

A well-balanced shock-capturing variational physics-informed neural method for shallow-water flow with topography

Sacha Cardonna¹ & Ali Haidar^{2,*}

¹Institut Montpelliérain Alexander Grothendieck (IMAG), UMR 5149, Université de Montpellier,
Place Eugène Bataillon, 34095 Montpellier Cedex 5, France

²Engineering College, Department of Mathematics, American University of the Middle East (AUM), Egaila, Kuwait

*Corresponding author, email: ali.haidar@aum.edu.kw

Abstract We introduce a variational physics-informed neural discretization for the nonlinear shallow-water equations with topography, in which a single global neural network represents the space–time trial state while the governing equations are enforced weakly in space through local cell moments and strongly in time at selected time nodes. Written in the free-surface variables $(\eta, \mathbf{q})$, the output map preserves the non-negative water depth by construction, inherits the initial condition without a penalty term, and embeds wet and wet–dry lakes at rest in the same trial manifold.

For shock-forming calculations, the conservative variational residual is supplemented with a one-sided mechanical-entropy loss, which penalizes only non-admissible positive energy production, and with a localized viscosity that concentrates dissipation near steep transitions. We establish two scheme-specific analytical results: an exact identity for the well-balanced defect generated by the positivity map, yielding distributional preservation for the positive part while quantifying the loss floor induced by a numerical film or a softplus map, and a two-regime consistency estimate for the nonlinear viscosity, which is $O(h^2)$ in smooth regions, reaches the $O(h)$ scale near an $O(h)$ transition layer, and vanishes in distributions along refined sequences with uniformly bounded spatial variation. Wet- and dry-bed dam breaks, fully wet and partially dry equilibria, and one- and two-dimensional moving-shoreline solutions then assess these mechanisms numerically.

Keywords Shallow-water balance laws; variational physics-informed neural networks; space–time neural trial manifolds; local Petrov–Galerkin residuals; entropy solutions; well-balancing; positivity; wet–dry interfaces.

1. Introduction

The nonlinear shallow-water equations are a prototype of hyperbolic balance laws for which several individually familiar requirements must be satisfied simultaneously. A useful approximation should conserve mass and momentum, preserve the admissible set $\eta \geq b$, maintain the cancellation defining a lake at rest, select entropy solutions after shock formation, and remain meaningful when the shoreline moves. Classical well-balanced finite-volume and discontinuous Galerkin constructions show that these requirements are tightly coupled: a positivity correction may destroy an exact steady state, while an inconsistent flux–source balance may generate waves whose amplitude exceeds a near-equilibrium perturbation [26–30, 34–37]. Wetting and drying make the issue sharper because strict hyperbolicity and the quotient $\mathbf{q}/H$ degenerate at vacuum. Recent subcell DG/FV strategies make the same coupling explicit by combining shock robustness, positivity and equilibrium preservation at the subcell level [31–33]. Physics-informed neural networks approximate the solution and enforce a differential operator through the training functional [1]. At a shock, the pointwise differential residual ceases to provide a natural consistency criterion for this global nonlinear approximation space. Conservative domain decompositions and variational PINNs replace part of the pointwise enforcement by integral information [10, 11, 19]; weak PINNs go further by using entropy formulations adapted to scalar conservation laws [20]. Rigorous PINN error estimates also emphasize that a small training error becomes informative only through stability and sampling arguments for the underlying PDE [21]. For the shallow-water system, positivity, well-balancing and wet–dry compatibility must consequently be encoded in the trial space and the discrete residual.

Michel-Dansac and coauthors have developed a complementary numerical-analysis viewpoint in which neural approximations are coupled to classical discretizations. Parametric PINN priors have been used to enrich discontinuous Galerkin bases for hyperbolic balance laws, with convergence estimates that quantify the influence of the prior [14]. The same principle has been extended to continuous finite-element spaces [15], while finite-basis PINNs with overlapping domain decomposition have been studied for the Helmholtz equation [16]. For transport-dominated equations, a Lagrangian construction based on neural characteristic maps restores the flow composition structure and avoids representing every transported scale directly in the network state [17]. These contributions place the neural approximation inside a discretization whose approximation space, decomposition and error mechanisms are stated explicitly.

Neural solvers for the shallow-water equations have been studied for flows on the sphere, one- and two-dimensional topographic configurations, flood propagation and tsunami inundation [5–9]. These works establish the relevance of physics-informed approximations in several geometries and flow regimes. To the best of our knowledge, a single shallow-water neural discretization has not yet combined a conservative variational residual, a positivity-compatible wet–dry state map and exact inheritance of lakes at rest with one-sided entropy discrimination and localized shock stabilization. We construct such a method and assess it on shocks, steady equilibria and moving shorelines in one and two space dimensions.

The numerical method comprises the network, the physical output map, the local test space, the quadrature rule and the nonlinear solver. We write the governing equations once in $d \in \{1, 2\}$ dimensions and use the pre-balanced free-surface variables introduced for well-balanced shallow-water discretizations [34]. The neural network defines a global space–time trial function, and the discretization is weak in space and strong in time: the residual is projected locally onto affine functions in one dimension and their bilinear tensor products in two dimensions, while time remains a collocation variable. This mixed viewpoint is related to spatial variational neural methods [11, 12, 19] and to neural Galerkin integrators for evolution equations [13]. The optimization acts on one global-in-time trial function over the complete interval. The term variational refers to the local spatial moments; time is treated by collocation and Gauss quadrature is restricted to the spatial integrals.

The construction starts from an admissible trial manifold. Its canonical map enforces non-negative height and vacuum-compatible discharge, while multiplicative carriers impose the initial data and include prescribed lakes at rest as exact trial states. The PDE operator is then measured by normalized local mass and momentum moments, which remain meaningful even when the strong derivative becomes large inside a shock layer. For nonsmooth solutions, the conservative moments are supplemented with the positive part of the mechanical-energy production and a height-gradient viscosity. These two terms provide, respectively, a one-sided entropy discriminator and a consistent localized stabilization.

This separation is essential because quadrature and test functions affect a VPINN independently of the expressivity of the neural trial set [18]. The analysis distinguishes algebraic identities of the output map, estimates for the quadrature-defined loss and consistency with a weak or entropy solution under refinement. The first two follow from the discrete construction. Convergence to an entropy solution additionally requires uniform compactness estimates as the spatial moments and trial spaces are refined.

We establish two analytical properties of the VPINN construction. We derive the exact residual created when the inherited lake-at-rest state is passed through an arbitrary scalar positivity map. This identity both proves wet–dry well-balancing for the positive part and exposes the non-zero loss floor created by a numerical film. It also yields an impossibility result: within this trial construction, smooth strict positivity and exact universal wet–dry well-balancing cannot both be obtained from a scalar regularization. For the solution-dependent shock viscosity, we prove the smooth-layer and shock-layer scalings and its distributional disappearance under a uniform bound on the spatial variation. Together, these results quantify the interaction between the admissible trial manifold and the residual stabilization.

2. Continuous model

The continuous problem is formulated in free-surface variables together with the regularity assumptions, weak balance law, steady states and entropy inequality used by the variational discretization.

2.1. Notation and functional setting

The domain, topography and state are described in the following functional setting, which also supports the distributional arguments of Sections 5 and 6.

If $D \subset \mathbb{R}^k$ is open, $\mathcal{C}^r(D; \mathbb{R}^m)$ and $\mathcal{C}_c^r(D; \mathbb{R}^m)$ denote, respectively, the spaces of r -times continuously differentiable functions and compactly supported functions of that class. We use the usual Lebesgue and Sobolev spaces $L^p(D; \mathbb{R}^m)$ and $W^{1,p}(D; \mathbb{R}^m)$. The distribution space is $\mathcal{D}'(D; \mathbb{R}^m) := (\mathcal{C}_c^\infty(D; \mathbb{R}^m))'$ and its duality pairing is denoted by $\langle \cdot, \cdot \rangle$. Moreover, $BV(D; \mathbb{R}^m)$ is the set of functions in $L^1(D; \mathbb{R}^m)$ whose distributional derivative is a finite Radon measure. Its variation seminorm is

$$|\mathbf{w}|_{BV(D)} := \sup \left\{ \int_D \mathbf{w} \cdot (\nabla \cdot \Phi) \, d\mathbf{x} \mid \Phi \in \mathcal{C}_c^1(D; \mathbb{R}^{m \times k}), \|\Phi\|_{L^\infty(D)} \leq 1 \right\},$$

where $\nabla \cdot \Phi$ acts rowwise. For a scalar z , we write $z_+ := \max(z, 0)$ and $\mathbb{R}_+ := [0, \infty)$. Unless a time derivative is shown explicitly, ∇ and $\nabla \cdot$ act only on the spatial variables. Finally, $A : B = \text{tr}(A^\top B)$ denotes the Frobenius product.

Let $d \in \{1, 2\}$, let $\Omega \subset \mathbb{R}^d$ be a bounded Lipschitz domain and set $\mathbb{Q} = \Omega \times (0, T)$ for a final time $T > 0$. The prescribed bottom elevation is assumed smooth, $b \in \mathcal{C}^\infty(\bar{\Omega})$, and the gravitational acceleration is denoted by $g > 0$. Since Ω is bounded, this assumption implies $b \in W^{1,\infty}(\Omega)$, which is sufficient for the balance identities and covers

the compactly supported $\mathcal{C}^1$ benchmark profiles used in the numerical section. We denote by H the water height, by $\mathbf{u} \in \mathbb{R}^d$ the depth-averaged velocity, by $\mathbf{q} = H\mathbf{u}$ the discharge and by $\eta = H + b$ the free-surface elevation. Thus

$$\mathbf{v} = (\eta, \mathbf{q})^\top, \quad H = \eta - b, \quad \mathbf{u} = \frac{\mathbf{q}}{H} \quad \text{on } \{H > 0\}. \quad (1)$$

The physical state set is

$$\mathcal{A}_b(\mathbf{x}) = \{(\eta, \mathbf{q}) \in \mathbb{R}^{d+1} \mid \eta > b(\mathbf{x})\} \cup \{(b(\mathbf{x}), \mathbf{0})\}.$$

At $H = 0$ we use the standard finite-energy convention $\mathbf{q} = \mathbf{0}$ and extend the kinetic energy by

$$\mathcal{K}(H, \mathbf{q}) := \begin{cases} |\mathbf{q}|^2/(2H), & H > 0, \\ 0, & (H, \mathbf{q}) = (0, \mathbf{0}), \\ +\infty, & H = 0 \text{ and } \mathbf{q} \neq \mathbf{0}. \end{cases}$$

Likewise, $\mathbf{q} \otimes \mathbf{q}/H$ is set to zero at the admissible vacuum state. $\mathbf{u}$ is set to $\mathbf{0}$ there. The canonical neural output map enforces this compatibility, and each alternative carrier is defined together with its admissibility properties.

2.2. The nonlinear shallow-water system

The shallow-water system is written in a dimension-independent form and then expressed in the pre-balanced free-surface variables used by the VPINN residual.

In height variables, the shallow-water system reads

$$\begin{cases} \partial_t H + \nabla \cdot \mathbf{q} = 0, \\ \partial_t \mathbf{q} + \nabla \cdot \left(\frac{\mathbf{q} \otimes \mathbf{q}}{H} + \frac{g}{2} H^2 \mathbb{I}_d \right) = -gH \nabla_x b. \end{cases} \quad (2)$$

Because b is independent of time, the mass equation is unchanged when H is replaced by η . Expanding

$$\frac{g}{2} \nabla_x (\eta^2 - 2\eta b) + g\eta \nabla_x b = g(\eta - b) \nabla_x \eta = gH \nabla_x (H + b)$$

shows that (2) is equivalent, on every wet smooth region, to the pre-balanced system

$$\partial_t \mathbf{v} + \nabla \cdot \mathbb{F}(\mathbf{v}, b) = \mathbf{S}(\mathbf{v}, b), \quad (3)$$

where

$$\mathbb{F}(\mathbf{v}, b) = \left(\mathbf{q}, \frac{\mathbf{q} \otimes \mathbf{q}}{H} + \frac{g}{2} (\eta^2 - 2\eta b) \mathbb{I}_d \right)^\top, \quad \mathbf{S}(\mathbf{v}, b) = (0, -g\eta \nabla_x b)^\top. \quad (4)$$

Here $\mathbb{I}_d$ is the identity matrix in $\mathbb{R}^d$, $\mathbb{F}$ is the physical flux and $\mathbf{S}$ is the bathymetric source. This is the single model used in both dimensions. When $d = 1$, bold vectors have one component.

2.3. Motionless steady states

For any constant level $\eta^c \in \mathbb{R}$, a lake-at-rest state is

$$\mathbf{q}^c = \mathbf{0}, \quad \eta_{\text{eq}}(\mathbf{x}) = b(\mathbf{x}) + (\eta^c - b(\mathbf{x}))_+.$$

In a wet region, $\eta_{\text{eq}} = \eta^c$ and the pressure-flux gradient balances the bathymetric source exactly. In a dry region, $\eta_{\text{eq}} = b$ and the shoreline is the free boundary $\{b = \eta^c\}$. The pre-balanced flux is useful because this cancellation can be tested without separately approximating ∇H and ∇b [34].

2.4. Weak solutions and entropy inequality

For $\mathbf{v} \in L^1_{\text{loc}}(\mathbb{Q}; \mathbb{R}^{d+1})$ with locally integrable flux and source, the balance law is understood in $\mathcal{D}'(\mathbb{Q}; \mathbb{R}^{d+1})$. Namely, for every $\varphi \in \mathcal{C}_c^\infty(\mathbb{Q}; \mathbb{R}^{d+1})$, a weak solution of (3) satisfies

$$\int_{\mathbb{Q}} (\mathbf{v} \cdot \partial_t \varphi + \mathbb{F}(\mathbf{v}, b) : \nabla_x \varphi + \mathbf{S}(\mathbf{v}, b) \cdot \varphi) \, d\mathbf{x} \, dt = 0,$$

This identity is exactly the statement $\partial_t \mathbf{v} + \nabla \cdot \mathbb{F}(\mathbf{v}, b) - \mathbf{S}(\mathbf{v}, b) = \mathbf{0}$ in $\mathcal{D}'(\mathbb{Q}; \mathbb{R}^{d+1})$. The usual initial and boundary terms are recovered when test functions are allowed to reach the parabolic boundary. Across shocks, conservation is understood through this distributional identity.

For smooth solutions, the mechanical energy and energy flux are

$$\mathcal{E}(H, \mathbf{q}, b) = \mathcal{K}(H, \mathbf{q}) + \frac{g}{2} H^2 + g H b, \quad \mathbf{F}_{\mathcal{E}} = \mathbf{u} \left(\mathcal{E} + \frac{g}{2} H^2 \right).$$

Whenever $\mathcal{E}$ and $\mathbf{F}_{\mathcal{E}}$ are locally integrable, an admissible weak solution satisfies $\partial_t \mathcal{E} + \nabla \cdot \mathbf{F}_{\mathcal{E}} \leq 0$ in distributions. Equivalently, for every non-negative $\zeta \in \mathcal{C}_c^\infty(\mathbb{Q})$,

$$\langle \partial_t \mathcal{E} + \nabla \cdot \mathbf{F}_{\mathcal{E}}, \zeta \rangle = - \int_{\mathbb{Q}} (\mathcal{E} \partial_t \zeta + \mathbf{F}_{\mathcal{E}} \cdot \nabla \zeta) \, d\mathbf{x} \, dt \leq 0.$$

This inequality supplies information not contained in conservation alone and motivates the one-sided entropy loss used for dam-break waves.

3. Space–time trial manifold and local test space

Let $\mathcal{T}_h = \{\omega_c\}_{c=1}^{N_{\text{el}}}$ be a partition of Ω into N_{el} cells. In one dimension, $\omega_c = (x_{c-\frac{1}{2}}, x_{c+\frac{1}{2}})$ may belong to a non-uniform mesh. In two dimensions,

$$\omega_{ij} = (x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}) \times (y_{j-\frac{1}{2}}, y_{j+\frac{1}{2}})$$

is a Cartesian cell. We set $h := \max_{\omega_c \in \mathcal{T}_h} \text{diam}(\omega_c)$. The neural trial state is global in space–time, while its residual is projected locally in space at selected time nodes.

We use the following local test functions. On a one-dimensional cell, with

$$\xi_c(x) := \frac{2x - x_{c-\frac{1}{2}} - x_{c+\frac{1}{2}}}{x_{c+\frac{1}{2}} - x_{c-\frac{1}{2}}},$$

they are the two affine functions

$$\psi_0^c(x) = \frac{1 - \xi_c(x)}{2}, \quad \psi_1^c(x) = \frac{1 + \xi_c(x)}{2}.$$

On a Cartesian cell, the four tests are the tensor products

$$\psi_{ab}^{ij}(x, y) = \psi_a^i(x) \psi_b^j(y), \quad (a, b) \in \{0, 1\}^2.$$

The scalar spatial test space is therefore

$$\mathcal{Y}_h := \text{span} \{ \psi_m^c \mathbf{1}_{\omega_c} \mid \omega_c \in \mathcal{T}_h, m \in \mathcal{M}_c \},$$

where $\mathcal{M}_c = \{0, 1\}$ in one dimension and $\mathcal{M}_c = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$ in two dimensions. Thus the residual is measured against fixed local linear or bilinear modes. The time interval is sampled at the nodes $0 = t^0 < \dots < t^{N_t} = T$, and we write

$$(f, g)_{\omega_c} := \int_{\omega_c} f g \, d\mathbf{x}.$$

Each residual degree of freedom is therefore indexed by a spatial cell, a local test and a time sample.

Global neural trial manifold

The trainable parameters belong to a finite-dimensional set $\Theta \subset \mathbb{R}^P$. For $\theta \in \Theta$, let $\mathcal{N}_\theta : \mathbb{Q} \rightarrow \mathbb{R}^{p_{\text{out}}}$ be a global feed-forward network of the standard multilayer-perceptron form used in neural PDE solvers [2]. The raw output dimension p_{out} is determined by the carrier and may differ from $d + 1$. The physical map always returns the state $\mathbf{v}_\theta \in \mathbb{R}^{d+1}$. Let $\chi : \mathbb{Q} \rightarrow \mathbb{R}^{p_{\text{in}}}$ be the prescribed input map, which contains normalized coordinates and, when appropriate, periodic features. For ℓ hidden layers of width n , set $n_0 = p_{\text{in}}$, $n_i = n$ for $i \in \{1, \dots, \ell\}$ and $n_{\ell+1} = p_{\text{out}}$. For each $i \in \{1, \dots, \ell + 1\}$, the trainable weight matrix and bias vector are $\mathbf{W}_i \in \mathbb{R}^{n_i \times n_{i-1}}$ and $\mathbf{b}_i \in \mathbb{R}^{n_i}$; the vector θ collects all these entries, so that $P = \sum_{i \in \{1, \dots, \ell+1\}} n_i(n_{i-1} + 1)$. With activation σ , the network layers are

$$\begin{aligned} \mathbf{z}_0(\mathbf{x}, t) &= \chi(\mathbf{x}, t), \\ \mathbf{z}_i(\mathbf{x}, t) &= \sigma(\mathbf{W}_i \mathbf{z}_{i-1}(\mathbf{x}, t) + \mathbf{b}_i), \quad i \in \llbracket 1, \ell \rrbracket, \\ \mathcal{N}_\theta(\mathbf{x}, t) &= \mathbf{W}_{\ell+1} \mathbf{z}_\ell(\mathbf{x}, t) + \mathbf{b}_{\ell+1}. \end{aligned}$$

The activation is applied componentwise; the calculations use smooth activations so that all derivatives entering the strong residual are obtained by automatic differentiation. The physical maps introduced in Section 4 transform the raw outputs into $\mathbf{v}_\theta = (\eta_\theta, \mathbf{q}_\theta)^\top$. The resulting trial set

$$\mathcal{X}_\Theta := \{\mathbf{v}_\theta \mid \theta \in \Theta, \mathbf{v}_\theta(\mathbf{x}, t) \in \mathcal{A}_b(\mathbf{x}), \forall (\mathbf{x}, t) \in \mathbb{Q}\}$$

is a nonlinear, globally continuous space–time manifold. The restriction $\mathbf{v}_\theta^c := \mathbf{v}_\theta|_{\omega_c \times (0, T)}$ supplies local residual moments from the same globally parametrized trial state. Case-dependent input features and carriers provide alternative coordinates for the same physical operator and are specified with each numerical experiment.

4. Admissible VPINN trial manifold

The physical constraints are enforced at the level of the trial manifold, before any residual is evaluated. We first construct a positive state that inherits the initial data and then specialize the same construction to known equilibria. Trial-function carriers have long been used to impose initial or boundary data exactly in neural PDE solvers [2, 3]. Non-negative output transformations and auxiliary positivity losses have both been examined in PINNs [4]; shallow-water PINNs have also treated non-negative depth and wet–dry compatibility through additional loss terms [6, 9]. Here the temporal carrier and the non-negative output map are combined directly in the trial manifold.

Hard inheritance of the initial data

Let $\hat{\eta}_\theta$ and $\hat{\mathbf{u}}_\theta$ denote raw network outputs, and let $\tau : [0, T] \rightarrow \mathbb{R}_+$ satisfy $\tau(0) = 0$. We set

$$\eta_\theta(\mathbf{x}, t) = b(\mathbf{x}) + \mathcal{P}_+(\eta^0(\mathbf{x}) - b(\mathbf{x}) + \tau(t)H_{\text{ref}}\hat{\eta}_\theta(\mathbf{x}, t)),$$

where $H_{\text{ref}} > 0$ is a characteristic height and $\mathcal{P}_+$ is a non-negative map. The computations use the exact positive part or $\mathcal{P}_+(z) = \max(z, H_\varepsilon)$ with a small numerical film $H_\varepsilon > 0$; the analysis also considers a softplus map to quantify the defect of a smooth strictly positive regularization. The value H_ε is the computational dry tolerance used to keep the flux quotients finite. When it is included in $\mathcal{P}_+$, an exactly dry point is represented by a layer of depth H_ε . The velocity and discharge are defined by

$$\mathbf{u}_\theta = \mathbf{u}^0 + \tau(t)U_{\text{ref}}\hat{\mathbf{u}}_\theta, \quad \mathbf{q}_\theta = (\eta_\theta - b)\mathbf{u}_\theta. \quad (5)$$

By construction, water height positivity is guaranteed, i.e. $\eta_\theta \geq b$ for every parameter vector. Assume that the prescribed initial state satisfies $\eta^0 \geq b$ and $\mathbf{q}^0 = (\eta^0 - b)\mathbf{u}^0$, with $\mathbf{u}^0 = \mathbf{0}$ at initially dry points. Since $\tau(0) = 0$, we have $\eta_\theta(\cdot, 0) = b + \mathcal{P}_+(\eta^0 - b)$ and $\mathbf{q}_\theta(\cdot, 0) = (\eta_\theta(\cdot, 0) - b)\mathbf{u}^0$. For the exact positive part $\mathcal{P}_+(z) = z_+ := \max(z, 0)$, the identity $\mathcal{P}_+(\eta^0 - b) = \eta^0 - b$ holds even where $\eta^0 = b$; hence $(\eta_\theta(\cdot, 0), \mathbf{q}_\theta(\cdot, 0)) = (\eta^0, \mathbf{q}^0)$ exactly and no initial-condition penalty is required. In contrast, the numerical-film choice $\mathcal{P}_+(z) = \max(z, H_\varepsilon)$ gives $\eta_\theta(\mathbf{x}, 0) = b(\mathbf{x}) + H_\varepsilon = \eta^0(\mathbf{x}) + H_\varepsilon$ at an initially dry point, while $\mathbf{q}_\theta(\mathbf{x}, 0) = \mathbf{0}$ there by the vacuum convention. More generally, the film changes the initial free surface by $(H_\varepsilon - (\eta^0 - b))_+ \leq H_\varepsilon$, giving a regularized approximation of the prescribed dry data with error at most H_ε . The carrier therefore preserves positivity in both cases, whereas exact initial-state inheritance at a dry point requires the unregularized positive part. This is the canonical admissible map. Some test cases use a bounded discharge carrier whose formula is given with the corresponding initial data.

Equilibrium inheritance

For a known lake at rest of level η^c , we incorporate the equilibrium into the trial map. We therefore write

$$\eta_\theta = b + \mathcal{P}_+((\eta^c - b)_+ + \tau(t)H_{\text{ref}}\hat{\eta}_\theta), \quad \mathbf{q}_\theta = (\eta_\theta - b)\tau(t)U_{\text{ref}}\hat{\mathbf{u}}_\theta. \quad (6)$$

For the exact positive part used in the lake-at-rest experiments, the zero correction belongs to the trial space and produces $\eta_\theta = b + (\eta^c - b)_+$ and $\mathbf{q}_\theta = \mathbf{0}$. Section 6 analyzes the corresponding zero-correction family when a general positivity map is applied directly to the level gap.

5. Local variational formulation

The trial state is global, while the governing equation is observed through a finite family of local spatial moments at selected time nodes. These moments, their physical normalization and the two shock-specific terms define the discrete objective.

5.1. Variational residual and normalized physical loss

This subsection specifies the Petrov–Galerkin part of the method. We distinguish exact spatial moments from their Gaussian-quadrature approximations.

Set

$$\mathfrak{R}(\mathbf{v}_\theta) := \partial_t \mathbf{v}_\theta + \nabla \cdot \mathbb{F}(\mathbf{v}_\theta, b) - \mathbf{S}(\mathbf{v}_\theta, b) = (\mathfrak{R}_\eta, \mathfrak{R}_\mathbf{q})^\top.$$

For every cell ω_c , test index $m \in \mathcal{M}_c$ and time index $n \in \{0, \dots, N_t\}$, the residual degrees of freedom are

$$\mathcal{R}_{\eta,c,m}^n(\theta) = (\mathfrak{R}_\eta(\mathbf{v}_\theta)(\cdot, t^n), \psi_m^c)_{\omega_c}, \quad \mathcal{R}_{\mathbf{q},c,m}^n(\theta) = (\mathfrak{R}_{\mathbf{q}}(\mathbf{v}_\theta)(\cdot, t^n), \psi_m^c)_{\omega_c}. \quad (7)$$

These quantities are spatial integrals at a fixed time. In particular, the formulation is weak in space in the weighted-residual sense and strong in time, with temporal collocation in (7). The spatial derivatives act on the neural trial state in the residual moments. For a sufficiently regular trial state, integration by parts exposes the cell-boundary flux:

$$(\mathfrak{R}(\mathbf{v}_\theta), \psi_m^c)_{\omega_c} = (\partial_t \mathbf{v}_\theta - \mathbf{S}, \psi_m^c)_{\omega_c} - (\mathbb{F}, \nabla \psi_m^c)_{\omega_c} + \int_{\partial \omega_c} \mathbb{F} \mathbf{n}_c \psi_m^c \, ds,$$

Here “weak in space” describes the local moment constraint, while evaluating the strong residual retains the spatial differentiability required of the neural output.

The derivatives remain on the neural trial function and the residual is measured through local moments. Since the two affine tests and the four bilinear tests form a partition of unity,

$$\sum_{m \in \mathcal{M}_c} \mathcal{R}_{\eta,c,m}^n = \int_{\omega_c} \mathfrak{R}_\eta(\mathbf{v}_\theta)(\mathbf{x}, t^n) \, d\mathbf{x}, \quad \sum_{m \in \mathcal{M}_c} \mathcal{R}_{\mathbf{q},c,m}^n = \int_{\omega_c} \mathfrak{R}_{\mathbf{q}}(\mathbf{v}_\theta)(\mathbf{x}, t^n) \, d\mathbf{x}.$$

Thus the mean cell balance is controlled by the same residual vector that is optimized. This identity follows from the partition of unity. It gives a computable local conservation defect and exact local conservation when the relevant moments vanish.

Quadrature and normalization

Let $\{(\mathbf{x}_j^c, w_j^c)\}_{j=1}^{N_q^c}$ be the affine image in ω_c of the tensor-product Gauss–Legendre rule on the reference cell, with the Jacobian included in the physical weights. Hence $\sum_{j=1}^{N_q^c} w_j^c = |\omega_c|$, and the quadrature approximation at a sampled time t^n is

$$\mathcal{R}_{c,m}^n(\theta) \simeq \sum_{j=1}^{N_q^c} w_j^c \mathfrak{R}(\mathbf{v}_\theta)(\mathbf{x}_j^c, t^n) \psi_m^c(\mathbf{x}_j^c).$$

The spatial cells, polynomial tests and quadrature nodes are fixed before training. The selection of time nodes and the automatic-differentiation procedure are described in Section 7.

Here T is the final physical time of the experiment, so $[0, T]$ is the training interval. Let H_{ref} and L_{ref} be characteristic height and length scales, set $C_{\text{ref}} = \sqrt{g H_{\text{ref}}}$, $Q_{\text{ref}} = H_{\text{ref}} C_{\text{ref}}$, and define

$$R_\eta = \frac{H_{\text{ref}}}{T} + \frac{Q_{\text{ref}}}{L_{\text{ref}}}, \quad R_q = \frac{Q_{\text{ref}}}{T} + \frac{Q_{\text{ref}} C_{\text{ref}}}{L_{\text{ref}}} + \frac{g H_{\text{ref}}^2}{L_{\text{ref}}}.$$

These fixed physical normalizations sum the characteristic magnitudes appearing in the mass and momentum residuals. The additional factor $|\omega_c|^{-1}$ removes the dependence of the normalized moment on cell volume. We therefore set

$$\widehat{\mathcal{R}}_{\eta,c,m}^n = \frac{\mathcal{R}_{\eta,c,m}^n}{|\omega_c| R_\eta}, \quad \widehat{\mathcal{R}}_{\mathbf{q},c,m}^n = \frac{\mathcal{R}_{\mathbf{q},c,m}^n}{|\omega_c| R_q}.$$

Let $\mathcal{N}_A \subset \{0, \dots, N_t\}$ be the time indices retained in one loss evaluation. Set $\alpha_1 = 1$, $\alpha_2 = 1/4$, and define $\varpi_c = |\omega_c|/|\Omega|$ on a one-dimensional mesh. In one dimension, the two squared affine moments are summed, which gives $\alpha_1 = 1$. In two dimensions, the four squared bilinear moments are averaged, which gives $\alpha_2 = 1/4$. On the uniform Cartesian meshes used below, set $\varpi_c = 1/N_{\text{el}}$. The physical loss is

$$\mathcal{L}_{\text{phy}}(\theta) = \frac{1}{|\mathcal{N}_A|} \sum_{n \in \mathcal{N}_A} \sum_{\omega_c \in \mathcal{T}_h} \varpi_c \alpha_d \sum_{m \in \mathcal{M}_c} \left(|\widehat{\mathcal{R}}_{\eta,c,m}^n|^2 + \|\widehat{\mathcal{R}}_{\mathbf{q},c,m}^n\|_2^2 \right). \quad (8)$$

On a non-uniform one-dimensional mesh, the cell-length weight prevents local refinement from changing the relative measure of a fixed physical region. The scales and cell weights remain fixed throughout one loss evaluation. The common scale R_q treats both Cartesian momentum components symmetrically in the 2D loss.

5.2. Entropy defect and localized shock viscosity

We use the mechanical energy pair introduced in Section 2. An entropy solution satisfies $\partial_t \mathcal{E} + \nabla \cdot \mathbf{F}_{\mathcal{E}} \leq 0$ in the distributional sense. At the same selected time slices and spatial quadrature points as the physical loss, the positive part of its normalized production is penalized,

$$\mathcal{L}_{\text{ent}} = \frac{1}{N_{\text{ent}}} \sum_{r \in \{1, \dots, N_{\text{ent}}\}} \left[\max \left(\frac{\partial_t \mathcal{E}_\theta + \nabla \cdot \mathbf{F}_{\mathcal{E}, \theta}}{R_{\mathcal{E}}}, 0 \right) \right]^2. \quad (9)$$

Here N_{ent} is the number of retained pairs formed by a selected time node and a spatial Gauss point. Time is collocated, and quadrature is applied only to the spatial integrals. On a rectangular domain with side lengths $L_1, \dots, L_d$, set

$$E_{\text{ref}} = \frac{3}{2} g H_{\text{ref}}^2, \quad R_{\mathcal{E}} = \frac{E_{\text{ref}}}{T} + C_{\text{ref}} E_{\text{ref}} \sum_{k \in \{1, \dots, d\}} \frac{1}{L_k}.$$

This scale is fixed before training. Only positive production is penalized: negative production is compatible with the entropy inequality and is not driven toward zero. The map $z \mapsto (z_+)^2$ is continuously differentiable, including at $z = 0$, and is therefore directly compatible with automatic differentiation. This term complements the conservative moments by discriminating against locally entropy-increasing states having comparable mass and momentum loss. It is activated only for the shock-forming tests. If a regular state has non-positive production almost everywhere, then the distributional entropy inequality of Section 2 follows by testing against non-negative functions. A zero value of (9) controls the entropy production at the retained points. Convergence to the entropy inequality over the continuum requires simultaneous refinement of the spatial quadrature and the temporal collocation set.

For dam-break waves, the residuals entering $\mathcal{L}_{\text{phy}}$ are replaced before projection by

$$\mathfrak{R}_\eta^\nu = \mathfrak{R}_\eta - \nabla \cdot (\nu_\theta \nabla (\eta_\theta - b)), \quad \mathfrak{R}_{\mathbf{q}}^\nu = \mathfrak{R}_{\mathbf{q}} - \nabla \cdot (\nu_\theta \nabla \mathbf{q}_\theta), \quad (10)$$

where the coefficient is defined through

$$G_\delta(r) := \sqrt{|\nabla r|^2 + \delta^2}, \quad \nu_\theta = C_\nu C_{\text{ref}} h_c \tanh \left(C_s h_c \frac{G_\delta(\eta_\theta - b)}{H_{\text{ref}}} \right). \quad (11)$$

The added terms are written in divergence form so that the stabilization remains a conservative spatial flux and admits the distributional estimate below. The mass equation diffuses the physical water height $\eta_\theta - b$, thereby avoiding artificial diffusion of the prescribed bathymetry, while the momentum flux is diffused componentwise. We take $\delta = 0$ in one dimension and $\delta = 10^{-6}$ in two dimensions. The parameter δ regularizes the Euclidean gradient norm at the origin. With $\delta = 0$ the sensor is locally Lipschitz and automatic differentiation uses the standard zero subgradient at a vanishing height gradient. The coefficient is negligible in smooth regions and bounded by $C_\nu C_{\text{ref}} h_c$ in steep layers, where h_c is the cell length in one dimension and the minimum Cartesian side in two dimensions. Diffusion is applied componentwise to $\eta_\theta - b$ and $\mathbf{q}_\theta$ before their residuals are tested. Hence the stabilization follows the learned wave and is independent of a prescribed shock location. In the experiments below, the coefficients are fixed before training and reported with each nonsmooth test; the viscosity is disabled for smooth and stationary experiments.

Proposition 1 (Consistency of the viscosity sensor). *Consider a sequence of uniform one-dimensional or Cartesian meshes of size $h \rightarrow 0$, and let ν_h be defined by (11). For the associated neural states $\mathbf{v}_h = \mathbf{v}_{\theta_h} = (\eta_h, \mathbf{q}_h)^\top$, set $\mathbf{z}_h = (\eta_h - b, \mathbf{q}_h)^\top$. Then*

$$0 \leq \nu_h \leq C_\nu C_{\text{ref}} \min \left\{ h, \frac{C_s}{H_{\text{ref}}} h^2 G_\delta(\eta_h - b) \right\}. \quad (12)$$

Consequently, $\|\nu_h\|_{L^\infty} = O(h^2)$ on any family of smooth regions where $\|G_\delta(\eta_h - b)\|_{L^\infty}$ is bounded. Suppose further that $\mathbf{z}_h \in L^1(0, T; W^{1,1}(\Omega; \mathbb{R}^{d+1}))$ and

$$M := \sup_{h>0} \int_{\mathbb{Q}} |\nabla_{\mathbf{x}} \mathbf{z}_h| \, d\mathbf{x} \, dt < \infty.$$

Then the componentwise viscous operator satisfies

$$|\langle \nabla \cdot (\nu_h \nabla \mathbf{z}_h), \boldsymbol{\varphi} \rangle| \leq C_\nu C_{\text{ref}} h M \|\nabla \boldsymbol{\varphi}\|_{L^\infty(\mathbb{Q})} \longrightarrow 0$$

for every $\boldsymbol{\varphi} \in \mathcal{C}_c^\infty(\mathbb{Q}; \mathbb{R}^{d+1})$. Thus the same sensor is second-order small in smooth regions, may reach the shock-capturing $O(h)$ scale in layers with $G_\delta(\eta_h - b) = O(h^{-1})$, and still disappears in $\mathcal{D}'(\mathbb{Q}; \mathbb{R}^{d+1})$.

Proof. Put $s_h = C_s h G_\delta(\eta_h - b)/H_{\text{ref}} \geq 0$. Since $0 \leq \tanh(s) \leq 1$ and $\tanh(s) \leq s$ for $s \geq 0$, we obtain

$$\nu_h \leq C_\nu C_{\text{ref}} h, \quad \nu_h \leq \frac{C_\nu C_s C_{\text{ref}}}{H_{\text{ref}}} h^2 G_\delta(\eta_h - b).$$

Their minimum is exactly (12). If $\|G_\delta(\eta_h - b)\|_{L^\infty(K)} \leq M_K$ on a smooth region K , the second bound gives $\|\nu_h\|_{L^\infty(K)} \leq C_\nu C_s C_{\text{ref}} M_K h^2 / H_{\text{ref}}$. In a transition layer of width proportional to h , one instead has $G_\delta(\eta_h - b) = O(h^{-1})$, hence $s_h = O(1)$ and the first bound allows the $O(h)$ shock-capturing scale.

Let $\varphi \in \mathcal{C}_c^\infty(\mathbb{Q}; \mathbb{R}^{d+1})$. By definition of the distributional divergence of the spatial flux,

$$\langle \nabla \cdot (\nu_h \nabla_{\mathbf{x}} \mathbf{z}_h), \varphi \rangle = - \int_{\mathbb{Q}} \nu_h \nabla_{\mathbf{x}} \mathbf{z}_h : \nabla_{\mathbf{x}} \varphi \, d\mathbf{x} \, dt,$$

where no boundary term occurs because φ is compactly supported. Consequently,

$$\begin{aligned} |\langle \nabla \cdot (\nu_h \nabla_{\mathbf{x}} \mathbf{z}_h), \varphi \rangle| &\leq \|\nu_h\|_{L^\infty(\mathbb{Q})} \int_{\mathbb{Q}} |\nabla_{\mathbf{x}} \mathbf{z}_h| \, d\mathbf{x} \, dt \|\nabla_{\mathbf{x}} \varphi\|_{L^\infty} \\ &\stackrel{(12)}{\leq} C_\nu C_{\text{ref}} h \int_{\mathbb{Q}} |\nabla_{\mathbf{x}} \mathbf{z}_h| \, d\mathbf{x} \, dt \|\nabla_{\mathbf{x}} \varphi\|_{L^\infty} \lesssim h \|\nabla_{\mathbf{x}} \varphi\|_{L^\infty}. \end{aligned}$$

The hidden constant in $\lesssim$ depends only on $C_\nu C_{\text{ref}} M$ and is independent of h . Hence the last expression tends to zero for every fixed test function, which is precisely convergence to zero in $\mathcal{D}'(\mathbb{Q}; \mathbb{R}^{d+1})$. $\square$

5.3. Boundary conditions and complete loss

The initial condition is already part of the trial manifold and is therefore absent from the objective in both dimensions. What remains is to impose the boundary operator and to collect the physical, entropy and boundary contributions into one normalized functional.

Let N_{bc} denote the number of boundary samples. Boundary losses are defined from an operator $\mathcal{B}(\mathbf{v}_\theta) = \mathbf{d}_{\text{bc}}$ on $\partial\Omega \times (0, T)$:

$$\mathcal{L}_{\text{bc}} = \frac{1}{N_{\text{bc}}} \sum_{r \in \{1, \dots, N_{\text{bc}}\}} \|\mathbb{D}_{\text{bc}}^{-1} [\mathcal{B}(\mathbf{v}_\theta)(\mathbf{x}_r, t_r) - \mathbf{d}_{\text{bc}}(\mathbf{x}_r, t_r)]\|_2^2.$$

The diagonal matrix $\mathbb{D}_{\text{bc}}$ contains H_{ref} for a height condition and Q_{ref} for a discharge condition. The samples cover the entire time interval. For the dam-break and well-balanced boxes, $\mathcal{B}$ is the reflective-wall operator $\mathbf{q}_\theta \cdot \mathbf{n} = 0$. In the parabolic-bowl tests, the square computational box lies outside every exact wet support; the same normal-discharge condition is supplemented by $\eta_\theta = b$ on its boundary. The Thacker test is periodic in time and uses the same wall boundary condition in space.

Complete training functional

The loss minimized in a run is

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{phy}}(\theta) + \lambda_{\text{ent}} \mathcal{L}_{\text{ent}}(\theta) + \lambda_{\text{bc}} \mathcal{L}_{\text{bc}}(\theta). \quad (13)$$

The physical term contains every mass and momentum component. The fixed coefficients used in one and two space dimensions are specified with the numerical settings. The shock viscosity enters $\mathcal{L}_{\text{phy}}$ through (10). The loss (13) depends exclusively on the governing equations, boundary conditions and prescribed initial state. Exact and finite-volume solutions are reserved for the subsequent error assessment.

6. Discrete structural properties of the VPINN

The following results characterize the algebraic properties of the neural state map and the equilibrium residual generated by its positivity transformation. The main result quantifies the exact equilibrium defect of a general scalar positivity map and leads to an obstruction for smooth strictly positive regularizations.

Well-balancing and the positivity regularization floor

We first isolate a property that is specific to the combination of the pre-balanced operator and the inherited positive trial map. Let $\mathcal{P} : \mathbb{R} \rightarrow \mathbb{R}_+$ be Lipschitz continuous and define the zero-correction state associated with a prescribed level $\eta^c \in \mathbb{R}$ by

$$\eta_{\mathcal{P}}^c(\mathbf{x}) := b(\mathbf{x}) + \mathcal{P}(\eta^c - b(\mathbf{x})), \quad \mathbf{q}_{\mathcal{P}}^c(\mathbf{x}) := \mathbf{0}. \quad (14)$$

This family isolates the zero-correction state obtained when the positivity map is applied directly to the unregularized level gap $\eta^c - b$. The positive-part member is exactly the lake-at-rest carrier used in the computations; the film and softplus members quantify two alternative regularizations of that carrier.

Theorem 1 (Well-balanced defect identity). *Let $b \in \mathcal{C}^\infty(\bar{\Omega})$ and use the convention that the kinetic flux vanishes when $H = 0$ and $\mathbf{q} = \mathbf{0}$. The proof uses only the consequence $b \in W^{1,\infty}(\Omega)$. The strong residual of the state $\mathbf{v}_{\mathcal{P}}^c = (\eta_{\mathcal{P}}^c, \mathbf{q}_{\mathcal{P}}^c)^\top$ in the pre-balanced system (3) satisfies, almost everywhere,*

$$\mathfrak{R}_\eta(\mathbf{v}_{\mathcal{P}}^c) = 0, \quad \mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}}^c) = g(\eta_{\mathcal{P}}^c - b)[1 - \mathcal{P}'(\eta^c - b)] \nabla_{\mathbf{x}} b. \quad (15)$$

In particular:

- (i) for $\mathcal{P}(z) = z_+ := \max(z, 0)$, the residual vanishes in $\mathcal{D}'(\Omega; \mathbb{R}^{d+1})$, including across wet-dry interfaces;
- (ii) for the numerical film $\mathcal{P}_\varepsilon(z) = \max(z, \varepsilon)$ and every $1 \leq p \leq \infty$,

$$\mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}_\varepsilon}^c) = g\varepsilon \mathbf{1}_{\{\eta^c - b < \varepsilon\}} \nabla_{\mathbf{x}} b, \quad \|\mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}_\varepsilon}^c)\|_{L^p(\Omega)} = g\varepsilon \|\nabla_{\mathbf{x}} b\|_{L^p(\{\eta^c - b < \varepsilon\})} \leq g\varepsilon \|\nabla_{\mathbf{x}} b\|_{L^p(\Omega)}; \quad (16)$$

- (iii) for the softplus map $\mathcal{P}_\varepsilon(z) = \varepsilon \log(1 + e^{z/\varepsilon})$,

$$\|\mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}_\varepsilon}^c)\|_{L^p(\Omega)} \leq \frac{g\varepsilon}{e} \|\nabla_{\mathbf{x}} b\|_{L^p(\Omega)}. \quad (17)$$

Consequently, for fixed b , a regularized lake at rest has $\|\mathfrak{R}_{\mathbf{q}}\|_{L^2(\Omega)}^2 = O(\varepsilon^2)$ and therefore an $O(\varepsilon^2)$ VPINN physical-loss contribution on a fixed mesh with fixed normalization. The exact positive part has zero physical and entropy loss, up to quadrature and round-off, when the shock viscosity is disabled. For the film map, the partition-of-unity combination of the VPINN tests on a spatial cell ω_c has the momentum moment

$$\int_{\omega_c} \mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}_\varepsilon}^c) \, d\mathbf{x} = g\varepsilon \int_{\omega_c \cap \{\eta^c - b < \varepsilon\}} \nabla_{\mathbf{x}} b \, d\mathbf{x}. \quad (18)$$

It therefore induces a genuine non-zero VPINN loss floor whenever the final integral does not vanish and the spatial quadrature resolves that moment.

Proof. Write $z = \eta^c - b$. Since $b \in \mathcal{C}^\infty(\bar{\Omega}) \subset W^{1,\infty}(\Omega)$ and $\mathcal{P}$ is Lipschitz, both $\eta_{\mathcal{P}}^c - b = \mathcal{P}(z)$ and $\eta_{\mathcal{P}}^c$ belong to $W^{1,\infty}$. The state is stationary and $\mathbf{q}_{\mathcal{P}}^c = \mathbf{0}$, so the mass residual and kinetic momentum flux vanish. The pressure flux is

$$\frac{g}{2}((\eta_{\mathcal{P}}^c)^2 - 2\eta_{\mathcal{P}}^c b) \mathbb{I}_d = \frac{g}{2}((\eta_{\mathcal{P}}^c - b)^2 - b^2) \mathbb{I}_d.$$

Rademacher's theorem states that every Lipschitz map on a Euclidean space is differentiable almost everywhere. Applied to $\mathcal{P}$, and then combined with the Sobolev chain rule, it gives $\nabla_{\mathbf{x}}(\eta_{\mathcal{P}}^c - b) = -\mathcal{P}'(z) \nabla_{\mathbf{x}} b$ almost everywhere. Since the residual is flux divergence minus source, the bathymetric contribution is $+g\eta_{\mathcal{P}}^c \nabla_{\mathbf{x}} b$. Therefore

$$\begin{aligned} \mathfrak{R}_{\mathbf{q}}(\mathbf{v}_{\mathcal{P}}^c) &= g(\eta_{\mathcal{P}}^c - b) \nabla_{\mathbf{x}}(\eta_{\mathcal{P}}^c - b) - gb \nabla_{\mathbf{x}} b + g\eta_{\mathcal{P}}^c \nabla_{\mathbf{x}} b \\ &= g(\eta_{\mathcal{P}}^c - b)[1 - \mathcal{P}'(\eta^c - b)] \nabla_{\mathbf{x}} b, \end{aligned}$$

which proves (15) almost everywhere. Here $\mathcal{P}'$ denotes any almost-everywhere representative of the derivative; the Sobolev chain rule makes the product independent of its arbitrary values on the non-differentiability set. Each term is the weak derivative of a $W^{1,\infty}$ function and therefore defines a regular distribution. The almost-everywhere identity consequently holds in $\mathcal{D}'(\Omega; \mathbb{R}^{d+1})$; no surface measure has been discarded at a wet-dry interface.

For $\mathcal{P}(z) = z_+$, on $\{z < 0\}$ one has $\eta_{\mathcal{P}}^c - b = 0$, whereas on $\{z > 0\}$ one has $\mathcal{P}' = 1$. The product in (15) hence vanishes almost everywhere, proving (i). For $\mathcal{P}_\varepsilon(z) = \max(z, \varepsilon)$, $(\eta_{\mathcal{P}}^c - b)(1 - \mathcal{P}')$ is zero on $\{z > \varepsilon\}$ and equals ε on $\{z < \varepsilon\}$. Substitution gives the exact film residual and taking its L^p norm proves (16).

For the softplus map, $\mathcal{P}'_\varepsilon(z) = e^{z/\varepsilon}/(1 + e^{z/\varepsilon})$. Setting $y = e^{z/\varepsilon} > 0$ gives

$$\mathcal{P}_\varepsilon(z)[1 - \mathcal{P}'_\varepsilon(z)] = \varepsilon \frac{\log(1 + y)}{1 + y} \leq \frac{\varepsilon}{e},$$

The derivative of $\log(1 + y)/(1 + y)$ has the sign of $1 - \log(1 + y)$; its maximum is attained at $y = e - 1$ and equals $1/e$. This proves (17). On a fixed mesh with fixed normalization, every VPINN moment is a bounded linear functional of the strong residual. Thus an $O(\varepsilon)$ residual gives an $O(\varepsilon^2)$ squared physical loss. Summing the affine or bilinear tests gives the constant function on ω_c ; inserting the film residual proves (18). Finally, for the positive-part equilibrium, $\mathbf{u} = \mathbf{0}$, the energy is stationary and its flux vanishes. Its positive entropy defect is therefore zero. $\square$

Corollary 1 (No smooth positive universal equilibrium map). *There exists no map $\mathcal{P} \in \mathcal{C}^1(\mathbb{R}; (0, \infty))$ such that the inherited state (14) is exactly well-balanced for every level $\eta^c \in \mathbb{R}$ and every non-constant bathymetry $b \in \mathcal{C}^\infty(\Omega)$. Thus, within the present trial construction, global smooth strict positivity and exact universal wet-dry well-balancing cannot be obtained simultaneously from a scalar output regularization.*

Proof. Assume that such a map exists and let $I = (\alpha, \beta)$ be any bounded interval. A Lipschitz domain contains a ball. On a smaller rectangular subset of that ball, choose an affine bathymetry with non-zero constant gradient and choose η^c so that $z(\mathbf{x}) = \eta^c - b(\mathbf{x})$ sweeps I ; extend this bathymetry smoothly to an element of $\mathcal{C}^\infty(\bar{\Omega})$. Exact well-balancing and (15) imply on this subset

$$\mathcal{P}(z) [1 - \mathcal{P}'(z)] \nabla b = \mathbf{0}.$$

The residual in (15) is an L^∞ function, so its vanishing as a distribution is equivalent to its almost-everywhere vanishing. Therefore the displayed identity holds almost everywhere on the affine subdomain. By assumption $\mathcal{P}(z) > 0$ and by construction $\nabla b \neq \mathbf{0}$. Hence $\mathcal{P}'(z) = 1$ for every $z \in I$, using the continuity of $\mathcal{P}'$. Since I was arbitrary, $\mathcal{P}' = 1$ on $\mathbb{R}$. The fundamental theorem of calculus gives $\mathcal{P}(z) = z + K$ for some constant K . Taking $z < -K$ contradicts $\mathcal{P}(z) > 0$. Therefore the assumed smooth strictly positive universal map cannot exist. $\square$

Remark 1 (Design consequence). Theorem 1 characterizes the compatibility between positivity and exact wet-dry well-balancing through the scalar output map. For a numerical film, the resolved moment (18) sets the attainable equilibrium-loss scale and must be included among the discretization parameters.

Scope of the structural results

The solution-dependent viscosity has the smooth- and shock-region scaling established in Proposition 1 and vanishes distributionally under a uniform bound on the spatial variation. The inherited pre-balanced trial map satisfies the explicit equilibrium identity (15), while $\eta_\theta \geq b$, $\mathbf{q}_\theta = (\eta_\theta - b)\mathbf{u}_\theta$ and inheritance of the initial state follow algebraically from the canonical parametrization. An entropy-convergence result additionally requires a uniform BV estimate and consistency of the refined family of residual moments. The entropy defect and localized viscosity provide the corresponding selection and stabilization at the discrete level.

7. Implementation and algebraic solution

We specify the sampling, optimization and floating-point parameters used in the numerical experiments.

Discrete nonlinear problem

After spatial quadrature, the target problem is

$$\min_{\theta \in \Theta} \mathcal{J}_h(\theta), \quad \mathcal{J}_h(\theta) := \mathcal{L}(\theta),$$

with $\mathcal{L}$ given by (13). The symbol θ_h denotes the final iterate returned by the algebraic solver and $\mathbf{v}_h := \mathbf{v}_{\theta_h}$. The parameter-to-state map, the dry-state projection and the solution-dependent viscosity make $\mathcal{J}_h$ non-convex. The algebraic solver seeks a low-residual stationary point of this objective. The strong imbalance and stiffness of PINN objectives are well documented, especially when different residual components evolve on different scales [24, 25]. Mesh size, test space, spatial quadrature, network architecture, sampling rule and stopping budget are therefore all part of the discrete resolution.

Adam and L-BFGS

The two-stage solver addresses two algebraic regimes of the same objective, following the Adam–L-BFGS continuation commonly used in PINN computations [1]. Adam accommodates the stochastic temporal batches and rescales heterogeneous gradient components during the initial descent. Once this stage has reached a low-loss region, L-BFGS uses curvature information from a fixed deterministic closure to refine the same parameter vector. The prescribed iteration budget determines the transition between the two stages.

Adam is an adaptive first-order method that rescales every parameter update using exponential moving averages of the gradient and its square [22]. We use the PyTorch defaults $(\beta_1, \beta_2) = (0.9, 0.999)$, set its algebraic denominator parameter to 10^{-7} and apply no weight decay. Global-gradient clipping is applied at every Adam iteration with the case-dependent norm specified in the numerical settings. The 1D learning rate is multiplied by 0.5 after 55%, 80% and 93% of the prescribed Adam budget. In 2D it is multiplied by 0.35 after 60% and 85%. These parameters control the size of the algebraic updates during the first-order stage.

The final Adam iterate is optionally refined by limited-memory BFGS, a quasi-Newton method that approximates inverse-Hessian information from recent secant pairs [23]. The initial line-search step is 0.5, the strong Wolfe condition is used, and the history contains 50 pairs in 1D and 30 pairs in 2D. The gradient and iterate-change

Table 1: Settings of the one-dimensional experiments. The Adam entry is the prescribed iteration budget and the L-BFGS entry is the maximum number of outer quasi-Newton iterations.

Test case	N_x	N_t	N_q	Adam/L-BFGS max.
Wet-bed dam break	180	80	8	60000/10000
Dry-bed dam break	180	80	8	90000/12500
Fully wet lake at rest	40	20	6	100/0
Partially dry lake	60	30	6	100/0
Thacker oscillation	80	160	8	100000/15000

tolerances are 10^{-9} and 10^{-12} , respectively. One L-BFGS iteration denotes one outer optimizer call; its line search may evaluate the fixed closure up to five times. No gradient clipping is applied during this phase. A fixed temporal sample set supplies the deterministic closure required by the line search and secant update.

Differentiation, quadrature and temporal sampling

Automatic differentiation evaluates the strong temporal and spatial derivatives at every retained pair consisting of a time node and a spatial Gauss point. Gauss–Legendre quadrature remains complete on every spatial cell at each selected time, while the temporal variable is collocated. During Adam, the 1D discretization draws a uniform subset of distinct time indices, whereas the 2D discretization draws one random index from each of a prescribed number of temporal strata. These batches are stochastic estimators of the loss over the complete time-node set. During L-BFGS, the indices are fixed once and are approximately equally spaced over $\{0, \dots, N_t\}$. Both phases use the same residual components, normalization and algebraic weights; Adam resamples the temporal batches and L-BFGS retains one fixed batch.

The entropy term is an arithmetic mean of the pointwise positive production over the retained time-node–Gauss-point pairs. Spatial quadrature weights enter the physical moments. Boundary values are sampled independently of the interior rule. In 1D, 200 equally spaced time values are used at both endpoints and the boundary contribution has weight $\lambda_{bc} = 1$. In 2D, each edge uses 24 spatial points at 8 time values and $\lambda_{bc} = 0.1$. The carrier imposes the initial condition strongly in both dimensions.

Equilibria, dry states and floating-point safeguards

For a lake at rest, the last affine layer of the correction network is initialized to zero, so the initial neural state is the inherited equilibrium. The well-balanced runs perform 100 Adam iterations at learning rate 10^{-7} and preserve the equilibrium within float32 roundoff.

All residuals are evaluated in float32 without mixed precision. The backend is selected in the order CUDA, Apple MPS and CPU, and the random seed is 43 for PyTorch, NumPy and the standard Python generator. Quotients by the derived height use $\max(\eta_\theta - b, H_\varepsilon)$, and every normalization denominator receives an additional 10^{-12} safeguard. The 2D sensor norm uses $\sqrt{|\nabla(\eta_\theta - b)|^2 + 10^{-12}}$, as represented by G_δ in (11). In both dimensions, the physical state is represented by $(\eta_\theta, \mathbf{q}_\theta)$ and the pre-balanced residual is evaluated through the derived depth $\eta_\theta - b$. The equivalent height and free-surface forms coincide almost everywhere on regular wet regions. The case-dependent values of H_ε , C_ν , C_s , λ_{ent} and the algebraic budgets are specified with the numerical experiments.

8. Numerical results

Tables 1 and 2 specify the spatial partitions, temporal sampling, quadrature and algebraic resolution of the experiments.

Here N_x and N_y denote the numbers of Cartesian cells, N_t is the number of temporal subintervals so that the residual is available at $N_t + 1$ time nodes, and N_q is the one-dimensional Gauss–Legendre order used in each cell. The two-dimensional rule therefore contains N_q^2 points. All networks use the hyperbolic tangent activation. The 1D calculations use six hidden layers of width 128, whereas the 2D calculations use five hidden layers of width 80; the final neural layer is followed by the physical map that produces the admissible variables. In the numerical results, θ denotes the final trained parameter vector of each run. Each plotted one-dimensional VPINN profile shows approximately 45–50 representative values evaluated directly at spatial Gauss–Legendre nodes used to assemble the variational residual. These nodes are regularly subsampled from the complete cellwise quadrature grid to keep the figures readable, and the thin blue line is a dense evaluation of the same trained network added only to guide the eye. Exact and finite-volume references are shown as unmarked red curves. The two-dimensional maps use the complete Cartesian evaluation grid. Their extracted profiles are taken along the tensor-product quadrature lines closest to the geometric center, so every displayed marker is an actual spatial quadrature point of the VPINN residual; circles and squares distinguish horizontal and vertical cuts when both are superposed.

The algebraic weights are part of these numerical settings. In 1D we use $\lambda_{bc} = 1$, whereas in 2D we use $\lambda_{bc} = 0.1$. No initial-condition loss is included in either dimension. For the wet- and dry-bed dam breaks, respectively,

Table 2: Settings of the two-dimensional experiments. The spatial quadrature uses N_q^2 tensor-product Gauss points per Cartesian cell; the optimizer entries have the same meaning as in Table 1.

Test case	N_x	N_y	N_t	N_q	Adam/L-BFGS max.
Circular wet-bed dam break	36	36	20	3	25000/500
Circular dry-bed dam break	40	40	24	3	35000/750
Fully wet lake at rest	24	24	8	3	100/0
Partially dry lake at rest	32	32	12	3	100/0
Planar Thacker oscillation	40	40	48	3	20000/1000

$(C_\nu, C_s, \lambda_{\text{ent}}) = (0.04, 8, 0.02)$ and $(0.06, 8, 0.03)$. The entropy defect and shock viscosity are disabled in the lake-at-rest and Thacker experiments. The Adam learning rate is 5×10^{-4} for both dam breaks, 10^{-7} for the lake-at-rest verifications, 10^{-4} for the 1D Thacker test and 2×10^{-4} for its 2D counterpart. Calculations use float32. The division floor is $H_\varepsilon = 10^{-6}$, except for the circular dam-break calculations where $H_\varepsilon = 10^{-5}$; the dry lake-at-rest state itself uses the exact positive part and the floor only regularizes quotients evaluated by the residual.

8.1. Wet-bed dam break

This experiment examines whether the same moment residual can represent a rarefaction and a compressive wave while the entropy and viscosity terms act only near the forming shock. We first consider the standard one-dimensional geometry and then a circular two-dimensional release.

One-dimensional configuration. We consider a flat bottom on $\Omega = (0, 2)$ and initialize the free surface by

$$\eta^0(x) = \eta_R + \frac{\eta_L - \eta_R}{2} \left[1 - \tanh\left(\frac{x-1}{\delta}\right) \right], \quad q^0(x) = 0, \quad (19)$$

with $(\eta_L, \eta_R, \delta) = (2, 1, 0.025)$. The narrow smooth transition provides classical derivatives for automatic differentiation while approximating the discontinuous Riemann datum. It subsequently evolves into the expected left-going rarefaction and right-going shock. To avoid retaining a stationary trace of this regularized jump, both dam-break tests use the decaying memory

$$m(t) = \exp\left(-\kappa \frac{t}{T}\right)$$

and the case-specific carrier

$$\eta_\theta = b + \mathcal{P}_+\left(m(t)(\eta^0 - b) + (1 - m(t)) \left[\eta_{\text{mid}} + \frac{3}{4} H_{\text{ref}} \tanh(\hat{\eta}_\theta) \right]\right), \quad \mathbf{q}_\theta = (1 - m(t)) \mathbf{q}_{\text{ref}} \tanh(\hat{\mathbf{q}}_\theta). \quad (20)$$

Here $\kappa = 12$, $\eta_{\text{mid}} = (\eta_L + \eta_R)/2$, $H_{\text{ref}} = \eta_L$ and $\mathbf{q}_{\text{ref}} = H_{\text{ref}} \sqrt{g H_{\text{ref}}}$. The carrier inherits $(\eta^0, \mathbf{0})$ at $t = 0$ and progressively releases the network from the smoothed profile. The entropy-defect penalty and localized viscosity are also active. Since the smoothed initial profile does not admit the classical discontinuous Riemann solution as an exact reference, the comparison profile is computed independently on 4000 cells with a first-order finite-volume discretization, the Rusanov numerical flux, SSP-RK2 time integration and CFL number 0.15.

Figure 1 superposes the solution at $t = 0, 0.05, 0.10$ and 0.20 . The neural free surface and discharge agree with the finite-volume reference throughout the wave fan at the displayed resolution. The rarefaction profile, the intermediate discharge plateau and the shock position are all recovered, while the stabilization confines visible discrepancies to the immediate shock neighborhood.

Two-dimensional configuration. Let $\Omega = (0, 3)^2$, $\mathbf{x}_c = (1.5, 1.5)$ and $r = |\mathbf{x} - \mathbf{x}_c|$. The circular counterpart of (19) is

$$\eta^0(\mathbf{x}) = \eta_R + \frac{\eta_L - \eta_R}{2} \left[1 - \tanh\left(\frac{r - R_0}{\delta}\right) \right], \quad \mathbf{q}^0(\mathbf{x}) = \mathbf{0}, \quad (R_0, \delta) = (0.5, 0.025),$$

with $(\eta_L, \eta_R) = (2, 1)$ in the fully wet case. The neural carrier is symmetry-adapted: its scalar correction depends on (t, r) and the discharge is reconstructed in the radial direction. More precisely, with $\sigma_{\log}(z) = (1 + e^{-z})^{-1}$, the evolution gate is

$$G(t, r) = (1 - m(t)) \sigma_{\log}\left(\frac{R_0 + 2C_{\text{ref}}t + 2\delta - r}{\delta}\right),$$




Figure 1: Smoothly initialized wet-bed dam break. Free surface (left) and discharge (right) at four selected times. VPINN values at representative spatial Gauss–Legendre nodes are shown as blue points joined by a thin line, and the finite-volume reference is shown in red.

and the radial velocity contains the additional regular factor $\tanh(r/(0.1R_0))$. The correction is initially zero, remains bounded at the center and is localized within the fastest shallow-water signal used in the carrier. The VPINN residual is evaluated on Cartesian cells for both momentum components and all four bilinear moments. The agreement of the two centerline slices checks preservation of the carrier’s radial symmetry by the Cartesian residual.

8.2. Dry-bed dam break

The dry-bed variant couples shock formation to the boundary of the admissible state set. Its 1D and circular 2D versions therefore test positivity and the vanishing of discharge near a moving dry front in addition to the balance-law moments.

One-dimensional configuration. The second experiment retains the same smooth initialization with $(\eta_L, \eta_R, \delta) = (1, 10^{-6}, 0.025)$. The small downstream film provides a numerically admissible approximation of the dry state and regularizes the flux quotient. The bounded discharge carrier is retained, while the entropy and shock-viscosity coefficients are slightly increased to $(C_\nu, C_s, \lambda_{\text{ent}}) = (0.06, 8, 0.03)$.

As shown in Figure 3, the method follows the rarefaction and the advancing wet–dry front over the displayed interval $0 \leq t \leq 0.15$. This problem is more demanding than the wet/wet Riemann problem because the characteristic speeds degenerate with the downstream height. The VPINN reproduces the front displacement and the discharge profile without generating negative depth, showing that the positivity-compatible parameterization and the localized stabilization remain effective in the same calculation. The reference is produced by the same independent finite-volume calculation as in the fully wet test.

Two-dimensional configuration. The circular dry-bed experiment uses the same domain, radius and smooth transition as the preceding wet-bed calculation, now with $(\eta_L, \eta_R) = (1, 10^{-5})$. The downstream film regularizes the dry state while retaining the positivity-compatible discharge map. The entropy-defect and localized-viscosity parameters are those of the one-dimensional dry-bed test.

8.3. Fully wet lake at rest

The fully wet lake isolates the flux–source cancellation without any interference from a shoreline. Both dimensional configurations start from the inherited equilibrium and use the short optimization verification described in Section 7.

One-dimensional configuration. On $\Omega = (0, 1)$, we prescribe the compactly supported bump

$$b_1(x) = \begin{cases} A \sin^2\left(\frac{\pi(x - x_1)}{x_2 - x_1}\right), & x_1 < x < x_2, \\ 0, & \text{otherwise,} \end{cases} \quad (x_1, x_2, A) = (0.125, 0.875, 4.75). \quad (21)$$

The exact equilibrium is $\eta_{\text{eq}} = 5$, with $\eta_{\text{eq}} - b_1 > 0$ and $q_{\text{eq}} = 0$. It is evolved over $0 \leq t \leq 100$. The equilibrium map (6) is initialized with zero correction; its initial normalized loss is 1.74×10^{-15} . We then perform 100 Adam iterations with learning rate 10^{-7} to verify that the optimizer leaves this discrete steady state invariant. Within float32 roundoff the maximum free-surface error is zero, and $\|q_\theta\|_{L^\infty} = 4.03 \times 10^{-11}$ over all selected times. *Two-*




Figure 2: Circular wet-bed dam break on a 36×36 Cartesian mesh. Free-surface maps (left) and horizontal and vertical centerline slices (right) are shown at $t = 0, 0.033, 0.067$ and 0.10 .




Figure 3: Smoothly initialized dam break onto a dry bed. Free surface (left) and discharge (right) at $t = 0, 0.05, 0.10$ and 0.15 , compared with the finite-volume reference.




Figure 4: Circular dam break onto a dry bed on a 40×40 Cartesian mesh. Free-surface maps (left) and the horizontal and vertical centerline slices (right) are shown at $t = 0, 0.067, 0.133$ and 0.20 . The advancing circular wet-dry front remains aligned in both Cartesian directions.




Figure 5: Fully wet lake at rest over the compact bump. The profiles at $t = 0, 25, 50$ and 100 coincide; the bottom is shown because it enters the flux-source cancellation.

dimensional configuration. On $\Omega = (0, 1)^2$, the bottom is the radial bump

$$b_2(\mathbf{x}) = \begin{cases} A \cos^2\left(\frac{\pi r}{2R}\right), & r < R, \\ 0, & r \geq R, \end{cases} \quad r = \left| \mathbf{x} - \left(\frac{1}{2}, \frac{1}{2}\right) \right|, \quad (A, R) = (4.75, 0.375). \quad (22)$$

Both derivatives $\partial_x b_2$ and $\partial_y b_2$ are non-zero away from the symmetry axes. The correction network uses the full input (t, x, y) , and its bounded momentum corrections are $q_{k,\theta} = (t/T)Q_{\text{ref}} \tanh(\widehat{q}_{k,\theta})$, $k \in \{x, y\}$. They vanish together with the zero-initialized final layer. The exact state is $\eta_{\text{eq}} = 5$, with $\eta_{\text{eq}} - b_2 > 0$ and $\mathbf{q}_{\text{eq}} = \mathbf{0}$. The minimum water height is 0.25 at the bump summit. After the prescribed 100 Adam iterations, the free-surface defect vanishes within float32 roundoff, while $\|q_{x,\theta}\|_{L^\infty} = 2.54 \times 10^{-12}$ and $\|q_{y,\theta}\|_{L^\infty} = 2.51 \times 10^{-12}$.

8.4. Partially dry lake at rest

This test adds an emerged region to the same equilibrium mechanism. It is the direct numerical counterpart of Theorem 1, because the positive-part carrier must preserve both the constant wet surface and the identity $\eta = b$ on the dry set.

One-dimensional configuration. We retain the domain and topography (21) and lower the constant water level to $\eta^c = 4.5$. The exact wet-dry equilibrium is

$$\eta_{\text{eq}} = b_1 + (4.5 - b_1)_+ = \max(4.5, b_1), \quad q_{\text{eq}} = 0.$$

The emerged set $\{b_1 > 4.5\}$ is dry and its boundary is a stationary shoreline. Over $0 \leq t \leq 100$, the initial normalized loss is 1.23×10^{-15} . After 100 Adam iterations with learning rate 10^{-7} , the maximum free-surface error vanishes within float32 roundoff and $\|q_\theta\|_{L^\infty} = 3.48 \times 10^{-11}$.

Two-dimensional configuration. For the radial topography (22), we set $\eta^c = 4.5$. Hence $\eta_{\text{eq}} = b_2 + (4.5 - b_2)_+ = \max(4.5, b_2)$, and the central cap $\{b_2 > 4.5\}$ is a dry island. The network is trained on a 32×32 Cartesian grid with the same 100 low-step-size Adam updates. The free-surface defect vanishes within float32 roundoff, while $\|q_{x,\theta}\|_{L^\infty} = 5.16 \times 10^{-12}$ and $\|q_{y,\theta}\|_{L^\infty} = 8.42 \times 10^{-12}$. Figure 8 zooms the centerline profiles around the emerged cap, where the exact free surface coincides with the topography. Across the four lake-at-rest computations, the remaining discharges lie between 10^{-11} and 10^{-12} and the free surfaces agree with equilibrium within float32 roundoff. Their small differences have the magnitude of the optimizer tolerance and float32 roundoff.

8.5. Periodic Thacker oscillation

The periodic experiment replaces stationary balance by a smooth moving wet-dry interface. The one-dimensional solution measures long-time phase preservation, while the planar solution activates both momentum components and transports the wet support in both Cartesian directions.

One-dimensional configuration. In one dimension, we use the Thacker solution in the parabolic basin [38, 39, 41],

$$b(x) = h_c \left(\frac{(x - L/2)^2}{a^2} - 1 \right), \quad (L, a, h_c, g) = (4, 1, 0.5, 9.81).$$

With $\omega = \sqrt{2gh_c}/a$, the exact fields used for training validation are

$$\begin{aligned} \eta(x, t) &= b(x) + h_c \left[1 - \left(\frac{x - L/2}{a} + \frac{\cos(\omega t)}{2a} \right)^2 \right]_+, \\ u(t) &= \frac{\sqrt{2gh_c}}{2a} \sin(\omega t), \quad q(x, t) = (\eta(x, t) - b(x))u(t). \end{aligned} \quad (23)$$

This solution has moving shorelines and period $T_0 = 2\pi/\omega = 2.00606$. In addition to the normalized coordinates x/L and t/T , the network input contains the Fourier features $\sin(\omega t)$ and $\cos(\omega t)$. The initial state is imposed by the case-specific carrier

$$\eta_\theta = b + \mathcal{P}_+(\eta^0 - b + (1 - \cos(\omega t))H_{\text{ref}}\widehat{\eta}_\theta), \quad u_\theta = u^0 + \sin(\omega t)C_{\text{ref}}\widehat{u}_\theta, \quad q_\theta = (\eta_\theta - b)u_\theta.$$

The two harmonic gates vanish at $t = 0$ and encode the known period without prescribing the spatial evolution. Unlike the lakes at rest, this state is genuinely transient and the network is trained up to $T = 5T_0 = 10.0303$, namely over five complete periods. The exact solution supplies a reference at every displayed time.

Figure 9 reports the free surface and discharge at $t = 0, 2.507575, 5.01515, 7.522725$ and 10.0303 . The computation recovers the alternating motion of the shoreline and the return toward the initial state after five periods.




Figure 6: Fully wet two-dimensional lake at rest over the radial bump. Free-surface maps with bathymetric contours (left), free-surface centerlines (upper right) and x -discharge centerlines (lower right) are shown at $t = 0, 25, 50$ and 100 . The one-dimensional panels resolve the scale of the numerical defects.




Figure 7: Partially dry lake at rest over the compact bump. The inherited equilibrium preserves the stationary shoreline and zero discharge up to $T = 100$.

This experiment combines smooth dynamics, long-time phase propagation, positivity and wetting–drying without invoking the shock-specific entropy viscosity.

The preservation of the oscillation period can be assessed more directly by tracking both moving wet–dry contacts. Their exact trajectories are

$$x_L(t) = \frac{L}{2} - a - \frac{1}{2} \cos(\omega t), \quad x_R(t) = \frac{L}{2} + a - \frac{1}{2} \cos(\omega t).$$

The numerical contacts are extracted from the connected wet component containing the maximum water height, using the threshold $\eta_\theta - b = 10H_\varepsilon$. As shown in Figure 10, the five successive oscillations remain in phase with both theoretical trajectories. For the left contact, the maximum and root-mean-square position errors are 5.43×10^{-3} and 2.33×10^{-3} ; for the right contact, they are 8.80×10^{-3} and 2.72×10^{-3} .

Two-dimensional configuration. The two-dimensional configuration uses the planar parabolic-bowl solution employed by Duran, Marche, Turpault and Berthon [40], which belongs to the moving-shoreline solutions of Thacker and Sampson–Easton–Singh [38,39]. We consider the square $\Omega = (0, 10000)^2$ and introduce the centered coordinates $X = x - 5000$, $Y = y - 5000$. The topography and parameters are

$$b(X, Y) = h_0 \frac{X^2 + Y^2}{a^2}, \quad (h_0, a, B, g) = (10 \text{ m}, 3000 \text{ m}, 5 \text{ m s}^{-1}, 9.81 \text{ m s}^{-2}).$$

With $s = \sqrt{2gh_0}/a$, the exact motion is

$$\begin{aligned} \tilde{\eta}(X, Y, t) &= h_0 - \frac{B^2}{2g} - \frac{Bs}{g} (\cos(st)X + \sin(st)Y), \\ u(t) &= B \sin(st), \quad v(t) = -B \cos(st), \\ \eta(X, Y, t) &= b(X, Y) + (\tilde{\eta}(X, Y, t) - b(X, Y))_+, \\ \mathbf{q}(X, Y, t) &= (\eta(X, Y, t) - b(X, Y))(u(t), v(t))^T. \end{aligned}$$

The period is $T_0 = 2\pi/s = 1345.710 \text{ s}$ and the half-period is 672.855 s , close to the 672 s value quoted in [40]. The affine extension $\tilde{\eta}$ defines a rotating normal, translates the wet disk in both coordinate directions and activates both momentum equations.

For this benchmark, define the normalized coordinates $\tau = t/T_0$, $\xi = X/5000$ and $v = Y/5000$. The network receives $(\tau, \xi, v, \sin(st), \cos(st))$ and predicts the free surface and velocity through the harmonic carrier

$$\begin{aligned} \tilde{\eta}_\theta &= \eta^0 + A_\eta [\sin(st) \tanh(\hat{\eta}_{\theta,s}) + (1 - \cos(st)) \tanh(\hat{\eta}_{\theta,c})], \\ u_\theta &= A_u \sin(st) \tanh(\hat{u}_\theta), \\ v_\theta &= -B + A_u (1 - \cos(st)) \tanh(\hat{v}_\theta), \\ \eta_\theta &= b + (\tilde{\eta}_\theta - b)_+, \\ \mathbf{q}_\theta &= (\eta_\theta - b) \mathbf{u}_\theta, \end{aligned}$$

where $\eta^0 = b + (\tilde{\eta}(\cdot, 0) - b)_+$, $A_\eta = 2H_{\text{ref}}$ and $A_u = 3B/2$. All five inputs enter every raw output, and the network class contains general non-radial corrections. Both harmonic factors vanish at $t = 0$ and the initial state




Figure 8: Partially dry two-dimensional lake at rest over the radial bump. Free-surface maps with bathymetric contours (left) cover the full domain, whereas the free-surface, topography and discharge centerlines (right) are restricted to a neighborhood of the dry island.

is inherited exactly. Their separation is necessary because the y -slope of η and u leave their initial values at order t , whereas the x -slope correction and $v + B$ are of order t^2 . The two harmonic gates represent these distinct first- and second-order temporal variations. The carrier encodes the known period, while the spatial coefficients remain trainable and the loss contains no exact field values.

Before training, automatic differentiation of the analytic fields on random interior wet points gives normalized strong-residual RMS 8.5×10^{-8} and maximum 4.5×10^{-7} in float32. This diagnostic is independent of the network and verifies that the exact solution, topography and sign conventions solve the same operator as the VPINN loss.

Figure 11 compares the trained state with the exact motion at the four quarter-periods and after one complete period. The wet disk moves successively in both Cartesian directions, while the horizontal and vertical cuts recover the two rotating slopes. On 101 equally spaced evaluation times, discrete norms over the union of the numerical




Figure 9: Periodic Thacker oscillation over five complete periods. Neural free surface (left) and discharge (right) are compared with the exact solution at five selected times.




Figure 10: Trajectories of the left contact x_L (left) and right contact x_R (right) for the Thacker test over five periods. The numerical shorelines extracted from the VPINN solution are compared with their exact periodic trajectories; markers correspond to the selected evaluation times.

and exact wet supports give

$$\frac{\|(\eta_\theta - b) - (\eta - b)\|_2}{\|\eta - b\|_2} = 5.30 \times 10^{-2}, \quad \frac{\|\mathbf{q}_\theta - \mathbf{q}\|_2}{\|\mathbf{q}\|_2} = 7.10 \times 10^{-2},$$

with height RMSE 3.05×10^{-1} m and maximum error 1.32 m. The largest pointwise defect lies in the shallow moving front, whereas the interior slopes and the phase of the translated disk are accurately reproduced.

9. Conclusion

We have formulated a shallow-water VPINN as a finite-dimensional nonlinear Petrov–Galerkin least-squares discretization: a global space–time neural trial manifold is paired with local affine or Cartesian bilinear tests, while the residual is written in pre-balanced free-surface variables. This viewpoint separates algebraic properties of the trial map from estimates based on quadrature-defined residual moments.

The constrained VPINN parameterization gives $\eta_\theta \geq b$; its canonical form sets $\mathbf{q}_\theta = (\eta_\theta - b)\mathbf{u}_\theta$, and it contains the initial data and lakes at rest exactly. Under compatible quadrature, the latter are zero-residual global minimizers of the VPINN loss. The scheme-specific identity (15) makes this compatibility quantitative: the exact positive part preserves a wet–dry lake in distributions, while a film of height ε creates an $O(\varepsilon)$ momentum residual and an $O(\varepsilon^2)$ squared consistency scale. Whenever the constant moment (18) is non-zero, this scale is an actual VPINN loss floor. Corollary 1 further proves that no smooth strictly positive scalar map can preserve every wet–dry lake within this trial construction. Positivity regularization and well-balancing therefore cannot be analyzed independently.

Proposition 1 gives the second method-specific result: the nonlinear viscosity is $O(h^2)$ where $\eta_h - b$ has a bounded gradient, may reach the shock-capturing $O(h)$ scale in an $O(h)$ transition layer, and vanishes in distributions along




Figure 11: Fully two-dimensional planar Thacker oscillation over one period. Free-surface maps on the wet support (left) include the exact shoreline (red dashed curve); the horizontal cuts (middle) and vertical cuts (right) compare the VPINN free surface with the exact solution at five selected times. Each cut is restricted to the wet support and a narrow dry neighborhood of its two shorelines; the parabolic topography is shown in gray.

refined sequences with uniformly bounded spatial variation. Convergence to an entropy solution further requires a uniform BV estimate and consistency of the refined residual moments. The positive entropy defect and localized viscosity remain explicit components of the numerical method.

The numerical results are consistent with these conclusions in one space dimension and on the completed Cartesian two-dimensional tests. In particular, the genuinely two-dimensional lake-at-rest calculation activates both bathymetric derivatives while preserving the equilibrium within float32 roundoff. The planar Thacker calculation activates both momentum components and follows the translating wet support over a complete rotation, with the discrepancies concentrated near the moving shallow front. Taken together, the experiments show that one VPINN construction can combine non-negative water depth, exact lakes at rest, entropy-regularized shock capture and moving wet-dry interfaces in one- and two-dimensional shallow-water configurations. This simultaneous treatment is the principal distinction from earlier neural shallow-water solvers that addressed subsets of these requirements. Proving a network-independent BV estimate, removing the symmetry-adapted carrier from the circular dam breaks and extending the analysis to dry fronts remain open problems.

References

- [1] 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.
- [2] I. E. Lagaris, A. Likas and D. I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5):987–1000, 1998.

- [3] S. Berrone, C. Canuto, M. Pintore and N. Sukumar. Enforcing Dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks. *Heliyon*, 9:e18820, 2023.
- [4] W. Zhang and M. Al Kobaisi. On the monotonicity and positivity of physics-informed neural networks for highly anisotropic diffusion equations. *Energies*, 15(18):6823, 2022.
- [5] A. Bihlo and R. O. Popovych. Physics-informed neural networks for the shallow-water equations on the sphere. *Journal of Computational Physics*, 456:111024, 2022.
- [6] S. Dazzi. Physics-informed neural networks for the augmented system of shallow water equations with topography. *Water Resources Research*, 60:e2023WR036589, 2024.
- [7] X. Qi, G. A. M. de Almeida and S. Maldonado. Physics-informed neural networks for solving flow problems modeled by the 2D shallow water equations without labeled data. *Journal of Hydrology*, 636:131263, 2024.
- [8] R. Brecht, E. Cardoso-Bihlo and A. Bihlo. Physics-informed neural networks for tsunami inundation modeling. *Journal of Computational Physics*, 536:114066, 2025.
- [9] Y. Tian, S. Ding, L. Huang, G. Su and J. Chen. Physics-informed neural networks for solving the two-dimensional shallow water equations with terrain topography and rainfall source terms. *Water Resources Research*, 61:e2025WR040052, 2025.
- [10] A. D. Jagtap, E. Kharazmi and G. E. 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.
- [11] E. Kharazmi, Z. Zhang and G. E. Karniadakis. Variational physics-informed neural networks for solving partial differential equations. arXiv:1912.00873, 2019.
- [12] R. Khodayi-Mehr and M. M. Zavlanos. VarNet: Variational neural networks for the solution of partial differential equations. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 298–307, 2020.
- [13] J. Bruna, B. Peherstorfer and E. Vanden-Eijnden. Neural Galerkin schemes with active learning for high-dimensional evolution equations. *Journal of Computational Physics*, 496:112588, 2024.
- [14] E. Franck, V. Michel-Dansac and L. Navoret. Approximately well-balanced discontinuous Galerkin methods using bases enriched with physics-informed neural networks. *Journal of Computational Physics*, 512:113144, 2024.
- [15] H. Barucq, M. Duprez, F. Faucher, E. Franck, F. Lecourtier, V. Lleras, V. Michel-Dansac and N. Victorion. Enriching continuous Lagrange finite element approximation spaces using neural networks. arXiv:2502.04947, 2025.
- [16] V. Dolean, D. Hrebenshchykova, S. Lanteri and V. Michel-Dansac. Neural network-driven domain decomposition for efficient solutions to the Helmholtz equation. arXiv:2511.15445, 2026.
- [17] M. Campos Pinto, V. Fournet, E. Franck, P. Krah and V. Michel-Dansac. A neural characteristic mapping method: Lagrangian PINNs based on flow maps for transport-dominated problems. arXiv:2610.12395, 2026.
- [18] S. Berrone, C. Canuto and M. Pintore. Variational physics informed neural networks: The role of quadratures and test functions. *Journal of Scientific Computing*, 92:100, 2022.
- [19] E. Kharazmi, Z. Zhang and G. E. Karniadakis. hp-VPINNs: Variational physics-informed neural networks with domain decomposition. *Computer Methods in Applied Mechanics and Engineering*, 374:113547, 2021.
- [20] T. De Ryck, S. Mishra and R. 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.
- [21] S. Mishra and R. Molinaro. Estimates on the generalization error of physics-informed neural networks for approximating PDEs. *IMA Journal of Numerical Analysis*, 43(1):1–43, 2023.
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015.
- [23] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989.
- [24] S. Wang, Y. Teng and P. Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.
- [25] A. S. Krishnapriyan, A. Gholami, S. Zhe, R. M. Kirby and M. W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems 34*, pages 26548–26560, 2021.
- [26] E. Audusse, F. Bouchut, M.-O. Bristeau, R. Klein and B. Perthame. A fast and stable well-balanced scheme with hydrostatic reconstruction for shallow water flows. *SIAM Journal on Scientific Computing*, 25(6):2050–2065, 2004.
- [27] Y. Xing, X. Zhang and C.-W. Shu. Positivity-preserving high-order well-balanced discontinuous Galerkin methods for the shallow water equations. *Advances in Water Resources*, 33(12):1476–1493, 2010.
- [28] Y. Xing and X. Zhang. Positivity-preserving well-balanced discontinuous Galerkin methods for the shallow water equations on unstructured triangular meshes. *Journal of Scientific Computing*, 57(1):19–41, 2013.
- [29] C. Berthon and F. Marche. A positive preserving high order VFRoe scheme for shallow water equations: a class of relaxation schemes. *SIAM Journal on Scientific Computing*, 30(5):2587–2612, 2008.
- [30] Q. Liang and F. Marche. Numerical resolution of well-balanced shallow water equations with complex source terms. *Advances in Water Resources*, 32(6):873–884, 2009.
- [31] A. Haidar, F. Marche and F. Vilar. A posteriori finite-volume local subcell correction of high-order discontinuous Galerkin schemes for the nonlinear shallow-water equations. *Journal of Computational Physics*, 452:110902, 2022.
- [32] S. Cardonna, A. Haidar, F. Marche and F. Vilar. Local subcell monolithic DG/FV methods for nonlinear shallow-water models with source terms. HAL open archive, hal-05434014, version 1, 2025.
- [33] S. Cardonna, F. Marche and F. Vilar. A high-order robust subcell monolithic DG/FV formulation for nonlinear shallow-water equations on unstructured grids. HAL open archive, hal-05630914, version 2, 2026.
- [34] A. Duran, Q. Liang and F. Marche. On the well-balanced numerical discretization of shallow water equations on unstructured meshes. *Journal of Computational Physics*, 235:565–586, 2013.
- [35] C. Berthon and C. Chalons. A fully well-balanced, positive and entropy-satisfying Godunov-type method for the shallow-water equations. *Mathematics of Computation*, 85(299):1281–1307, 2016.
- [36] C. Berthon, S. Bulteau, F. Foucher, M. M’Baye and V. Michel-Dansac. A very easy high-order well-balanced reconstruction for hyperbolic systems with source terms. *SIAM Journal on Scientific Computing*, 44(4):A2506–A2535, 2022.
- [37] C. Berthon and V. Michel-Dansac. A fully well-balanced hydrodynamic reconstruction. *Journal of Numerical Mathematics*, 32(3):275–299, 2024.
- [38] W. C. Thacker. Some exact solutions to the nonlinear shallow-water wave equations. *Journal of Fluid Mechanics*, 107:499–508, 1981.
- [39] J. Sampson, A. Easton and M. Singh. Moving boundary shallow water flow above parabolic bottom topography. *ANZIAM*

- Journal*, 47:C373–C387, 2006.
- [40] A. Duran, F. Marche, R. Turpault and C. Berthon. Asymptotic-preserving scheme for the shallow water equations with source terms on unstructured meshes. *Journal of Computational Physics*, 287:184–206, 2015.
- [41] O. Delestre, C. Lucas, P.-A. Ksinant, F. Darboux, C. Laguerre, T. N. T. Vo, F. James and S. Cordier. SWASHES: a compilation of shallow water analytic solutions for hydraulic and environmental studies. *International Journal for Numerical Methods in Fluids*, 72(3):269–300, 2013.