{\rho_L}{s_\rho}, \frac{u_L}{s_u}, \frac{p_L}{s_p}, \frac{\rho_R}{s_\rho}, \frac{u_R}{s_u}, \frac{p_R}{s_p}, \gamma, \log \frac{p_L}{p_R}, \log \frac{\rho_L}{\rho_R}, \frac{u_R - u_L}{\bar{a}}\right) \in \mathbb{R}^{10}. \quad (14)$$

A multilayer perceptron $\text{MLP}_\theta : \mathbb{R}^{10} \rightarrow \mathbb{R}^2$ (widths $10 \rightarrow 64 \rightarrow 64 \rightarrow 2$, last layer zero-initialised; 4994 parameters) produces $(n_p, n_u) = \tanh(\text{MLP}_\theta(\phi))$. The weak-PVRs and Newton-star controls later establish that this residual parameterization can represent and optimize useful star corrections; the reported non-improvement over Q_0 is therefore not attributable solely to zero-initialization, contact gating, or insufficient output capacity. The star is

$$p_* = p_{\text{guess}} \exp(8 g n_p), \quad (15)$$

$$u_* = u_{\text{guess}} + g(a_L + a_R + |u_L - u_R|) n_u. \quad (16)$$

At zero last-layer weights, $n_p = n_u = 0$, so the model reduces exactly to the fixed baseline Q_0 for every IC. When $g = 0$, the residual cannot fire: contacts stay at the guess ($u_* = u_L = u_R$) even if MLP_θ outputs large logits. The factor 8 still gives a multiplicative range $e^{\pm 8}$ when the gate is open, enough for Toro 123 and colliding shocks if the guess is poor.

3.4 Star densities and Lax speeds

Given a prescribed p_* , write $\Gamma = (\gamma + 1)/(\gamma - 1)$. If $p_* > p_k$ the k -wave is a shock and

$$\rho_k^* = \rho_k \frac{p_*/p_k \Gamma + 1}{\Gamma + p_*/p_k}, \quad k \in \{L, R\}. \quad (17)$$

If $p_* \leq p_k$ the wave is a rarefaction and

$$\rho_k^* = \rho_k (p_*/p_k)^{1/\gamma}. \quad (18)$$

These densities are thermodynamically consistent with the *learned* p_* ; they are not copied from Toro's Newton output.

Sound speeds in the star are $a_L^* = \sqrt{\gamma p_*/\rho_L^*}$ and $a_R^* = \sqrt{\gamma p_*/\rho_R^*}$. Left and right waves follow the Lax formulae

$$\text{left shock } (p_* > p_L) : \quad s_L = u_L - a_L \sqrt{\frac{\gamma+1}{2\gamma} \frac{p_*}{p_L} + \frac{\gamma-1}{2\gamma}}, \quad (19)$$

$$\text{left fan } (p_* \leq p_L) : \quad s_L^{\text{head}} = u_L - a_L, \quad s_L^{\text{tail}} = u_* - a_L^*, \quad (20)$$

$$\text{right shock } (p_* > p_R) : \quad s_R = u_R + a_R \sqrt{\frac{\gamma+1}{2\gamma} \frac{p_*}{p_R} + \frac{\gamma-1}{2\gamma}}, \quad (21)$$

$$\text{right fan } (p_* \leq p_R) : \quad s_R^{\text{tail}} = u_* + a_R^*, \quad s_R^{\text{head}} = u_R + a_R. \quad (22)$$

The contact is the entropy wave

$$s_c = u_*, \quad x_c(t) = x_d + u_* t. \quad (23)$$

3.5 Self-similar sample $\mathcal{S}$

For $t > 0$ the similarity coordinate is $\xi = (x - x_d)/t$. The sampler is Toro's sampling procedure [5] with $(p_*, u_*, \rho_L^*, \rho_R^*)$ prescribed by H_θ .

Left of the contact ($\xi < u_*$). If the left wave is a shock,

$$(\rho, u, p) = \begin{cases} (\rho_L, u_L, p_L), & \xi < s_L, \\ (\rho_L^*, u_*, p_*), & \xi \geq s_L. \end{cases} \quad (24)$$

If it is a rarefaction and $\xi \in [s_L^{\text{head}}, s_L^{\text{tail}}]$,

$$u = \frac{2}{\gamma + 1} \left(a_L + \frac{\gamma-1}{2} u_L + \xi \right), \quad (25)$$

$$a = \frac{2}{\gamma + 1} \left(a_L + \frac{\gamma-1}{2} (u_L - \xi) \right), \quad (26)$$

$$\rho = \rho_L (a/a_L)^{2/(\gamma-1)}, \quad p = p_L (a/a_L)^{2\gamma/(\gamma-1)}. \quad (27)$$

Outside the fan one recovers either the left IC or the left star.

Right of the contact ($\xi \geq u_*$). The shock branch is the symmetric jump at s_R . Inside a right rarefaction,

$$u = \frac{2}{\gamma + 1} \left(-a_R + \frac{\gamma-1}{2} u_R + \xi \right), \quad (28)$$

$$a = \frac{2}{\gamma + 1} \left(a_R - \frac{\gamma-1}{2} (u_R - \xi) \right), \quad (29)$$

with the same isentropic lifts (27) from the right state.

At $t = 0$ one recovers the IC step about x_d . Hard IC/Dirichlet (3) are applied after $\mathcal{S}$. No unreached, star, or fan pin is used. No field-operator model is used in the present experiments.

The reconstruction is analytic and piecewise differentiable with respect to $(\log p_*, u_*)$ away from branch-switch surfaces ($p_* = p_k$) and events in which a wave front crosses a grid node. Gradients at those measure-zero boundaries are not defined in the classical sense; automatic differentiation follows the active $\xi \leq s$ branch. Supplementary Appendix E reports centred finite-difference checks ($\varepsilon = 10^{-4}$) at the frozen Q_0 star: relative errors on $\partial L / \partial \log p_*$ are 10^{-5} – 10^{-2} on Sod, the blasts, and the stationary shock, while collision $\partial L / \partial u_*$ sits on a kink (autograd 1.4×10^{-3} , centred difference 0). Contacts have a vanishing field gradient because the acoustic gate and the skip are inactive.

If H_θ outputs the *true* Toro star, $\mathcal{S}$ reproduces the exact field (Sod Rel- $L^2 < 5 \times 10^{-3}$ on 128×17). Training is therefore equivalent to learning a residual $\text{IC} \mapsto (\Delta \log p_*, \Delta u_*)$; the three waves are then written exactly and cannot smear.

Algorithm 2 Structured map Q_θ : analytic wave reconstruction with a gated residual

- 1: $\phi \leftarrow$ scaled endpoints and log ratios (14)
 - 2: $(n_p, n_u) \leftarrow \tanh(\text{MLP}_\theta(\phi))$
 - 3: $(p_{\text{guess}}, u_{\text{guess}}) \leftarrow$ PVRs/TR/Picard (11)–(12), detached
 - 4: $g \leftarrow$ acoustic gate (13)
 - 5: $(p_*, u_*) \leftarrow$ residual (15)–(16)
 - 6: $(\rho_L^*, \rho_R^*) \leftarrow$ Hugoniot or isentrope (17)–(18)
 - 7: speeds $\leftarrow$ Lax formulae (19)–(23)
 - 8: $(\rho, u, p) \leftarrow \mathcal{S}(\xi; p_*, u_*)$
 - 9: overwrite $t = 0$ and $x \in \{0, L\}$ by (3)
 - 10: **return** (ρ, u, p)
-

3.6 Model boundary and primary comparisons

Ingredient	Source	Role in this study
Residual (n_p, n_u)	trained $H_\theta(\text{IC})$	learned correction
$(p_{\text{guess}}, u_{\text{guess}})$	PVRs/TR/6-pass Picard	fixed solver Q_0
Zero-residual field	Q_0 sampled by $\mathcal{S}$	primary baseline
ρ_L^*, ρ_R^* , speeds, fans	Euler structure of $\mathcal{S}$	fixed reconstruction
Newton p_* written in	Toro root of $f(p)$	excluded (oracle)
Arithmetic midpoint	WS-1 construction	ablation
Hard IC / Dirichlet ends	IBVP data	imposed data

4 Training protocol

Let $\hat{q} = Q_\theta(q_0)$ and q^* the exact label. Field Rel- L^2 is cheap to miss on the Toro 123 well and is dominated by $p \sim 10^3$ on blasts. The data term is

$$\mathcal{L}_{\log} = \frac{1}{3} \left(\overline{\|\log \hat{p} - \log \rho^*\|^2} + \overline{\|(\hat{u} - u^*)/s_u\|^2} + \overline{\|\log \hat{p} - \log p^*\|^2} \right), \quad (30)$$

plus the same three terms restricted to *label plateaus* (cells whose $|\partial_x q^*|$ sits below the 0.7-quantile of that sample). The plateau mask comes from the label gradient, not from Toro geometry. Total variation is off ($w_{\text{tv}} = 0$). The exact-star teacher L_{star} stays at weight 0: it overfits Sod and wrecks blasts.

A soft star-identity residual

$$\mathcal{L}_{\text{id}} = \left\| \frac{f_L(p_*) + f_R(p_*) - (u_R - u_L)}{a_L + a_R + |u_L - u_R|} \right\|^2 + \left\| \frac{u_* - u_{\text{id}}(p_*)}{a_L + a_R + |u_L - u_R|} \right\|^2 \quad (31)$$

is included and left at weight 0 in the mixed-weight table. It is the Euler star equation as a physics loss, not a Newton solve. It was not needed for the stationary shock: the PVRs skip already holds that jump. Section 6.5 later uses (31) as a sole objective on the same Q_0 skip.

The critic is the WENO-Z/Rusanov residual of (1), a strong-form physics-informed field residual in the sense of Raissi et al. [6], not the conservative control-volume PINN of Jagtap et al. [8]. Section 6.5 tests a trapezoid control-volume residual and a strong-form entropy inequality as sole replacements, not as that mixed critic. The default weight $w_{\text{pde}} = 10^{-4}$ is a hypothesis under test, not a validated regularizer. A field-only ablation sets that weight to zero and keeps (30). Adam updates *only* MLP_θ (learning rate 3×10^{-3} , cosine, 40 epochs, CPU). The unused lift and spectral-block weights are not in the graph. Training uses the frozen protocol-v2 split: 120 train ICs, 9 validation, 9 classic holdouts never seen in training; grid 128×33 ; exact labels; inactive T . Checkpoint selection is separated into an original study and a confirmatory selection study. In the original five-seed field-residual study, validation was recorded only at epochs 1, 10, 20, 30, 40, and the selected checkpoint minimized the validation overall median $\text{Rel-}L^2$. That sparse cadence produced two epoch-1 selections. The original experiment is retained as a historical diagnostic and is not the basis of the strengthened selection claim. In the confirmatory study of Section 5.3, validation is evaluated at every epoch for seeds $0, \dots, 9$. Both selection rules are applied to the same stored trajectories: the original overall-median rule and a prospectively specified composite rule. The nine held-out canonical anchors are never used for checkpoint selection. Seeds change the network initialization, Adam random state, and per-epoch order of the fixed training cases; they do not redraw the protocol-v2 split.

5 Diagnostic results

5.1 Fixed-budget Pareto

Table 4 is the primary comparison on the nine classic anchors. All unlearned rows use the same sampler $\mathcal{S}$. Rows labelled PVRs, two-rarefaction, and two-shock freeze that constituent for every case, irrespective of wave regime. The branched map Q_0 applies (11): PVRs on mild pressure ratios, two-rarefaction when the PVRs seed lies below p_{min} , and otherwise a six-pass two-shock Picard update. On the nine anchors that rule selects two-shock six times (Sod, both blasts, collision, stationary shock, slow shock), two-rarefaction once (Toro 123), and PVRs twice (both contacts). On the locked protocol-v2 split the same rule selects two-shock / TR / PVRs on 96/8/16 of the 120 training ICs, 6/1/2 of the nine validation ICs, and 12/2/4 of each 18-member in-family and detuned-family set.

On this nine-anchor set, three Picard passes already match Q_0 on the hard median; six passes reduce the worst-case collision error ($0.186 \rightarrow 0.038$). Two-shock with $N = 6$ alone is not Q_0 : it misses the Toro 123 rarefaction branch (field error 0.007) and therefore has a slightly larger classic median. The exact Newton star is an oracle; its collision floor 0.027 is the one-cell sampling discrepancy of Section 2.3, not a root-finding miss. Reported times are forward CPU evaluations of all nine anchors on a 128×33 grid, including reconstruction. They are a pilot cost check: GPU throughput, batch scaling, memory, and reverse-mode differentiation remain unmeasured. Within that unbatched forward-CPU pilot the times are comparable.

Table 4 Fixed-budget Pareto on the nine classic anchors. Hard families are 3–5. Constituent rows freeze one formula for every case; Q_0 uses the branch rule (11) with six Picard passes on the two-shock branch. Learned rows attach the same gate to that Q_0 base. Times are forward CPU milliseconds per case, including reconstruction; GPU, batch, memory, and reverse-mode costs are not reported.

Method	Classic	Soft	Hard	Worst	ms/case
Midpoint	0.353	0.242	0.861	0.861	7.5
PVRS only	0.077	0.006	0.203	0.540	9.8
Two-rarefaction only	0.031	0.006	0.617	0.693	6.4
Two-shock only, $N = 0$	0.077	0.006	0.203	0.540	6.9
Two-shock only, $N = 1$	0.035	0.004	0.048	0.454	5.5
Two-shock only, $N = 3$	0.0004	0.000	0.034	0.186	7.1
Two-shock only, $N = 6$	0.0001	0.000	0.034	0.038	5.5
Q_0 (branched, $N = 6$)	0.000	0.000	0.034	0.038	5.2
Exact Newton star (oracle)	0.000	0.000	0.000	0.027	4.7
Field residual on Q_0 (WS-2)	0.014	0.003	0.068	0.083	6.6
Field residual on Q_0 (field only)	0.002	0.001	0.068	0.070	5.8

On this diagnostic set the residual attached to Q_0 does not improve that base and is worse than a one-pass two-shock map on the hard median. Removing the Rusanov critic improves the classic median ($0.014 \rightarrow 0.002$) but leaves the hard median at 0.068. The log-field loss, not only the PDE critic, walks the star off Q_0 over training, even though the first gradient at the frozen Q_0 star is not anti-aligned with a Newton teacher (Supplementary Section 2.1). Gated residuals on weaker PVRS and one-pass bases are Table 12.

5.2 Five-seed residual

Table 5 and Figure 1 lock the same split, loss, and checkpoint rule (validation median) across five seeds. Mean classic $\text{Rel-}L^2$ is 0.009 ± 0.013 ; mean hard is 0.079 ± 0.050 . On this diagnostic set no seed improves on the Q_0 hard median 0.034. Two seeds finish near Q_0 (hard 0.034). The original seed-0 checkpoint used in the per-anchor tables (classic 0.014, hard 0.068) is a typical degraded run, not an outlier; the five-seed seed-0 rerun selected epoch 40 (classic 0.004, hard still 0.068).

Table 5 Gated residual, five seeds, locked protocol-v2 split. Checkpoint is validation median. Q_0 hard median is 0.034.

Seed	Best epoch	Classic	Soft	Hard
0	40	0.004	0.001	0.068
1	1	0.032	0.016	0.150
2	1	0.005	0.002	0.108
3	40	0.001	0.001	0.034
4	40	0.001	0.000	0.034
mean $\pm$ std		0.009 ± 0.013	0.004 ± 0.007	0.079 ± 0.050

Seeds 1 and 2 select epoch 1 under the original locked validation-median rule. This observation belongs to the original sparse-cadence analysis, in which validation was recorded only at epochs 1, 10, 20, 30, 40. It does not describe the confirmatory analysis. In the latter, validation was evaluated at every epoch, neither selection rule selected epoch 1, and the overall-median rule still yielded no hard-anchor gains over Q_0 .




Fig. 1 Classic and hard medians for the original five-seed sparse-cadence study. Dashed line: Q_0 hard median 0.034. Seeds 1 and 2 stop at epoch 1 under that sparse log; the confirmatory full-epoch study does not select epoch 1.




Fig. 2 Original sparse-cadence seed-0 training log, retained for historical diagnosis; it is not the basis of the confirmatory selection result. Locked validation overall, soft, and hard medians versus epoch. The dotted line is the selected checkpoint (epoch 30).

5.3 Confirmatory hard-aware checkpoint selection

The original five-seed study used validation records only at epochs 1, 10, 20, 30, 40, so it could not distinguish selection-rule behavior from a sparse-logging artifact. We therefore conducted a confirmatory experiment specified before training. The same field-residual formulation around Q_0 , locked protocol-v2 split, optimizer schedule, and 40-epoch budget were used for seeds 0, $\dots$, 9. Validation was evaluated at every epoch. The nine held-out canonical anchors were never used for selection. Both rules were applied to the same stored trajectory from each seed:

$$J_{\text{overall}} = \text{median}(e_{\text{all}}), \quad (32)$$

$$J_{\text{composite}} = \frac{1}{2} \text{median}(e_{\text{all}}) + \frac{1}{2} \text{median}(e_{\text{hard}}). \quad (33)$$

Table 6 Confirmatory checkpoint-selection study for the field residual around Q_0 . For each seed, the hard-anchor score is the median over left blast, right blast, and collision. The table reports the mean over the ten seed-wise paired differences Δ_s ; the protocol and bootstrap construction are given in Section 5.3.

Selection rule	Cadence	Mean hard-anchor median	Mean Δ [95% CI]	Practical gains
Overall median J_{overall}	every epoch	0.063	+0.028 [+0.018, +0.039]	0/10
Composite $J_{\text{composite}}$	every epoch	0.044	+0.010 [−0.001, +0.020]	2/10
Fixed baseline Q_0	—	0.034	0	—

For the composite rule, the two practical gains occur at seeds 0 and 3. Eight of ten seed-wise differences remain positive, and the seed-wise median paired difference is +0.013.

Here e_{all} is the per-case validation score over all nine validation members and e_{hard} is the score over its three predefined hard-family members. A practical gain was predefined as $\Delta_s < -0.01$, where

$$\Delta_s = \text{median}_{c \in \{3,4,5\}} \left(e_{s,c}^{\text{holdout}} \right) - \text{median}_{c \in \{3,4,5\}} \left(e_c^{Q_0} \right) \quad (34)$$

for seed s . The second term is deterministic and equals 0.034 on the three hard held-out anchors.

Table 6 and Figure 3 show the outcome. Both rules select from the same stored training trajectory for each seed. The reported paired point estimate is the arithmetic mean of the ten seed-wise differences Δ_s . Its 95% interval is a percentile bootstrap interval obtained from $B = 10,000$ resamples, with replacement, of these ten paired seed values. The random seed (20260910) and all ten Δ_s values are provided in Supplementary Table 1 and the reproducibility archive. The deterministic Q_0 hard-anchor reference is 0.034. A practical gain was predefined as $\Delta_s < -0.01$. Full-epoch validation removes the earlier epoch-1 artifact: neither rule selects epoch 1. The overall-median rule nevertheless produces no practical hard-anchor gains (0/10); its mean seed-wise paired difference is +0.028, with bootstrap interval [+0.018, +0.039]. The composite rule reduces mean hard-anchor error from 0.063 to 0.044 and yields two practically improved seeds (0 and 3, with $\Delta = -0.020$ and -0.022). However, the remaining eight seeds are worse than Q_0 , the seed-wise median paired difference is +0.013, and the bootstrap interval for the mean paired difference is [−0.001, +0.020] and therefore includes zero. No seed improves the overall median on the nine held-out anchors under either rule. Thus, selection aligned with hard-family validation error reduces average degradation but does not make the tested field-residual method repeatably beneficial.

5.4 Finite-volume star embedding

For a finite-volume CFD application, a star correction can affect the solution only through the interface flux. We therefore embed the same (p_*, u_*) maps in a first-order Godunov-type finite-volume scheme whose reconstruction, CFL, and time integrator stay fixed: donor faces, CFL 0.4, SSP-RK2, and Dirichlet ends equal to the initial states. The only varying ingredient is the interface star. The interface flux is the exact Euler flux of the three-wave state sampled at $\xi = 0$ from that prescribed star,

$$\widehat{\mathbf{F}}_{i+1/2} = \mathbf{F}(\mathcal{S}_{\text{interface}}[p_*, u_*; \mathbf{U}_i, \mathbf{U}_{i+1}]), \quad (35)$$

where $\mathcal{S}_{\text{interface}}$ is the $\xi = 0$ sample of the same three-wave reconstruction used in the analytic study. The protocol was written before evaluation. The 100–400-cell range was selected as a low-cost diagnostic resolution range in which a first-order donor scheme remains truncation dominated; it is not intended as a grid-convergence study of a proposed flux. Grids are 100, 200, and 400 cells. Labels are the exact Riemann solution on cell centers at the inactive horizon.




Fig. 3 Confirmatory hard-aware checkpoint selection. Both rules are applied to the same ten training trajectories; validation is evaluated at every epoch and the held-out anchors are not used for selection. Panel (a) shows seed-wise hard-anchor differences Δ_s relative to Q_0 : the solid line denotes parity and the dashed line denotes the prospectively specified practical-gain threshold -0.01 . Panel (b) shows selected epochs. Full-epoch validation removes the earlier epoch-1 logging artifact. The composite rule reduces degradation and recovers two practically improved seeds, but does not yield a dependable aggregate improvement.

Methods are Q_0 , the locked-rule residual from seed 0, the two composite case-study maps (seeds 0 and 3), and a Newton-star oracle. Seeds 0 and 3 are not selected on finite-volume error.

Table 7 and Figure 4 report hard-family and nine-anchor medians of primitive $\text{Rel-}L^2$ at $t = T$. At every grid, the maximum inter-method difference in the reported median errors is at most 10^{-3} . Q_0 and the Newton-star oracle are indistinguishable at this first-order resolution: donor-cell truncation masks the remaining star error that the sampler study still sees on the right blast. The locked residual and both composite case studies stay on the same curves. At the saved final time, across all tested maps, grids, and nine anchors, density and pressure remain positive; the observed minima are $\rho_{\min} \approx 8 \times 10^{-3}$ and $p_{\min} \approx 3 \times 10^{-3}$. Because fixed Dirichlet end states permit net boundary fluxes, global domain integrals are not expected to remain constant. We therefore do not use constancy of domain-integrated conserved variables as a diagnostic in this test. At the tested first-order resolutions, therefore, donor-cell truncation masks the differences among the evaluated star maps. This result does not assess whether those differences persist under higher-order reconstruction or finer grids.

Table 7 First-order Godunov-type finite-volume embedding of the same star maps. Entries are nine-anchor / hard-family median $\text{Rel-}L^2$ at $t = T$. Composite seeds 0 and 3 are case studies, not a selected solver.

Star map	$n = 100$	$n = 200$	$n = 400$
Q_0	0.080/0.087	0.067/0.072	0.044/0.051
Locked seed 0	0.080/0.087	0.067/0.072	0.044/0.051
Composite seed 0	0.080/0.087	0.067/0.072	0.044/0.051
Composite seed 3	0.080/0.087	0.067/0.071	0.044/0.051
Newton star	0.080/0.087	0.067/0.072	0.044/0.051




Fig. 4 Hard-family median Rel- L^2 of the first-order Godunov-type rollout versus grid size. Only the star map changes; reconstruction, CFL, and SSP-RK2 are shared.

5.5 Nine classic holdouts

Table 8 Per-anchor Rel- L^2 for Q_0 and the seed-0 field residual on Q_0 (WS-2), with stars versus the Newton exact state.

Case	Q_0	WS-2	p_{Q_0}	p_θ	p^{ex}	u_{Q_0}	u_θ	u^{ex}
Sod	0.000	0.014	0.303	0.301	0.303	0.927	0.927	0.927
Toro 123	0.000	0.004	0.002	0.002	0.002	0	0.098	0
Left blast	0.000	0.068	462	455	461	19.59	19.59	19.60
Right blast	0.034	0.056	46.2	45.6	46.1	-6.19	-6.18	-6.20
Collision	0.038	0.083	1685	1672	1692	8.69	8.68	8.69
Stat. shock	0.000	0.002	0.804	0.800	0.804	0.375	0.376	0.375
Stat. contact	0.000	0.000	0.400	0.400	0.400	0	0	0
Slow contact	0.000	0.000	1.00	1.00	1.00	0.100	0.100	0.100
Slow shock	0.000	0.034	10.33	10.25	10.33	-0.811	-0.810	-0.811

The trained walk-off from the detached Q_0 star is small (Supplementary Table 2): contacts remain at zero by the gate, the largest velocity correction is Toro 123 ($|\Delta u_*| = 0.098$), and the largest pressure correction is the left blast ($|\Delta \log p_*| = 0.013$). Primitive and front-location errors are Supplementary Table 3.

5.6 Representative holdout fields

Figures 5 and 6 are IC-only predictions of the trained residual (red dashed) against the exact Toro label (black) at four fractions of the inactive horizon. The other seven anchors are in the Supplementary Information (Supplementary Appendix A). Residual error is a star bias, not a Fourier ramp.

Sod (soft). Q_0 is already exact on this grid (Rel- L^2 0.000). The trained residual walks to 0.014 ($p_* = 0.301$ versus 0.303).




Fig. 5 Sod (family 1). Jumps are sharp; the trained star is a small walk-off from Q_0 .

Collision (hard). Two shocks and a contact. $Q_0 = 0.038$ ($p_* = 1685$); the exact-star oracle floor is 0.027. Trained p_* is 1672 versus 1692. Trained $\text{Rel-}L^2$ 0.083; shock x -err 0.007. Training does not tighten the Picard guess.




Fig. 6 Collision (family 5). Two-shock structure is sharp; u matches well (0.049). The remaining error is a small star-pressure bias.

Contacts remain at machine-zero field error because the acoustic gate vanishes. The remaining seven anchors, physics residuals, and the field-residual holdout bar are in the Supplementary Information (Supplementary Appendix A).

6 Interpretation

The experiments give two clear findings (Table 9). First, the residual can learn useful star corrections when the base is weak PVRs. Second, it does not provide a repeatable improvement over the stronger Q_0 starter under either predeclared full-epoch selection rule. The remaining Q_0 errors are real. The right-blast error is 0.034, while the exact-star oracle is at 0.000. The collision error is 0.038, of which 0.027 is a known one-cell assignment floor (Table 8).

The evidence is empirical. It comes from matched baselines, a weak-base rescue, loss ablations, multiple seeds, a star-supervision control, and front-insensitive metrics. The Newton-star teacher is only a control. It is not a deployable solver (Section 6.4). The field loss is a usual pointwise physics-informed objective used in PINN-style training [6]. Section 6.5 tests three replacements on the same Q_0 skip. Under the locked overall-median selection protocol used for that objective-comparison experiment, none of the four tested objectives yielded a dependable hard-anchor

improvement over Q_0 . For the field-residual objective, the confirmatory full-epoch composite-selection study reduced degradation but did not produce repeatable improvement (Section 5.3). They are not dual-norm weak residuals [10, 11] or locally linearized shock treatments [12].

Table 9 Main findings on the locked nine-anchor set. The PVRs and original field-residual results use five seeds. The objective comparison uses ten seeds.

Finding	Evidence	Interpretation
Residual learning can help a weak base	PVRs hard median $0.203 \rightarrow 0.045 \pm 0.007$ on all five seeds	The residual and optimiser can learn useful star corrections.
Four objectives do not help a strong base under the locked rule	Across ten seeds, field, star-identity, entropy, and control-volume objectives give no dependable hard-family gain over Q_0 under the locked overall-median protocol	That protocol is the objective-comparison experiment, not the later field-residual composite-selection study.
Hard-aware selection reduces harm, not enough to deploy	Composite (33) yields 2/10 practical gains; eight seeds remain worse; paired interval includes zero	Occasional residual capacity, not a repeatable improvement over Q_0 .

6.1 Label convention does not invert the Toro-protocol headline

The primitive space-time Rel- L^2 is front-sensitive. On the collision anchor, the Newton-star reconstruction and the Toro label assign opposite sides of a shock at one intermediate grid node; hence the Newton-star oracle has nonzero primitive error under this convention. Conserved-state and plateau-only metrics remove this sampling ambiguity. The ordering relevant to this paper is unchanged: the representative field-residual correction remains worse than Q_0 on the hard-anchor median under primitive, conserved, and plateau-aware scores. Detailed node-level diagnostics are provided in the supplement (Supplementary Appendix C). Table 10 is the main robustness check.

The mismatch is localised to the collision front band. Under conserved-state and plateau-only metrics the seed-0 Toro-trained residual remains worse than Q_0 (hard medians 0.103 versus 0.042 conserved, 0.076 versus 0.028 plateau). The same residual also degrades non-collision anchors (Sod, left blast, stationary shock, slow shock; Table 8). The five-seed hard median is over left blast, right blast, and collision, and no Toro-trained seed beats Q_0 's hard median 0.034 under the locked validation rule. Rescoring those published checkpoints against $\mathcal{S}(\text{Newton})$ labels leaves the same classic / hard pair 0.014/0.068.

A stronger control retrain the field residual on $\mathcal{S}(\text{Newton})$ labels, keeping the locked split and hyperparameters, and evaluates on both conventions (Table 11). That retrain stops the walk-off: five-seed hard median 0.00024 ± 0.00006 versus Q_0 0.034, driven by a real right-blast repair ($0.034 \rightarrow 4 \times 10^{-5} - 4 \times 10^{-4}$ on every seed; the right-blast Q_0 error is identical on both label conventions). Conserved and plateau hard medians at seed 0 fall to $\approx 2 \times 10^{-4}$. The same map still does not beat Q_0 on the nine-anchor median (0.00014 ± 0.00005 versus 8×10^{-6}) or on sampler-consistent protocol-v2 medians, where Q_0 collapses to $\sim 10^{-5}$ and the retrain remains $10^{-4} - 10^{-3}$. Toro 123 can worsen (seed 2 primitive 0.012). Collision improves but does not reach Newton (0.014 versus Toro, 0.034 versus $\mathcal{S}$). On an independent 200-member draw scored against Toro labels, seed-0 of the retrain has median 0.008 versus Q_0 6×10^{-5} .

The Toro-protocol walk-off is therefore a real training outcome, not a scoring artefact of one collision node. Changing the training target to $\mathcal{S}(\text{Newton})$ is a different experiment: it can repair the remaining right-blast error without producing a deployable improvement over Q_0 on self-consistent field medians. It is reported as a control and is not used to rewrite Table 8 or Supplementary Table 7.

Table 10 Nine-anchor medians under front-insensitive metrics (seed-0 residual). Hard entries use families $\{3, 4, 5\}$. The Newton-star oracle is not zero on the primitive metric because of a one-cell collision assignment; conserved and plateau scores fall to the floor. Node-level forensics are in Supplementary Appendix C.

Metric	Scope	Q_0	Oracle	Residual
Primitive Rel- L^2	classic / hard	0.000/0.034	0.000/0.000	0.014/0.068
Conserved Rel- L^2	classic / hard	0.000/0.042	0.000/0.000	0.011/0.103
Plateau Rel- L^2	classic / hard	0.000/0.028	0.000/0.000	0.005/0.076
Front-band Rel- L^2	classic / hard	0.000/0.0002	0.000/0.000	0.002/0.121

Table 11 Self-consistent-label control on the nine anchors. The Toro-trained residual is the published seed-0 map; the retrain is five-seed mean $\pm$ std except the per-case notes. Q_0 hard 0.034 is the right blast, identical on both conventions.

Map (classic / hard)	vs Toro labels	vs $\mathcal{S}$ (Newton)
Q_0	$8 \times 10^{-6}/0.034$	$8 \times 10^{-6}/0.034$
Exact Newton + $\mathcal{S}$	$8 \times 10^{-8}/9 \times 10^{-8}$	0/0
Toro-trained field residual	0.014/0.068	0.014/0.068
$\mathcal{S}$ -trained field residual	0.00014/0.00024	0.00014/0.00024

6.2 Results within the present protocol

- Pareto of fixed bases.** Midpoint, PVRS, two-rarefaction, and two-shock budgets $N \in \{0, 1, 3, 6\}$ are reported in Table 4. On the nine anchors, accuracy of this branch family plateaus at three to six Picard passes on the 128×33 grid. The Newton star is an oracle, not a runtime peer.
- Confirmatory checkpoint selection.** A ten-seed, full-epoch study removes the prior sparse-cadence epoch-1 artifact. The overall-median rule yields no practical hard-anchor gains. The predeclared composite hard-aware rule reduces average degradation and produces two practical gains, but eight of ten seeds remain worse than Q_0 and the aggregate paired interval includes zero (Section 5.3).
- Field-only ablation.** Setting $w_{\text{pde}} = 0$ improves the classic median but leaves the hard median at 0.068. With the log-field loss alone, the prediction moves away from Q_0 .
- Matched per-anchor table.** Table 8 reports Q_0 and the residual together with star states.
- Weak bases.** Five seeds of the PVRS residual give classic / hard 0.021 ± 0.013 / 0.045 ± 0.007 (frozen PVRS 0.077/0.203). One-pass two-shock hard medians stay at the frozen plateau (Table 13).
- Star-aware losses.** Field, identity, and mixed Newton-teacher objectives still degrade Q_0 . Five seeds of a Newton-star teacher with no field term improve the nine-anchor hard median (0.002 ± 0.004); collision is not robust (Tables 14, 13).
- Replacement physics objectives.** The ten-seed, pre-registered comparison tests field, star-identity, entropy, and control-volume objectives on the same Q_0 skip. None improves Q_0 under the locked rule (Table 15). The entropy-only objective is unstable.
- Selection and hard-regime performance.** Overall validation medians can underrepresent hard-family performance. Nevertheless, the confirmatory composite rule shows that aligning checkpoint selection with hard-family validation reduces damage without producing repeatable benefit beyond Q_0 in this protocol.

9. **Tests at new γ values and near vacuum.** The field residual does not improve the Q_0 hard error at $\gamma = 1.2$ or $\gamma = 5/3$ (Table 16). The new near-vacuum anchor is already resolved by the two-rarefaction branch at the Q_0 accuracy floor ($\sim 10^{-8}$). It is a control, not an equally demanding error-reduction test.
10. **Broader splits, Monte Carlo, and cost.** On the small locked in-family / detuned-family sets and on the locked 1800-member Monte Carlo draw (seed-0 learned maps), Q_0 is close to the oracle under paired median metrics. On the tested Apple M2 batch, reverse-mode differentiation through the residual is about $3 \times Q_0$ (Supplementary Tables 5, 7 and Table 17).

6.3 Residual on weaker bases

Table 12 attaches the same gate and field-plus-Rusanov recipe to PVRS and to one-pass two-shock, each trained for 40 epochs on the locked split. The displayed cells in Table 12 are the seed-0 runs used in the open-protocol eval. Table 13 repeats those two weak-base recipes, and the Newton-star-only teacher, over seeds 0–4. On the nine anchors the residual improves PVRS substantially ($0.077/0.203 \rightarrow 0.025/0.048$ at seed 0; five-seed classic / hard 0.021 ± 0.013 / 0.045 ± 0.007). All five PVRS seeds beat the frozen PVRS hard median. The one-pass two-shock residual improves the classic median on four of five seeds (five-seed 0.016 ± 0.013) and leaves the hard median at the frozen plateau (0.048 ± 0.011 versus frozen 0.048); seed 2 does not beat the frozen classic median. Neither trained weak-base map reaches Q_0 ($0.000/0.034$). Validation-median checkpointing selects epoch 25 for two PVRS seeds and one two-shock seed. These results are the weak-base half of Table 9: the residual can improve PVRS, but the trained weak-base maps stay above Q_0 .

Table 12 Same gated residual and field-plus-Rusanov objective on three fixed bases (locked split). Seed-0 cells used in Supplementary Table 5. Five-seed weak-base / teacher studies are Table 13; the Q_0 residual five-seed study is Table 5.

Base	Base classic	Base hard	+residual classic	+residual hard
PVRS only	0.077	0.203	0.025	0.048
Two-shock only, $N = 1$	0.035	0.048	0.020	0.048
Q_0 (branched, $N = 6$)	0.000	0.034	0.014	0.068

Table 13 Five-seed means $\pm$ sample standard deviation of *seed-level* classic / hard / worst medians on the nine anchors (seeds 0–4, locked split, validation-median checkpoint). Not a pooled seed-case standard deviation. Frozen bases are deterministic. Per-seed values are Supplementary Table 8.

Recipe	Classic	Hard	Worst
Frozen PVRS	0.077	0.203	0.540
Residual on PVRS	0.021 ± 0.013	0.045 ± 0.007	0.080 ± 0.031
Frozen two-shock, $N = 1$	0.035	0.048	0.454
Residual on two-shock, $N = 1$	0.016 ± 0.013	0.048 ± 0.011	0.060 ± 0.009
Frozen Q_0	0.000	0.034	0.038
Newton-star teacher only	0.0001 ± 0.0001	0.002 ± 0.004	0.022 ± 0.019




Fig. 7 Five-seed classic and hard medians for residual learning on PVRs, one-pass two-shock, the branched starter Q_0 , and Newton-star-only supervision. Horizontal lines are the corresponding frozen bases. Every PVRs seed beats frozen PVRs; no field residual on Q_0 beats the Q_0 hard median.

6.4 Star-aware losses

Table 14 keeps the Q_0 skip and changes only the objective. Field, identity, and mixed-teacher rows are seed 0. Field-only and field-plus-Rusanov both leave the hard median at 0.068. Adding the soft identity $\mathcal{L}_{id}$ at weight 1 worsens the classic median and produces a worst-case 0.575. Adding a Newton-star teacher *together with* the field loss selects epoch 1 (validation median) and yields hard median 0.175. The same teacher with the field and PDE weights set to zero ($\mathcal{L}_{star}$ only) reaches classic / hard 0.000/0.000 on the nine anchors at seed 0 by correcting the right blast ($0.034 \rightarrow 3 \times 10^{-5}$). Five seeds give classic / hard 0.0001 ± 0.0001 / 0.002 ± 0.004 ; every seed beats the Q_0 hard median. Collision is not robust (0.013 ± 0.020 , range 3×10^{-5} to 0.047). Scores below the 0.027 oracle floor are shock-side assignment artefacts, not a solver that beats Newton (Supplementary Appendix C). That run uses the Newton root as a training label only; inference still evaluates $\mathcal{S} \circ H_\theta$. It is not a field-objective residual.

Table 14 Loss ablation on the Q_0 skip, nine anchors. Field / identity / mixed-teacher rows are seed 0. Star-sup. only is a Newton-star teacher with no field or Rusanov term.

Objective	Classic	Hard	Worst
Field + Rusanov	0.014	0.068	0.083
Field only	0.002	0.068	0.070
Field + $\mathcal{L}_{id}$	0.044	0.048	0.575
Field + Newton-star teacher	0.042	0.175	0.204
Newton-star teacher only	0.000	0.000	0.047

Role of Newton-star supervision. The Newton-star-only run is a control, not a competing solver. It establishes four facts. The residual parameterisation can represent useful corrections (five-seed hard median 0.002 ± 0.004 versus Q_0 0.034). Direct star targets contain information that field losses do not convert reliably into better corrections. The comparison is not a fair deployable baseline, because it assumes Newton-root labels that are cheap in this exact 1D setting. The same map does not exceed the Newton oracle on in-family, detuned-family, or Monte Carlo draws, and collision scores remain sensitive to one-cell assignment. A learned

substitute would matter only if amortized over a much more expensive downstream workflow or embedded in a broader differentiable computation.

6.5 Replacement physics objectives

Table 14 changes terms inside the field-loss experiment or uses a Newton-star teacher. A separate, pre-registered test is reported here. It asks whether a single physics objective can improve Q_0 when the field loss does not.

Everything except the objective is fixed. The network H_θ still corrects (p_*, u_*) around the same six-pass Q_0 starter. The sampler $\mathcal{S}$ is unchanged. The test uses the same locked split, exact Toro labels, inactive time horizon, 128×33 grid, Adam cosine schedule, 40 epochs, and validation overall-median checkpoint. The protocol was written before training. Each recipe uses only its named objective. It does not use the field loss or a Newton-star teacher. Seeds 0 to 9 are evaluated. Checkpoint selection never uses the nine classic anchors.

The three replacement objectives are: (i) the star identity in Eq. (31); (ii) the squared positive part of the entropy inequality $\partial_t \eta + \partial_x(u\eta)$, where $\eta = -\rho(\ln p - \gamma \ln \rho)$; and (iii) a trapezoid space-time control-volume residual of $\partial_t \mathbf{U} + \partial_x \mathbf{F} = 0$. The star identity is an algebraic Euler residual. It is not an exact Newton-star label. The entropy and control-volume objectives are computed from the sampled field. They are not a dual-norm weak residual or an embedded Godunov flux [8, 10].

Table 15 and Figure 8 show the results. The reported 95% intervals use a normal approximation over the ten seed-level medians; with $n = 10$ that interval is only a rough scale of seed-to-seed spread, not a formal decision test. The paired value is $\text{Rel-}L^2(\text{recipe}) - \text{Rel-}L^2(Q_0)$. A fully negative interval would indicate a consistent seed-wise gain.

Frozen Q_0 has classic / hard medians $8 \times 10^{-6}/0.034$. The field loss again walks away from Q_0 . Its paired hard interval is fully positive. The control-volume loss is also clearly worse on the hard cases. The entropy-only loss is unstable. Finite-difference entropy spikes near captured shocks can be reduced by changing the star state in the wrong direction. A checkpoint can then look acceptable on soft validation cases but fail badly on the hard cases.

The star-identity loss is the closest replacement. Two of ten seeds are better than Q_0 on the hard median. However, its paired interval includes zero, its classic median is worse, and some seeds give a large worst-case error. It is therefore not a deployable improvement under this protocol. This is also different from the Newton-teacher control.

The result is narrow but useful. Under the locked overall-median selection protocol used for this objective-comparison experiment, none of the four tested objectives yielded a dependable hard-anchor improvement over Q_0 . Mixed weights, dual-norm weak residuals, and other entropy treatments remain untested. Section 6.6 tests whether this conclusion holds after a γ change and reports a near-vacuum control. Section 5.4 embeds the same stars in a first-order Godunov-type flux; mixed objectives and HLLC remain untested.

Table 15 Sole-objective residuals around Q_0 on the nine classic anchors, using seeds 0–9. Entries are the ten-seed mean with a normal-approximation 95% interval. Paired Δ is the hard median of the recipe minus frozen Q_0 ; negative is better. Checkpoint selection uses the validation overall median only.

Objective	Classic mean [95% CI]	Hard mean [95% CI]	Paired Δ hard [95% CI]	Better seeds
Frozen Q_0	8×10^{-6}	0.034	—	—
Field $\mathcal{L}_{\log}$	0.016 [0.008, 0.024]	0.075 [0.049, 0.101]	+0.041 [+0.014, +0.067]	1/10
Star identity (31)	0.055 [0.029, 0.080]	0.050 [0.027, 0.073]	+0.015 [−0.008, +0.038]	2/10
Entropy inequality	0.155 [0.107, 0.202]	356 [158, 555]	+356 [+158, +555]	0/10
Control volume	0.038 [0.026, 0.050]	0.391 [0.377, 0.406]	+0.357 [+0.343, +0.371]	0/10




Fig. 8 Ten-seed means and normal-approximation 95% intervals for sole-objective residuals around Q_0 . The dashed line shows frozen Q_0 . The entropy result is omitted from the hard panel because its hard mean is 356 (95% interval [158, 555]).

6.6 Tests at new γ values and a near-vacuum control

The four-loss comparison uses only $\gamma = 1.4$. A second pre-registered study has two parts. The first recreates the nine families at a new γ and asks whether the field residual can reduce the remaining Q_0 error. The two γ retrains were specified before training. The second part is a near-vacuum control. When Q_0 is already at the accuracy floor, there is no measurable error for the residual to reduce. This control is therefore not an equally demanding error-reduction test.

The map, split ICs, exact labels, inactive T , 128×33 grid, Adam cosine schedule, and overall-median checkpoint are unchanged. Only γ , or the extra control family, changes. Seeds are 0–9. Family 10 is not included in the published hard set $\{3, 4, 5\}$. The same primitive ICs are kept when γ changes, so family 6 is no longer an RH-stationary shock.

The control family is a stronger double rarefaction than Toro 123: $(\rho, u, p)_L = (1, -3, 0.4)$ and $(\rho, u, p)_R = (1, 3, 0.4)$. The Newton star pressure is about $400\times$ smaller than Toro 123 and remains finite. It is not a vacuum-limit redraw of family 2.

Table 16 and Figure 9 give the ten-seed results. At $\gamma = 1.2$, the frozen Q_0 hard error grows to 0.053 (left blast 0.083, collision 0.053). Zero-shot evaluation of the published $\gamma = 1.4$ field maps does not reduce it (hard 0.087 ± 0.061 ; two seeds of ten beat Q_0 ; the paired interval includes zero). Retraining at $\gamma = 1.2$ gives a clearer negative result (hard 0.111 ± 0.051 ; paired $\Delta = +0.059$ with interval $[+0.027, +0.090]$; no seed beats Q_0). At $\gamma = 5/3$, Q_0 is almost exact on the hard median (0.0007). The residual is worse both zero-shot and after retrain (retrain hard 0.026 ± 0.011 ; paired $+0.026$ with interval $[+0.019, +0.033]$; 0/10).

On the locked near-vacuum anchor, Q_0 already has $\text{Rel-}L^2 \sim 10^{-8}$. That is the accuracy floor of this grid and sampler. The two-rarefaction branch already resolves this case, so the residual has no measurable error to reduce. Published $\gamma = 1.4$ field maps and a ten-seed retrain that adds twelve in-family vacuum members stay slightly above that floor (4×10^{-4} and 1×10^{-4}). This is a small degradation on an already solved control, not a failed attempt to improve a hard case. Detuned vacuum members leave a probe median of 0.020 for Q_0 ; the residual does not reduce it. The vacuum retrain also leaves the original nine-anchor hard median unimproved (paired $\Delta = +0.015$ with interval $[-0.002, +0.032]$; two seeds better).

The two γ retrains therefore give no hard-family gain over Q_0 . The near-vacuum row is a control, not a third hard-error test. The study still does not test mixed weights or a dual-norm weak residual. A first-order Godunov-type embedding of the same stars is Section 5.4.

Table 16 Field residual around Q_0 , seeds 0–9. The two γ rows test whether the residual can reduce the remaining Q_0 error. The vacuum row is a control at the Q_0 accuracy floor, not an equally demanding error-reduction test. Eval uses the published $\gamma = 1.4$ field checkpoints. Retrain uses the same locked rule on recreated labels. Paired Δ on the γ rows is hard Rel- L^2 minus frozen Q_0 .

Setting	Q_0	Field eval	Field retrain	Retrain Δ	Beats
$\gamma = 1.2$ hard	0.053	0.087 ± 0.061	0.111 ± 0.051	$+0.059$ [$+0.027, +0.090$]	0/10
$\gamma = 5/3$ hard	0.0007	0.036 ± 0.029	0.026 ± 0.011	$+0.026$ [$+0.019, +0.033$]	0/10
Near-vacuum control	$\sim 10^{-8}$	0.0004 ± 0.0004	0.0001 ± 0.0001	control	—




Fig. 9 Ten-seed field residual around Q_0 . Left: tests at $\gamma = 1.2$ and $\gamma = 5/3$ of whether the residual reduces the remaining Q_0 error. Right: near-vacuum control. Eval: published $\gamma = 1.4$ maps. Retrain: new labels, same overall-median rule. Error bars are normal-approximation 95% intervals. The vacuum panel is logarithmic because Q_0 is already at the accuracy floor ($\sim 10^{-8}$).

6.7 Why field training moves away from Q_0

The first gradient of the log-field loss at the frozen Q_0 star is not anti-aligned with a Newton teacher (Supplementary Table 4 and Figure 1). On every non-contact anchor the two-vector cosine is at least 0.93. Alignment with the Newton direction is nevertheless poor on the blasts for both losses (left blast ≈ 0.19 ; right blast ≈ 0.40). Walk-off is therefore a trajectory, mixed-batch, later-iterate, and validation-selection effect, not a local sign flip at initialisation. Finite-difference checks of $\mathcal{S}$ are Supplementary Table 9.

6.8 Broader Riemann draws and Monte Carlo

On the locked in-family and detuned-family members, Q_0 is already close to the Newton-star oracle under paired median metrics (Supplementary Tables 5–6 and Figure 2). The same seed-0 maps on the locked 1800-member Monte Carlo draw (seed 20260819) leave Q_0 slightly better than the field residual (Supplementary Table 7 and Figures 3–4). Those draws are wider local perturbations of the same nine ideal-gas regimes, not a new PDE class.

6.9 Forward and reverse-mode cost

Table 17 reports a batched map of the nine anchors (128×33) on the tested Apple M2 CPU. No practical forward advantage is observed in this 1D batch: Q_0 and the Newton-star oracle have essentially the same forward time, and reconstruction dominates the star algebra. Reverse-mode

differentiation through the residual costs about three times Q_0 (1.37 versus 0.41 ms per case). These numbers replace the unbatched sequential milliseconds of Table 4 for any cost-aware claim. They are not a GPU throughput study.

Table 17 Batched nine-anchor cost, including reconstruction. Times are milliseconds per case on an Apple M2 CPU. Backward differentiates a field MSE through $\mathcal{S}$ (and through H_θ for the residual).

Method	CPU fwd	CPU bwd
Q_0	0.25	0.41
PVRS only	0.26	0.39
Two-shock, $N = 1$	0.23	0.35
Exact Newton star	0.24	0.37
Residual on Q_0	0.30	1.37

Nine-anchor density overlays of exact, Q_0 , and the seed-0 residual at $t = T$ are Supplementary Figure 5.

7 Conclusion

The residual parameterization can represent useful star-state corrections, as shown by the weak-PVRS and direct-star controls. Nevertheless, under the locked overall-median selection protocol used for the objective-comparison study, none of four tested field- and physics-based objectives yields a dependable hard-anchor improvement over the stronger Q_0 starter. Full-epoch validation eliminates the earlier sparse-logging epoch-1 artifact. For the field-residual objective, a predeclared composite selector reduces average degradation and identifies two practically improved seed trajectories, but eight of ten selected models remain worse than Q_0 , the median seed-wise paired difference remains positive, and the bootstrap interval for the mean paired difference includes zero. Thus, hard-aware selection improves model selection behavior without making the tested residual-learning procedure reliably beneficial.

The first-order finite-volume embedding of Section 5.4 shows only that donor-cell truncation masks the evaluated star-map differences at 100–400 cells. The first-order finite-volume result should be interpreted only as a truncation-dominated masking diagnostic: it does not test a higher-order learned flux construction, and it does not establish finite-volume equivalence among the star maps. A higher-order flux embedding is needed to determine whether occasional local star-state gains persist at integrated CFD-solution level. The result is robust to the reported front-aware checks: the representative field-residual map remains worse than Q_0 under primitive, conserved, and plateau metrics. The durable contribution is therefore a controlled benchmarking protocol and a diagnostic characterization, not a replacement numerical flux or a general claim about physics-informed learning.

The conclusions concern this two-scalar residual, the specified one-dimensional ideal-gas Euler family suite, the tested objectives, and the two predeclared checkpoint-selection rules. They do not establish improved finite-volume flux performance, multidimensional interface treatment, non-ideal-EOS performance, or general learned operator performance.

Statements and Declarations

Funding. The author did not receive support from any organization for the submitted work.

Competing Interests. The author has no relevant financial or non-financial interests to disclose.

Author Contributions. Ramesh Kolluru: conceptualization; methodology; software; validation; formal analysis; investigation; data curation; visualization; writing—original draft; writing—review and editing.

Data Availability. A reproducibility archive is available at submission and will be permanently archived upon publication. It contains the protocol files and hashes, locked data split, deterministic label-generation code, training configurations, seed lists, raw per-epoch validation logs, selected checkpoints, raw holdout metrics, bootstrap script and random seed ($B = 10,000$, seed 20260910), figure and table generation scripts, the Godunov embedding script, and automated tests. The study uses Python 3.11.14 and PyTorch 2.3.0 on an Apple M2 system; the archive includes the complete environment specification. Unless a caption says otherwise, reported seeds are $\{0, 1, 2, 3, 4\}$. The objective comparison, γ / near-vacuum study, and confirmatory selection study use $\{0, \dots, 9\}$.

Ethics approval. Not applicable.

References

- [1] P. L. Roe. Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, 43(2):357–372, 1981. doi: 10.1016/0021-9991(81)90128-5.
- [2] Amiram Harten, Peter D. Lax, and Bram van Leer. On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM Review*, 25(1):35–61, 1983. doi: 10.1137/1025002.
- [3] Bernd Einfeldt, Claus-Dieter Munz, P. L. Roe, and Björn Sjögreen. On Godunov-type methods near low densities. *Journal of Computational Physics*, 92(2):273–295, 1991. doi: 10.1016/0021-9991(91)90211-3.
- [4] E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4:25–34, 1994. doi: 10.1007/BF01414629.
- [5] Eleuterio F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer, Berlin, Heidelberg, 3 edition, 2009.
- [6] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.
- [7] 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*, 2024.
- [8] Ameya D. Jagtap, Ehsan Kharazmi, and George Em Karniadakis. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020. doi: 10.1016/j.cma.2020.113028.
- [9] Ravi G. Patel, Indu Manickam, Nathaniel A. Trask, Mitchell A. Wood, Myoungkyu Lee, Ignacio Tomas, and Eric C. Cyr. Thermodynamically consistent physics-informed neural networks for hyperbolic systems. *Journal of Computational Physics*, 449:110754, 2022. doi: 10.1016/j.jcp.2021.110754.

- [10] Tim De Ryck, Siddhartha Mishra, and Roberto Molinaro. wPINNs: Weak physics-informed neural networks for approximating entropy solutions of hyperbolic conservation laws. *SIAM Journal on Numerical Analysis*, 62(2):811–841, 2024. doi: 10.1137/22M1522504.
- [11] Aidan Chaumet and Jan Giesselmann. Improving weak PINNs for hyperbolic conservation laws: Dual norm computation, boundary conditions and systems. *SMAI Journal of Computational Mathematics*, 10:373–401, 2024. doi: 10.5802/smai-jcm.116.
- [12] Jiahao Liu, Supei Zheng, Xueli Song, and Doudou Xu. Locally linearized physics-informed neural networks for Riemann problems of hyperbolic conservation laws. *Physics of Fluids*, 36(11):116135, 2024. doi: 10.1063/5.0238865.
- [13] Tim De Ryck and Siddhartha Mishra. Numerical analysis of physics-informed neural networks and related models in physics-informed machine learning. *Acta Numerica*, 33: 633–713, 2024. doi: 10.1017/S0962492923000089.
- [14] Ramesh Kolluru, N. V. Raghavendra, S. V. Raghurama Rao, and G. N. Sekhar. Simple and robust contact-discontinuity capturing central algorithms for high speed compressible flows. *Applied Mathematics and Computation*, 414:126660, 2022. ISSN 0096-3003. doi: 10.1016/j.amc.2021.126660. URL <https://www.sciencedirect.com/science/article/pii/S009630032100744X>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [JOURNAL1DEULERPINOSI.pdf](#)