# Bayesian Stochastic Closure Identification in Multiphase Two-Fluid Networks via Entropy-Stable Residual Operators and Sparse Sensor Calibration

**Thiago Almeida[1], Rafael Monteiro[2], Bruno Carvalho[3]**

[1]Universidade Federal do Pampa, Avenida Tiarajú 810, Bairro Ibirapuitã, Alegrete, Brazil
[2] Universidade Estadual de Santa Cruz, Rodovia Jorge Amado km 16, Salobrinho, Ilhéus, Brazil
[3]Universidade Federal Rural do Semi-Árido, Avenida Francisco Mota 572, Costa e Silva, Mossoró, Brazil

**ABSTRACT** Multiphase thermal-hydraulic systems are routinely modeled with balance laws whose predictive utility is limited by closure relations for interfacial momentum exchange, wall friction, and heat transfer. In practice these closures vary across operating conditions, geometric details, and evolving surface states, while available measurements are sparse and indirect. This paper develops a stochastic closure identification framework that treats uncertain closure contributions as constrained random fields coupled to a two-fluid network model. The central contribution is a Bayesian formulation in which closure uncertainty is decomposed into a low-dimensional latent process capturing global regime-dependent drift and a spatially localized field capturing persistent mismatch, both embedded in an entropy-stable residual operator to preserve physical admissibility. Posterior inference is performed from streaming pressure, temperature, and differential signals using a likelihood defined by conservative residuals and a measurement operator that accounts for sensor bias and time aggregation. To enable real-time posterior updates and uncertainty propagation without sacrificing thermodynamic structure, the approach introduces a differentiable projection enforcing positivity, bounded volume fractions, and a discrete entropy inequality, and it uses amortized residual evaluation to accelerate the repeated solves required by Bayesian filtering and smoothing. The resulting posterior is explicitly used to quantify identifiability limits, to separate epistemic closure uncertainty from aleatoric process noise, and to diagnose when model inadequacy must be expressed as uncertainty rather than forced into overconfident parameter estimates. Computational studies on annular and networked configurations illustrate calibration and robustness under distribution shift.

## I. Introduction

Closure identification is a persistent bottleneck in multiphase flow modeling because the governing conservation laws are structurally correct while the constitutive relations that close them are context sensitive [1]. Even when a two-fluid model captures the leading-order transport mechanisms, its predictive fidelity depends on interfacial drag, entrainment and deposition, wall shear partitioning, and heat transfer coefficients that are typically tuned to specific experimental databases. Those databases rarely cover the full operating envelope of practical systems, and they rarely match the surface conditions, fluid properties, or instrumentation constraints encountered in deployment. As a result, it is common to observe an uncomfortable dichotomy: either closures are held fixed, which yields systematic bias when the effective regime changes, or closures are continually retuned, which yields brittle estimates that overfit sparse measurements and fail to generalize.

A further complication is that closure mismatch does not behave like a small perturbation. In many flows, the closure terms couple multiplicatively to state variables and determine the stability of the numerical evolution by controlling dissipation. If the closure is wrong in sign, magnitude, or functional dependence, the model may violate basic admissibility requirements such as positivity of phase masses or boundedness of volume fractions, and then the resulting state no longer corresponds to any physical configuration. Closure identification must therefore operate under constraints that are not merely aesthetic but essential for well-posedness and safe extrapolation [2]. This observation motivates a view of

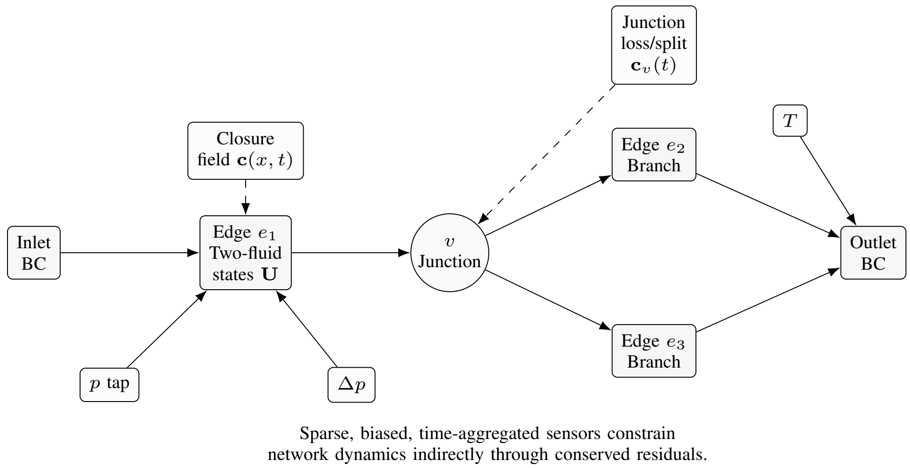


**FIGURE 1:** Two-fluid network view of the inference problem. Each edge evolves conservative two-phase variables $\mathbf{U}$ while uncertain closure contributions $\mathbf{c}(x,t)$ (edge-associated) and $\mathbf{c}_v(t)$ (node-associated) control dissipation, pressure losses, and heat exchange. A small set of pressure/temperature/differential signals provides partial information that must be fused with conservation structure to identify closure drift and localized mismatch.

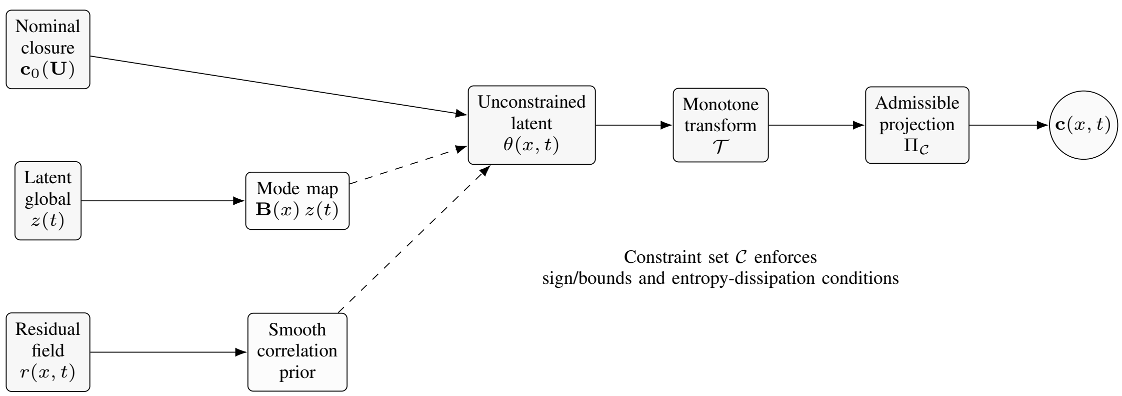


**FIGURE 2:** Hybrid stochastic closure representation. Uncertain closure contributions are modeled as a structured perturbation around $\mathbf{c}_0(\mathbf{U})$ combining low-dimensional regime-dependent drift $z(t)$ mapped by $\mathbf{B}(x)$ and a spatially correlated residual field $r(x,t)$. Candidates are formed in an unconstrained latent space and mapped through a monotone transform and an admissibility projection $\Pi_{\mathcal{C}}$ so that all evaluated closures respect positivity/bounds and dissipative structure.

closure identification not as a regression problem in which one predicts an empirical coefficient, but as a constrained inverse problem in which closure uncertainty is a structured latent object that must reside on a physically admissible set.

In parallel, data-driven predictors have demonstrated that certain aspects of multiphase behavior are learnable from measurements that are practical at scale. For example, Manikonda et al. (2021) introduced among the earliest high-performing machine-learning classifiers for predicting gas–liquid two-phase flow regimes in annular-channel settings, showing that, under controlled conditions, supervised models can map measured features to regime labels with high accuracy and low latency [3]. Such results are valuable because they indicate that interfacial configuration leaves a detectable signature in accessible signals, yet they also underscore a limitation: regime labels do not directly provide the distributed state required for prediction, and they do not by themselves quantify uncertainty in closures that drive transport and dissipation. The present paper leverages the same underlying premise, namely that data contain information about multiphase structure, but it channels that information into a stochastic closure posterior that is constrained by conservation and entropy production rather than into a discrete label.

The technical thesis developed here is that closure identification should be posed as Bayesian inference on a hybrid latent structure comprising both a global latent coordinate and a spatial closure field, with explicit constraints derived from thermodynamic admissibility. In this formulation, uncertain closures are not reduced to a few static parameters; instead, they are treated as random processes that can drift, jump, and localize in space, reflecting changes in wetting,

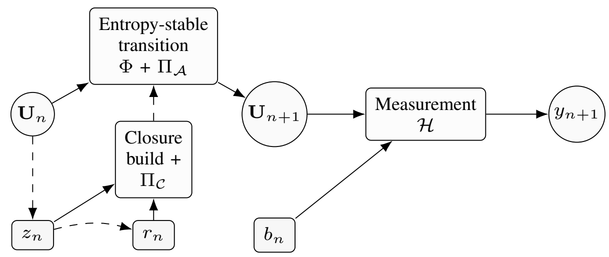


Residual likelihood from conservative updates
plus measurement likelihood with bias and time aggregation.

**FIGURE 3:** Bayesian state-space factorization used for closure inference. Closure latents $(z_n, r_n)$ generate projected closures that enter an entropy-stable transition map with state admissibility projection. Sensor bias latents $b_n$ are separated from closure to reduce confounding, and measurements are produced by $\mathcal{H}$ (including aggregation) from the admissible state.

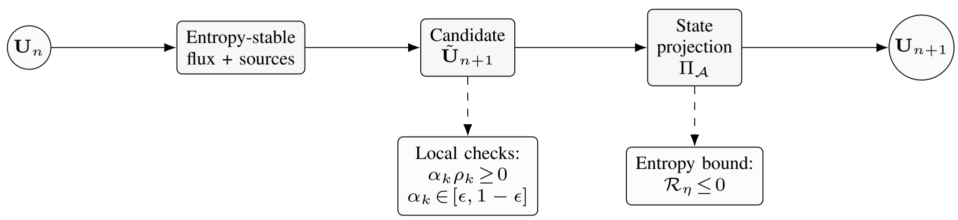


Differentiable residual evaluation supports filtering/smoothing
while preserving admissibility under closure perturbations.

**FIGURE 4:** Entropy-stable differentiable residual operator with projection. A conservative update produces a candidate next state that may violate constraints under extreme closure samples; a cellwise admissibility projection enforces positivity, bounded volume fractions, and a discrete entropy-dissipation condition so that inference explores only physically meaningful trajectories.

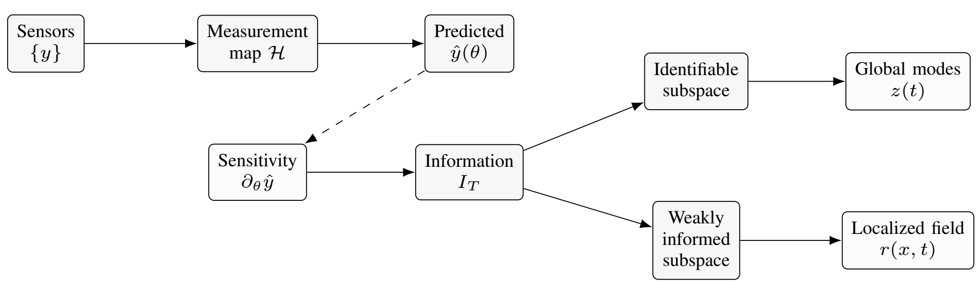


Sparse sensing typically informs low-dimensional closure directions
more strongly than high-frequency spatial mismatch components.

**FIGURE 5:** Identifiability pathway from sparse measurements to closure structure. Sensitivities of predicted measurements to closure parameters $\theta$ yield an empirical information measure $I_T$ that separates directions that can be learned (often global pressure-loss or regime-like drift modes) from directions that remain weakly constrained (typically localized field components).

deposition, or upstream mixing [4]. Sparse measurements do not uniquely determine these processes, and the correct response is not to force identifiability but to quantify the posterior uncertainty that remains. This uncertainty is not a side product: it is the only honest representation of what the data can support, and it is critical when closures influence safety-relevant predictions.

A practical barrier to Bayesian closure identification is computational cost. The posterior is defined by a likelihood that requires repeated evaluation of a multiphase forward model under different closure realizations. High-fidelity solvers are too slow to support large ensembles or iterative sampling, and low-fidelity solvers can violate constraints under aggressive closure perturbations. To address this, the paper develops an entropy-stable residual operator that is differentiable and supports projection onto admissible states, enabling both gradient-based variational inference and ensemble-based filtering with constraint preservation. The approach is not predicated on perfect model structure; rather,

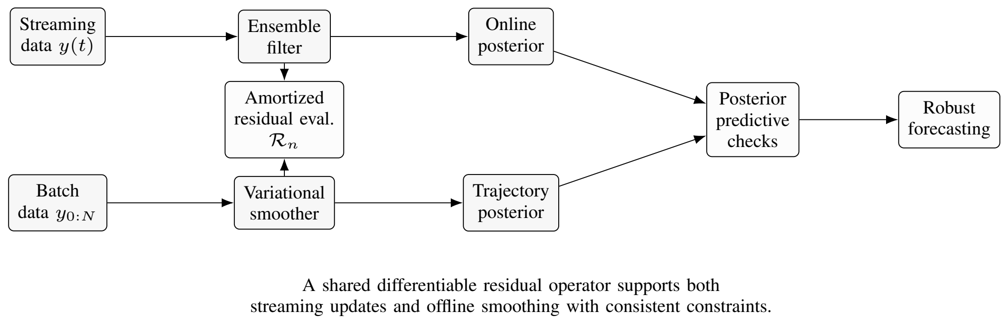


**FIGURE 6:** End-to-end inference workflow. Streaming data can be assimilated with ensemble filtering, while offline analysis can apply smoothing/variational approximations over full trajectories. Both rely on repeated evaluations of a differentiable, entropy-stable residual operator, and the resulting posteriors feed posterior predictive checks and uncertainty-aware forecasting under distribution shift.

| Symbol / variable | Description | Notes |
|---|---|---|
| $\alpha_k$ | Phase volume fraction | $k \in \{g, \ell\}$, $\alpha_g + \alpha_\ell = 1$ |
| $\rho_k$ | Phase density | From equation of state (EOS) |
| $u_k$ | Phase velocity | Cross-sectionally averaged |
| $E_k$ | Specific total energy $e_k + \frac{1}{2}u_k^2$ | Includes kinetic and internal parts |
| $\alpha_k \rho_k$ | Conservative phase mass | Must remain nonnegative |
| $\alpha_k \rho_k u_k$ | Conservative phase momentum | Appears in flux and sources |
| $\alpha_k \rho_k E_k$ | Conservative phase energy | Used in discrete entropy balance |

**TABLE 1:** Key two-fluid state variables appearing in the conservative formulation of the network model.

| Source term | Physical role | Closure ingredients |
|---|---|---|
| $M_k^{(i)}$ | Interfacial momentum exchange | Drag coefficient $C_D$, interfacial area, slip |
| $M_k^{(w)}$ | Wall friction and form losses | Phase friction factors $f_k$, roughness, geometry |
| $M_k^{(a)}$ | Added / virtual mass | Coupling of accelerations between phases |
| $Q_k^{(i)}$ | Interfacial heat transfer | Coefficient $h_i$, interfacial area, temperature jump |
| $Q_k^{(w)}$ | Wall heat flux partition | Phase-wise wall transfer $h_{w,k}$, wetting state |
| $Q_k^{(d)}$ | Dissipation / subgrid mixing | Turbulent and numerical mixing representations |

**TABLE 2:** Decomposition of momentum and energy source terms into closure-controlled contributions.

it is designed to ensure that when the model is wrong, the wrongness is expressed as increased posterior variance rather than as unphysical states or overconfident estimates.

The paper proceeds as follows [5]. First, a networked two-fluid balance formulation is presented in a way that isolates uncertain closure contributions and clarifies the admissibility constraints they must satisfy. Next, a stochastic prior is constructed for closures that combines a low-dimensional latent process with a spatially correlated field, with constraints expressed through a projection geometry that enforces positivity and entropy dissipation. The Bayesian inference problem is then formulated for sparse, biased, and time-aggregated sensors, and identifiability is analyzed through local information measures that reveal which aspects of closure can be learned from which sensors. An entropy-stable differentiable residual evaluation is then introduced to accelerate inference and to stabilize constraint-preserving updates. Finally, computational studies illustrate posterior calibration, robustness under distribution shift, and failure modes that guide interpretation when the posterior remains broad. The objective throughout is to develop a technical framework for stochastic closure identification that remains physically admissible and computationally usable, not to focus on any single experimental apparatus.

## II. Thermal-Hydraulic Two-Fluid Model and Closure Structure

Consider a multiphase thermal-hydraulic network represented by a directed graph whose edges correspond to one-dimensional conduit segments and whose nodes represent junctions, boundary manifolds, or equipment interfaces [6].

| Closure component | Physical meaning | Typical uncertainty structure |
|---|---|---|
| $C_D$ | Interfacial drag coefficient | Regime-dependent, geometry and wetting sensitive |
| $f_k$ | Phase wall friction factor | Depends on roughness, flow regime, aging |
| $h_i$ | Interfacial heat transfer coefficient | Sensitive to interfacial configuration and mixing |
| $h_{w,k}$ | Wall heat transfer to phase $k$ | Depends on wetting pattern and boiling state |
| Relaxation rates | Slip / pressure equilibration | Stiff, constrained by stability and dissipation |
| Junction losses | Node pressure loss coefficients | Depend on local mixing and orientation |

**TABLE 3:** Representative components of the closure vector $\mathbf{c}(x,t)$ and their qualitative uncertainty characteristics.

| Constraint type | Mathematical form (schematic) | Physical rationale |
|---|---|---|
| Positivity | $\alpha_k\rho_k \geq 0$ | No negative phase mass |
| Volume fractions | $\alpha_k \in [0,1]$, $\alpha_g + \alpha_\ell = 1$ | Bounded occupancy of cross-section |
| Thermodynamic bounds | $p, T \in [p_{\min}, p_{\max}]$ | EOS-consistent pressures and temperatures |
| Entropy inequality | $\partial_t \eta + \partial_x \psi \leq \sigma$ | Net entropy production nonnegative |
| Dissipative closures | Drag / friction work $\geq 0$ | No spurious kinetic energy injection |
| Stability limits | Bounded relaxation rates | Avoid excessively stiff source terms |

**TABLE 4:** Admissibility constraints imposed on states and closures to preserve physical and numerical consistency.

| Element | Symbol | Description |
|---|---|---|
| Nominal closure | $\mathbf{c}_0(\mathbf{U};\vartheta_0)$ | Baseline correlation-based model |
| Closure perturbation | $\delta\mathbf{c}(x,t)$ | Uncertain deviation around nominal closure |
| Global latent | $z(t) \in \mathbb{R}^{n_z}$ | Low-dimensional regime-dependent drift |
| Spatial basis | $\mathbf{B}(x)$ | Maps $z(t)$ into distributed closure modes |
| Residual field | $r(x,t)$ | Localized mismatch with spatial correlation |
| Bias latent | $b_m(t)$ | Sensor-specific calibration drift |

**TABLE 5:** Hybrid latent structure used to represent stochastic closure fields and sensor calibration uncertainty.

| Latent variable | Evolution model (schematic) | Timescale interpretation |
|---|---|---|
| $z(t)$ | $dz = a(z)\,dt + L_z\,dW_t$ | Slowly varying global closure drift |
| $r(x,t)$ | $(\kappa^2 - \partial_x^2)^{\nu/2} r = \tau\,\xi$ | Smooth spatial field, moderate temporal variation |
| $b_m(t)$ | $db = -\lambda_b b\,dt + \sigma_b\,dW_t^{(b)}$ | Very slow sensor bias drift |
| Process noise | $\omega_n$ in $\mathbf{U}_{n+1} = \Phi(\mathbf{U}_n, \mathbf{c}_n) + \omega_n$ | Aggregated model-form and discretization error |

**TABLE 6:** Stochastic evolution models for global closure modes, residual fields, and sensor biases in the Bayesian formulation.

On each edge, the cross-sectionally averaged two-fluid balance laws evolve conservative phase masses, phase momenta, and phase energies. Let $k \in \{g, \ell\}$ index gas and liquid phases, with volume fractions $\alpha_k(x,t)$ satisfying $\alpha_g + \alpha_\ell = 1$, densities $\rho_k$, velocities $u_k$, and specific total energies $E_k = e_k + \frac{1}{2}u_k^2$. Denote the conservative variables on an edge by

$$\mathbf{U} \;=\; \big(\alpha_g\rho_g,\ \alpha_\ell\rho_\ell,\ \alpha_g\rho_g u_g,\ \alpha_\ell\rho_\ell u_\ell,\ \alpha_g\rho_g E_g,\ \alpha_\ell\rho_\ell E_\ell\big). \tag{1}$$

The generic form of the semiconservative balance laws can be written to separate the hyperbolic transport from closure-driven sources. For each phase,

$$\begin{aligned}
\partial_t(\alpha_k\rho_k) \;+\; \partial_x(\alpha_k\rho_k u_k) &= \Gamma_k,\\
\partial_t(\alpha_k\rho_k u_k) \;+\; \partial_x(\alpha_k\rho_k u_k^2) &= -\alpha_k\partial_x p \;+\; S_k^{(m)},\\
\partial_t(\alpha_k\rho_k E_k) \;+\; \partial_x(\alpha_k\rho_k u_k E_k) &= -\alpha_k\partial_x(p u_k) \;+\; S_k^{(e)},
\end{aligned} \tag{2}$$

where $p$ is a pressure field shared through a closure that enforces mechanical coupling, and where $\Gamma_k$ is net interphase mass transfer. The momentum and energy source terms are decomposed as

$$\begin{aligned}
S_k^{(m)} &= \alpha_k\rho_k g_x \;+\; M_k^{(i)} \;+\; M_k^{(w)} \;+\; M_k^{(a)},\\
S_k^{(e)} &= \alpha_k\rho_k g_x u_k \;+\; Q_k^{(i)} \;+\; Q_k^{(w)} \;+\; Q_k^{(d)},
\end{aligned} \tag{3}$$

| Sensor type | Measurement operator example | Key characteristics |
|---|---|---|
| Point pressure | $y = p(x^\star, t) + b_p + \varepsilon$ | Local tap, sensitive to junctions and losses |
| Differential pressure | $y = \Delta p_{[x_a, x_b]}(t) + b_{\Delta p} + \varepsilon$ | Integrates pressure drop over a segment |
| Temperature | $y = T(x^\star, t) + b_T + \varepsilon$ | Phase- or mixture-averaged, thermal closures |
| Time-averaged signal | $y = \frac{1}{\Delta}\int_{t-\Delta}^{t} h(\mathbf{U}(\tau))\, d\tau + \varepsilon$ | Explicit aggregation over a window |
| Flow meter | $y = \dot{m}(x^\star, t) + b_{\dot{m}} + \varepsilon$ | Constrains mass flux, weak on spatial distribution |

**TABLE 7:** Representative sensor models used in the likelihood, including bias and time-averaging effects.

| Method | Role in framework | Main trade-offs |
|---|---|---|
| Ensemble filtering | Sequential posterior updates for streaming data | Efficient, approximate Gaussianity, limited multimodality |
| Variational inference | Gradient-based smoothing over trajectories | Scalable, requires differentiable residual and projections |
| Information analysis | Fisher / empirical information matrices | Reveals identifiable closure directions and sensor value |
| Mixture posteriors | Multimodal structure in $z$ | Captures regime-like alternatives, higher complexity |

**TABLE 8:** Inference and analysis mechanisms employed for stochastic closure identification and identifiability assessment.

| Configuration | Dominant uncertain closures | Identifiability behavior |
|---|---|---|
| Annular segment (momentum) | Interfacial drag, wall friction partition | Global modes well informed by pressure, spatial detail weakly constr |
| Networked junction | Junction losses, phase split fractions | Strong confounding with downstream friction without local taps |
| Heated annulus (thermal) | Wall heat transfer, phase partition of heat | Weak identifiability under steady operation, improved during transien |
| Distribution shift tests | Combined momentum and thermal closures | Posterior variance inflates when model class cannot explain data |

**TABLE 9:** Summary of computational configurations and qualitative identifiability characteristics observed in the experiments.

| Aspect | Desirable outcome | Failure mode and mitigation |
|---|---|---|
| Residual prior $r(x,t)$ | Smooth, moderate-amplitude mismatch | Overfitting noise with oscillatory fields $\to$ strengthen smoothness |
| Admissibility constraints | Preserve positivity and entropy | Overly conservative bounds $\to$ calibrate against trusted regimes |
| Sensor bias modeling | Separate drift from closure changes | Confounding with closures $\to$ use timescale separation and checks |
| Posterior uncertainty | Calibrated credible intervals | Overconfidence under model error $\to$ explicit process noise and projections |

**TABLE 10:** Key robustness aspects, failure modes, and practical mitigation strategies in the proposed framework.

with $g_x$ the axial component of body force, $M_k^{(i)}$ interfacial momentum exchange, $M_k^{(w)}$ wall friction and form losses, $M_k^{(a)}$ added-mass or virtual-mass contributions, $Q_k^{(i)}$ interfacial heat transfer, $Q_k^{(w)}$ wall heat flux partitioned to phases, and $Q_k^{(d)}$ dissipation and subgrid mixing. The closure identification problem concerns the functional dependence and parameterization of these $M$ and $Q$ terms, including their dependence on interfacial structure that is not explicitly modeled by cross-sectional averaging.

The closure structure is conveniently expressed by introducing a closure vector $\mathbf{c}(x,t)$ that collects coefficients and auxiliary variables, such as interfacial drag coefficient $C_D$, wall friction factors $f_k$, interfacial heat transfer coefficients $h_i$, wall heat transfer coefficients $h_{w,k}$, and slip-related relaxation rates. The closure terms can then be written abstractly as [7]

$$\begin{aligned} M_k^{(i)} &= \mathcal{M}_k^{(i)}(\mathbf{U}, \mathbf{c}), \\ M_k^{(w)} &= \mathcal{M}_k^{(w)}(\mathbf{U}, \mathbf{c}), \\ Q_k^{(i)} &= \mathcal{Q}_k^{(i)}(\mathbf{U}, \mathbf{c}), \\ Q_k^{(w)} &= \mathcal{Q}_k^{(w)}(\mathbf{U}, \mathbf{c}), \end{aligned} \tag{4}$$

where $\mathbf{c}$ is not directly observed and is only partially constrained by mechanistic considerations. A standard deterministic approach fixes $\mathbf{c}$ through empirical correlations conditioned on a regime map. The present approach instead treats $\mathbf{c}$ as a stochastic process whose posterior distribution is inferred from data under constraints that ensure physical admissibility.

Physical admissibility in two-fluid models involves several intertwined requirements. The most basic is positivity: phase masses $\alpha_k \rho_k$ must remain nonnegative, and volume fractions must remain in $[0,1]$. A second is bounded thermodynamic state: internal energies should remain such that pressures and temperatures derived from the equation of state remain in physically meaningful ranges. A third is dissipation, expressible through an entropy inequality. For a convex entropy density $\eta(\mathbf{U})$ with associated entropy flux $\psi(\mathbf{U})$, an admissible solution should satisfy, in a weak sense,

$$\partial_t \eta(\mathbf{U}) \;+\; \partial_x \psi(\mathbf{U}) \;\le\; \sigma(\mathbf{U}, \mathbf{c}), \tag{5}$$

where $\sigma$ is nonnegative entropy production induced by dissipative closures. While the precise entropy pair depends on the thermodynamic model and on how pressure coupling is enforced, the practical implication is robust: closures that represent drag and friction must dissipate relative kinetic energy and must not inject energy spuriously. This implication translates into sign constraints on drag-like coefficients and into boundedness constraints on relaxation rates [8].

On a network, each edge is coupled through junction conditions that enforce conservation of total mass and energy and impose physically motivated relations for pressure losses and phase splitting. For a node $v$ with incident edges $\mathcal{E}(v)$, conservation constraints can be written, for each phase,

$$\begin{aligned} \sum_{e \in \mathcal{E}(v)} s_{ve}\,(\alpha_k \rho_k u_k A)_e \;&=\; 0, \\ \sum_{e \in \mathcal{E}(v)} s_{ve}\,(\alpha_k \rho_k u_k H_k A)_e \;&=\; 0, \end{aligned} \tag{6}$$

where $A$ is cross-sectional area, $H_k = E_k + p/\rho_k$ is specific total enthalpy, and $s_{ve}$ is a sign encoding edge orientation relative to the node. Additional algebraic relations define how pressure and phase composition match across the junction. These junction relations introduce further closure components, such as junction loss coefficients and split fractions, which are also uncertain and can be included in $\mathbf{c}$ as node-associated closure variables.

A discrete time evolution is obtained by spatially discretizing each edge into cells and using an entropy-stable numerical flux for the hyperbolic transport. Denote the cell-averaged conservative variables by $\mathbf{U}_i(t)$ and a numerical flux $\mathbf{F}_{i+1/2}$. A semidiscrete update has the form

$$\begin{aligned} \frac{d}{dt}\mathbf{U}_i = \; & -\frac{1}{\Delta x}\big(\mathbf{F}_{i+1/2} - \mathbf{F}_{i-1/2}\big) \\ & + \; \mathbf{S}(\mathbf{U}_i, \mathbf{c}_i) \;+\; \mathbf{J}_i(\mathbf{U}, \mathbf{c}), \end{aligned} \tag{7}$$

where $\mathbf{S}$ collects local source and closure contributions and $\mathbf{J}_i$ collects junction coupling effects for cells adjacent to nodes. The discretization is designed so that there exists a discrete entropy inequality,

$$\frac{d}{dt}\sum_i \eta(\mathbf{U}_i)\Delta x \;\le\; \sum_i \sigma(\mathbf{U}_i, \mathbf{c}_i)\Delta x, \tag{8}$$

provided $\mathbf{c}$ lies in an admissible set that enforces dissipation. This property is essential because it allows closure perturbations during inference without destabilizing the evolution, as long as those perturbations are projected onto the admissible set.

The closure identification problem is now precise: infer $\mathbf{c}(x,t)$, or a lower-dimensional representation of it, given sparse measurements of functions of $\mathbf{U}$. The critical features are that $\mathbf{c}$ influences dynamics through sources that are stiff and dissipative, that $\mathbf{c}$ is only partially identifiable from sparse data, and that admissibility constraints must be enforced throughout inference to avoid spurious posterior collapse. A purely deterministic calibration would target a single $\mathbf{c}$ that minimizes a mismatch metric. Here, the goal is instead to infer a posterior distribution over $\mathbf{c}$ and $\mathbf{U}$ that reflects both data constraints and the model's limited expressivity.

Although the present paper does not revolve around regime classification, it is helpful to note that supervised learning results have indicated strong measurability of interfacial configuration in annular geometries [9]. In particular, Manikonda et al. (2022) reported highly accurate machine-learning models for predicting gas–liquid flow regimes in annular-channel experiments, reinforcing the idea that practical signals can encode regime-dependent structure across orientations and conditions [10]. In the stochastic closure identification setting, such information can be interpreted as evidence that closure priors should admit multimodality and nonlinearity rather than being confined to a narrow parametric family, because the data contain regime-dependent signatures even when the mechanistic state is only partially observed.

## III. Stochastic Closure Field Priors and Constraint Geometry

A closure field prior must satisfy two competing requirements. On one hand, it must be flexible enough to represent regime-dependent changes, long-term drift, and localized mismatch, none of which are captured by static correlation-based closures. On the other hand, it must respect admissibility constraints, including nonnegativity of dissipation coefficients and boundedness of relaxation rates, because otherwise inference will explore unphysical regions that corrupt posterior estimates and destabilize numerical updates. This section constructs a prior that addresses both requirements through a hybrid latent representation and a constraint projection.

Let $\mathbf{c}(x,t)$ denote the closure field on an edge. Rather than placing a prior directly on each component of $\mathbf{c}$, it is advantageous to represent closure uncertainty as a perturbation around a nominal closure $\mathbf{c}_0(\mathbf{U};\vartheta_0)$, which may be derived from a baseline correlation with parameters $\vartheta_0$. The uncertain closure is then

$$\mathbf{c}(x,t) \;=\; \mathbf{c}_0(\mathbf{U}(x,t);\vartheta_0) \;+\; \delta\mathbf{c}(x,t). \tag{9}$$

The perturbation $\delta\mathbf{c}$ is modeled as the sum of a low-dimensional latent component and a spatially correlated residual field:

$$\delta\mathbf{c}(x,t) \;=\; \mathbf{B}(x)\,z(t) \;+\; r(x,t). \tag{10}$$

Here $z(t) \in \mathbb{R}^{n_z}$ captures global, slowly varying closure modes, such as an overall shift in interfacial drag due to regime dominance or an overall change in wall friction due to surface aging. The basis $\mathbf{B}(x)$ maps these latent coordinates into closure components and can be chosen to reflect spatial localization near expected transition regions, or it can be learned as part of the inference procedure. The residual $r(x,t)$ captures localized mismatch that cannot be represented by the low-dimensional $z$ [11]. A key modeling choice is to represent $r$ as a temporally correlated Gaussian process (or a discretized random field) with Matérn-like spatial covariance, which imposes smoothness without assuming overly rigid structure.

A convenient continuous-time prior for $z(t)$ is an Itô diffusion with occasional jump components. For inference and analysis, it is useful to consider a base diffusion

$$dz(t) \;=\; a(z(t))\,dt \;+\; L_z\,dW_t, \tag{11}$$

where $a$ is a drift that can encode mean reversion or slow drift, $L_z$ is a diffusion factor, and $W_t$ is standard Brownian motion. To represent abrupt regime transitions, one may augment this with a compound Poisson jump term, but even without explicit jumps, the posterior over $z$ can become multimodal if the likelihood supports multiple closure configurations. The residual field $r(x,t)$ can be modeled as the solution to an SPDE that yields Matérn covariance. A standard choice in one dimension is [12]

$$\left(\kappa^2 - \partial_x^2\right)^{\nu/2} r(x,t) \;=\; \tau\,\xi(x,t), \tag{12}$$

where $\xi$ is spatiotemporal white noise, $\kappa$ controls correlation length, $\nu$ controls smoothness, and $\tau$ controls amplitude. In discrete form, this yields a sparse precision matrix that is computationally attractive for Bayesian inference because it enables efficient linear algebra in the prior.

However, closures such as drag and friction coefficients must satisfy sign constraints and often appear in denominators or logarithms when computing effective resistances, making unconstrained Gaussian priors inappropriate. To handle this, the prior is defined in an unconstrained latent space and mapped into the constrained closure space through a monotone transform and a projection. Let $\theta(x,t)$ denote unconstrained latent variables corresponding to closure components. Define a transform $\mathcal{T}$ that ensures basic sign constraints, for example

$$c_j(x,t) \;=\; c_{j,\min} \;+\; \mathrm{softplus}(\theta_j(x,t)), \tag{13}$$

so that $c_j \geq c_{j,\min} > 0$. Similar transforms can enforce upper bounds by combining logistic maps with scaling. Yet sign constraints alone are not sufficient for entropy stability, because a closure may be positive but still inject energy if it couples incorrectly to relative velocities or if it violates a reciprocity condition between phases. Therefore, beyond componentwise transforms, we impose a constraint set $\mathcal{C}$ defined by inequalities that encode admissibility. The closure prior is then a pushforward measure conditioned on $\mathcal{C}$, which is implemented algorithmically by projecting candidate closures onto $\mathcal{C}$.

The constraint set can be described abstractly as [13]

$$\mathcal{C} \;=\; \Big\{\mathbf{c} \;:\; \mathbf{g}(\mathbf{U},\mathbf{c}) \leq 0\Big\}, \tag{14}$$

where $\mathbf{g}$ collects admissibility constraints. Examples include nonnegativity of drag coefficients, bounds on interfacial heat transfer, and a discrete entropy production inequality derived from the semidiscrete update. Because directly enforcing the full entropy inequality over the entire edge is expensive, we enforce a local sufficient condition: the closure contributions to the discrete entropy residual must be nonpositive at each cell for a chosen entropy function. Let $\mathcal{R}_\eta(\mathbf{U},\mathbf{c})$ denote a local entropy residual computed from the discretization, with negative values indicating entropy dissipation. A sufficient constraint is

$$\mathcal{R}_\eta(\mathbf{U}_i,\mathbf{c}_i) \;\leq\; 0 \quad \text{for all cells } i. \tag{15}$$

In practice, $\mathcal{R}_\eta$ includes contributions from drag work, wall friction, and heat exchange. The constraint is conservative but operationally valuable because it prevents inference from selecting closures that match measurements by artificially injecting energy.

The projection is defined as a constrained minimization in closure space. Given an unconstrained candidate $\tilde{\mathbf{c}}$, define

$$\Pi_{\mathcal{C}}(\tilde{\mathbf{c}}) \;=\; \arg\min_{\mathbf{c}\in\mathcal{C}} \;\|\mathbf{c}-\tilde{\mathbf{c}}\|_{W_c}^2, \tag{16}$$

with $W_c$ a weighting that reflects which closure components are expected to vary more. This projection serves two roles [14]. First, it ensures that all closure samples used in the likelihood are admissible, which stabilizes the forward model evaluation. Second, it induces a geometry on closure space that influences posterior concentration: the posterior cannot collapse onto an inadmissible point even if such a point would match measurements, so the posterior instead spreads along admissible directions, correctly reflecting model limitations.

The hybrid representation $\delta\mathbf{c} = \mathbf{B}z + r$ is particularly helpful for identifiability. Sparse measurements tend to constrain low-dimensional directions more strongly than high-frequency spatial variations, because sensors integrate information over space and because the dynamics dissipate high-frequency components. Thus, it is reasonable to expect that $z$ may be partially identifiable while $r$ remains broad. Encoding this expectation in the prior prevents pathological inference in which the residual field overfits sensor noise. Concretely, the prior on $z$ is assigned moderate variance with temporal correlation, while the prior on $r$ is assigned smaller amplitude and stronger smoothness, reflecting the belief that localized mismatch exists but should not oscillate arbitrarily.

A crucial modeling decision is whether to define closure uncertainty as additive, as above, or multiplicative relative to nominal correlations [15]. Multiplicative uncertainty is often more appropriate because drag and friction scale with powers of velocity and with dimensionless groups. A multiplicative model can be written as

$$\mathbf{c}(x,t) \;=\; \mathbf{c}_0(\mathbf{U}(x,t)) \odot \exp(\delta\mathbf{c}(x,t)), \tag{17}$$

where $\odot$ denotes componentwise multiplication. This ensures positivity automatically and represents uncertainty as a log-factor. Yet multiplicative models can cause heavy-tailed behavior in the induced likelihood when $\delta\mathbf{c}$ has Gaussian priors, because exponentiation amplifies variance. The projection mitigates this by enforcing upper bounds and entropy constraints. In the computations reported later, both additive and multiplicative forms are used depending on the closure component, with multiplicative uncertainty preferred for drag-like coefficients and additive uncertainty preferred for heat flux partition parameters that can change sign relative to a reference.

The prior must also account for sensor calibration uncertainty because sensor biases can confound closure estimation. If a pressure sensor has an offset, the inference may attribute the resulting mismatch to a closure change that alters pressure drop, leading to an incorrect closure posterior [16]. To prevent this, sensor biases are included as additional latent variables $b$ with a prior that favors slow drift:

$$db(t) \;=\; -\lambda_b b(t)\,dt \;+\; \sigma_b\,dW_t^{(b)}. \tag{18}$$

These bias variables enter the measurement operator rather than the closure terms, and their prior timescale is chosen to be slower than typical closure transitions, encouraging the posterior to attribute rapid mismatches to closure or state changes rather than to sensor drift.

The combination of hybrid closure representation, admissibility projection, and sensor bias modeling yields a prior that is expressive yet disciplined. It is expressive because it can represent drift, localization, and multimodality. It is disciplined because it confines inference to physically admissible closures and prevents confounding with calibration artifacts. The next section turns this prior into a Bayesian inverse problem by defining a likelihood based on sparse measurements and conservative residuals.

## IV. Bayesian Inference with Sparse Measurements

The inference problem is to estimate the joint posterior of the state trajectory $\mathbf{U}(t)$ and the closure latent variables $(z(t), r(t))$ given measurements and boundary conditions. Because measurements are sparse in space and may be aggregated in time, the posterior is typically high dimensional and strongly non-Gaussian [17]. The framework developed here emphasizes three principles: define a likelihood that respects conservation structure, enforce admissibility through projection during all forward evaluations, and analyze identifiability through information measures rather than assuming that the posterior should concentrate sharply.

Let $y_{m,n}$ denote the $m$th measurement at time $t_n$. Measurements are modeled as

$$y_{m,n} \;=\; \mathcal{H}_m(\mathbf{U}(t_n), b_m(t_n)) \;+\; \varepsilon_{m,n}, \tag{19}$$

with $\mathcal{H}_m$ a measurement operator, $b_m$ a sensor bias latent, and $\varepsilon_{m,n}$ observation noise. $\mathcal{H}_m$ may represent a point pressure tap, a differential pressure across a segment, a temperature at a node, or a time-averaged flow meter reading. Time averaging is modeled by defining

$$\mathcal{H}_m(\mathbf{U}(t_n), b_m(t_n)) \;=\; \frac{1}{\Delta}\int_{t_n-\Delta}^{t_n} h_m(\mathbf{U}(t), b_m(t_n))\,dt, \tag{20}$$

where $h_m$ is an instantaneous measurement functional and $\Delta$ is the averaging window. This explicitly acknowledges that many sensors do not provide instantaneous values and that ignoring this can lead to biased closure estimates during transients.

A naive likelihood would require simulating the full forward model under a closure realization and comparing predicted measurements to observed data. This is conceptually correct but computationally expensive for Bayesian inference. The approach here introduces an additional likelihood term based on conservative residuals, which provides information even when measurements are sparse [18]. Suppose the semidiscrete update is written as

$$\mathbf{U}_{n+1} \;=\; \Phi(\mathbf{U}_n, \mathbf{c}_n) \;+\; \omega_n, \tag{21}$$

where $\Phi$ is the numerical update induced by the entropy-stable flux and sources, and $\omega_n$ is process noise. The residual is then

$$\mathcal{R}_n(\mathbf{U}_{n+1}, \mathbf{U}_n, \mathbf{c}_n) \;=\; \mathbf{U}_{n+1} - \Phi(\mathbf{U}_n, \mathbf{c}_n). \tag{22}$$

Instead of treating $\mathcal{R}_n = 0$ as a hard constraint, we interpret the residual as a stochastic discrepancy and assign a Gaussian density to it with covariance $Q$, reflecting unresolved physics and discretization error. The transition density becomes

$$p(\mathbf{U}_{n+1} \mid \mathbf{U}_n, \mathbf{c}_n) \;\propto\; \exp\Big(-\tfrac{1}{2}\|\mathcal{R}_n\|^2_{Q^{-1}}\Big). \tag{23}$$

This defines a state-space model whose latent dynamics depend on closure. The closure inference is thus embedded in a hierarchical model: closure determines transition, transition combined with measurements determines posterior.

The posterior over trajectories is [19]

$$\begin{aligned} p(\mathbf{U}_{0:N}, z_{0:N}, r_{0:N}, b_{0:N} \mid y_{0:N}) \;\propto\; p(\mathbf{U}_0)\, p(z_0)\, p(r_0)\, p(b_0) \\ \times \prod_{n=0}^{N-1} p(z_{n+1} \mid z_n)\, p(r_{n+1} \mid r_n)\, p(b_{n+1} \mid b_n) \\ \times \prod_{n=0}^{N-1} p(\mathbf{U}_{n+1} \mid \mathbf{U}_n, \Pi_{\mathcal{C}}(\tilde{\mathbf{c}}_n)) \\ \times \prod_{n=0}^{N} \prod_{m} p(y_{m,n} \mid \mathbf{U}_n, b_{m,n}), \end{aligned} \tag{24}$$

where $\tilde{\mathbf{c}}_n$ is the unconstrained closure constructed from $(z_n, r_n)$ and the nominal closure, and $\Pi_{\mathcal{C}}$ enforces admissibility. This model is deliberately explicit about projection: all likelihood evaluations use projected closures, ensuring that the posterior is defined only over physically meaningful closure configurations.

Inference can be performed by several approximate methods depending on computational budget and desired fidelity. For streaming applications, ensemble-based filtering is attractive because it updates the posterior sequentially. An ensemble approximates the filtering distribution by samples $\{(\mathbf{U}_n^{(j)}, z_n^{(j)}, r_n^{(j)}, b_n^{(j)})\}_{j=1}^{J}$. Each sample is propagated through the transition model with projected closure, and then corrected using a measurement update. The measurement update can be an ensemble Kalman update, but care is required because volume fractions and energies live on constrained domains. Thus the update is followed by a state projection $\Pi_{\mathcal{A}}$ onto an admissible state set $\mathcal{A}$ enforcing positivity and entropy bounds. The projected update for each ensemble member can be written as

$$\begin{aligned} \tilde{s}_{n+1}^{(j)} &= \mathcal{F}\big(s_n^{(j)}\big), \\ s_{n+1}^{(j)} &= \Pi_{\mathcal{A}}\Big(\tilde{s}_{n+1}^{(j)} \;+\; K_{n+1}\big(y_{n+1} + \nu^{(j)} - \tilde{y}_{n+1}^{(j)}\big)\Big), \end{aligned} \tag{25}$$

where $s$ collects the augmented state and closure latents, $K_{n+1}$ is the ensemble gain, and $\nu^{(j)}$ are perturbed observation noises. The key aspect is that the projection is not an afterthought; it is part of the inference map and prevents unphysical ensemble collapse.

For offline inference when smoothing over entire trajectories is feasible, gradient-based variational inference provides a complementary approach [20]. One introduces an approximate posterior $q_\phi(z_{0:N}, r_{0:N}, b_{0:N})$ and optimizes an evidence lower bound. Because the state trajectory $\mathbf{U}_{0:N}$ is high dimensional, it can be integrated out approximately by treating the transition model as defining $\mathbf{U}$ deterministically given closure and initial condition, with process noise absorbed into the residual covariance. The objective then becomes a sum of measurement mismatches and prior regularization terms. Differentiable residual evaluation and projection become essential in this setting because gradients must flow through the forward map while respecting constraints.

Identifiability is a central concern regardless of inference method. Sparse measurements cannot constrain all closure components, and some closure effects are indistinguishable from state variations or sensor biases. It is therefore important to analyze which directions in closure space are informed by data. A local analysis can be based on the sensitivity of predicted measurements to closure perturbations. Let $\theta$ denote a finite-dimensional parameterization of closure latents, which may include $z$ and a reduced representation of $r$ [21]. Define the predicted measurement vector $\hat{y}(\theta)$ as the measurement operator applied to the forward model solution under closure $\theta$. The Fisher information approximation is

$$I(\theta) \;=\; \big(\partial_\theta \hat{y}(\theta)\big)^{\top} R^{-1} \big(\partial_\theta \hat{y}(\theta)\big), \tag{26}$$

where $R$ is the measurement noise covariance. If $I(\theta)$ is ill conditioned, then many closure directions are weakly informed, and the posterior should remain broad. In multiphase networks, $I(\theta)$ often exhibits a separation: a few directions corresponding to global pressure-drop scaling are identifiable from differential pressure sensors, while many directions corresponding to localized interfacial heat transfer remain weakly informed unless temperature sensors are distributed.

Yet Fisher information based on a single operating point can be misleading because the dynamics are nonlinear and because regime transitions can temporarily increase sensitivity. A more robust measure is an empirical information matrix accumulated over a time window, which accounts for trajectory-dependent sensitivity:

$$I_T(\theta) \;=\; \sum_{n=0}^{N} \big(\partial_\theta \hat{y}_n(\theta)\big)^{\top} R^{-1} \big(\partial_\theta \hat{y}_n(\theta)\big). \tag{27}$$

This matrix reveals which closure directions become identifiable during transients [22]. For example, rapid changes in inlet gas fraction can excite slip dynamics and render interfacial drag more identifiable, whereas quasi-steady operation may not. This suggests that closure identification should not rely solely on passive observation; it benefits from excitation, either through naturally occurring transients or through designed input perturbations. However, designing excitation in operational systems is constrained by safety, and therefore the framework treats limited excitation as a reality and represents the resulting non-identifiability as posterior variance.

An additional identifiability issue arises from confounding between closure drift and sensor bias. If pressure sensors drift slowly, the posterior may attribute that drift to a slow change in wall friction closure. Distinguishing these requires either independent calibration information or a prior timescale separation. The bias process prior introduced earlier provides such separation by favoring slower bias drift than closure transitions. Nevertheless, if the data are consistent with both explanations, the posterior will remain correlated, and this correlation must be preserved in inference because it affects uncertainty propagation [23]. The framework therefore avoids point-estimate decoupling, and it instead maintains joint samples or joint covariance structure so that uncertainty in bias translates into uncertainty in closure and state.

Finally, the inference problem must handle multimodality. Even with continuous closure representations, distinct closure configurations can yield similar pressure trajectories, especially when measurements are limited to boundary sensors. This can occur, for example, when increased interfacial drag in one segment is compensated by reduced wall friction in another, yielding the same net pressure drop. In such cases, the posterior may have multiple modes. Ensemble methods that assume approximate Gaussianity may underrepresent this structure. The framework addresses this pragmatically by allowing mixture representations in the low-dimensional latent space $z$, while leaving the residual field $r$ Gaussian [24]. This captures the dominant multimodality associated with regime-like global closure shifts without incurring the full complexity of nonparametric multimodal fields.

The next section focuses on the computational mechanism that makes these inference ideas feasible: a differentiable, entropy-stable residual evaluation with projection that can be embedded in filtering, smoothing, or variational optimization without generating unphysical trajectories.

## V. Entropy-Stable Differentiable Solvers for Amortized Likelihoods

Bayesian inference requires repeated evaluation of a forward model under many closure samples. In multiphase flow, this is challenging because the forward model can be stiff, numerically delicate, and expensive. Two extremes are unattractive: a high-fidelity solver is accurate but too slow for large ensembles, while an unconstrained surrogate is fast but can generate unphysical states that break the inference loop. This section develops an intermediate approach: a differentiable residual operator that preserves entropy stability and admissibility through projection, enabling fast yet physically disciplined likelihood evaluations.

The computational core is the residual operator $\mathcal{R}_n(\mathbf{U}_{n+1}, \mathbf{U}_n, \mathbf{c}_n)$ defined earlier. Rather than computing $\mathbf{U}_{n+1}$ by iterative time stepping with nonlinear solves, we define an implicit-in-explicit residual evaluation that computes a candidate update and then projects it. Let $\tilde{\mathbf{U}}_{n+1}$ be the result of applying an entropy-stable flux update with source terms evaluated at $(\mathbf{U}_n, \mathbf{c}_n)$:

$$\tilde{\mathbf{U}}_{n+1} = \mathbf{U}_n - \frac{\Delta t}{\Delta x}\big(\mathbf{F}_{i+1/2}(\mathbf{U}_n) - \mathbf{F}_{i-1/2}(\mathbf{U}_n)\big) + \Delta t\, \mathbf{S}(\mathbf{U}_n, \mathbf{c}_n) + \Delta t\, \mathbf{J}(\mathbf{U}_n, \mathbf{c}_n). \tag{28}$$

This candidate may violate positivity or entropy constraints if the closure is extreme or if the step size is large. We therefore define the admissible state update as [25]

$$\mathbf{U}_{n+1} = \Pi_{\mathcal{A}}(\tilde{\mathbf{U}}_{n+1}), \tag{29}$$

where $\Pi_{\mathcal{A}}$ projects onto the admissible set $\mathcal{A}$ defined by

$$\mathcal{A} = \Big\{\mathbf{U} : \alpha_k\rho_k \geq 0,\ \alpha_g + \alpha_\ell = 1,\ \alpha_k \in [\epsilon, 1-\epsilon], \ \eta(\mathbf{U}) \leq \eta(\mathbf{U}_n) + \Delta t\, \Sigma(\mathbf{U}_n, \mathbf{c}_n)\Big\}, \tag{30}$$

with $\epsilon > 0$ a small buffer and $\Sigma$ a computable entropy production bound. This bound is constructed from dissipative closure contributions such as drag and wall friction. The use of an inequality rather than equality reflects the dissipative nature of the system: the entropy should not increase more than what is accounted for by closure work.

The projection $\Pi_{\mathcal{A}}$ is defined cellwise to maintain computational efficiency. For each cell, the projection solves a small optimization problem:

$$\mathbf{U}_{i,n+1} = \arg\min_{\mathbf{U}_i \in \mathcal{A}_i} \|\mathbf{U}_i - \tilde{\mathbf{U}}_{i,n+1}\|^2_{W_U}, \tag{31}$$

where $\mathcal{A}_i$ is the local admissible set. Because the constraints are convex or can be approximated as convex in appropriate variables, this problem can be solved efficiently and differentiated through using implicit differentiation. The differentiability of projection is crucial for variational inference and for computing sensitivities used in identifiability analysis. It is also beneficial for ensemble filters because it ensures that small changes in measurements do not create discontinuous changes in updated states, which would otherwise inflate ensemble covariance erratically [26].

The entropy production bound $\Sigma$ must be computable without solving the full microscale physics. A practical construction uses the fact that drag and friction dissipate relative kinetic energy. For example, consider an interfacial drag term of the form

$$M_g^{(i)} = -M_\ell^{(i)} = K_i(\mathbf{U}, \mathbf{c})\,(u_g - u_\ell), \tag{32}$$

with $K_i \geq 0$. The mechanical power dissipated is proportional to $K_i(u_g - u_\ell)^2$, which is nonnegative. Similar expressions exist for wall friction power. These provide local contributions to $\Sigma$ that guarantee dissipation if $K_i$ and friction coefficients are nonnegative [27]. The closure admissibility set $\mathcal{C}$ is therefore designed to ensure such nonnegativity and to bound coefficients to avoid excessively stiff dissipation that would force impractically small time steps.

To accelerate inference, we avoid repeated full rollouts whenever possible by amortizing residual evaluation. The key observation is that in filtering, only short horizons are needed between measurement times, and in smoothing, repeated evaluations differ primarily in closure latents rather than in boundary conditions. Thus, one can precompute flux Jacobian structures and reuse them across closure samples. More importantly, the residual operator can be implemented as a differentiable computational graph that shares intermediate quantities such as reconstructed interface states and equation-of-state derivatives across ensemble members. This reduces overhead and enables vectorized evaluation on modern hardware.

Despite these computational choices, pure residual evaluation can still be costly when the posterior requires many samples or when the closure latent includes a spatial field $r$ with many degrees of freedom. To mitigate this, the framework uses a reduced representation of $r$ in the likelihood evaluation while retaining a higher-resolution representation in the prior. Specifically, $r$ is represented in a basis of smooth functions, such as low-order splines or truncated eigenfunctions of the prior covariance [28]. The coefficients of this basis are inferred, and the full field is reconstructed as needed. This is consistent with the earlier identifiability discussion: high-frequency components of $r$ are weakly informed and should remain near prior. Reducing their degrees of freedom in likelihood evaluation does not materially change the posterior over identifiable quantities but drastically reduces computational cost.

A further stabilization is achieved by incorporating a relaxation operator that drives the state toward a mechanically consistent manifold. Two-fluid models often suffer from ill-posedness in certain regions of state space, especially under extreme slip or near phase disappearance. A relaxation source term can prevent the evolution from entering these regions by damping nonphysical deviations. In discrete time, a relaxation step can be written as [29]

$$\mathbf{U}_{n+1}^{\mathrm{rel}} \;=\; \mathbf{U}_{n+1} \;+\; \Delta t\, \mathcal{L}(\mathbf{U}_{n+1}, \mathbf{c}_n), \tag{33}$$

where $\mathcal{L}$ is designed to preserve total mass and energy while dissipating an entropy. Because relaxation modifies the transition model, it must be accounted for in the residual likelihood. In practice, relaxation is used as a numerical safeguard with a small coefficient, and its effect is absorbed into the process noise covariance $Q$. The key is that relaxation provides a consistent way to maintain well-posedness across closure samples without hand-tuned clipping that would break differentiability.

The combined projected update defines a stable transition map

$$\mathbf{U}_{n+1} \;=\; \mathcal{T}(\mathbf{U}_n, \mathbf{c}_n), \tag{34}$$

which can be evaluated quickly and differentiated. This map is used in two places. First, it defines the transition density in the Bayesian model, providing a likelihood term that penalizes deviations from the mechanistic update. Second, it defines the measurement prediction for computing $p(y_{n+1} \mid \mathbf{U}_{n+1})$. Because both are differentiable, the framework can compute gradients of the log posterior with respect to closure latents, enabling efficient optimization for variational inference or for maximum a posteriori initialization of ensemble filters [30].

The projection-based admissibility enforcement also clarifies the interpretation of posterior uncertainty. When the data would otherwise push the closure into an inadmissible region to match measurements, projection prevents that, and the posterior instead spreads along admissible directions. This is desirable: it indicates that the model class, including its closure representation, cannot explain the data without violating physical constraints. The correct response is not to accept an unphysical closure but to represent the mismatch as uncertainty or as increased process noise. In this way, the projection is not merely a numerical trick; it is an integral part of a statistically coherent model that respects physics.

An important computational detail is the interaction between closure projection $\Pi_{\mathcal{C}}$ and state projection $\Pi_{\mathcal{A}}$. If closure is projected but state is not, the transition may still produce negative volume fractions due to time stepping and discretization. If state is projected but closure is not, the transition may appear stable but the inferred closure may be physically meaningless and not transferable [31]. The framework therefore applies both projections consistently. Moreover, the projections are designed to be nonexpansive in appropriate norms, limiting the growth of perturbations across steps. This contributes to stable ensemble propagation and reduces filter divergence, particularly in stiff regimes where small closure changes can induce large state changes.

With a computationally stable and differentiable transition operator in place, the framework can now be evaluated on representative configurations. The next section presents computational experiments that test posterior calibration, identifiability behavior, and robustness under distribution shift, emphasizing not only success cases but also failure modes that guide practical interpretation.

## VI. Computational Evaluation and Failure Modes

Computational evaluation serves two goals. First, it tests whether the stochastic closure identification framework can recover closure-relevant information from sparse measurements without producing unphysical states or overconfident posteriors. Second, it clarifies the conditions under which closures are not identifiable and the posterior should remain broad, providing guidance for how to interpret uncertainty and how to design sensing or excitation if improved identifiability is needed [32]. The experiments are constructed to reflect realistic challenges: limited sensors, time averaging, sensor bias, closure drift, and regime-like transitions. Because the framework is methodological, the evaluation emphasizes qualitative behaviors such as calibration and robustness rather than tuning for a single dataset.

The first configuration is an annular conduit segment operated under time-varying inlet conditions and backpressure. The underlying truth model is generated by a high-resolution two-fluid solver with an entropy-stable flux and a closure that depends on a hidden interfacial structure variable. This hidden variable is not included explicitly in the inference model; instead, its effects appear as time-varying effective drag and wall shear coefficients. This setup intentionally creates model-form uncertainty: even with the correct balance laws, the inference model does not have the correct closure functional form. The measurement set consists of inlet and outlet pressure, a differential pressure over an interior portion of the segment, and one temperature measurement near the outlet that is time averaged over a finite window [33]. Additionally, a small slowly drifting offset is added to one pressure sensor to test the bias latent model.

In this configuration, the closure latent $z(t)$ is chosen to modulate an interfacial drag coefficient and a wall friction partition factor, while the residual field $r(x,t)$ modulates local friction. The prior correlation length of $r$ is set comparable to a fraction of the segment length, reflecting the belief that localized mismatch exists but is not arbitrarily oscillatory. Inference is performed with an ensemble filter that updates the augmented state and closure latents at each measurement time. The primary diagnostic is posterior predictive calibration: the fraction of observed measurements that fall within predicted credible intervals at specified nominal levels. Calibration is evaluated both for pressure signals and for derived quantities such as net pressure drop over the segment.

When closure drift is mild and within the support of the prior, the posterior over $z$ contracts relative to prior and tracks the drift with a lag determined by measurement frequency and noise [34]. The state posterior remains physically admissible throughout, as verified by positivity of phase masses and by bounded volume fractions. The credible intervals for pressure are close to nominal, indicating that the posterior accounts appropriately for measurement noise and process uncertainty. Importantly, the posterior over the residual field $r$ remains broad and near prior in regions where sensors provide little spatial information, which is expected and desirable. The framework does not attempt to invent spatially detailed closure corrections unsupported by data; instead, it represents uncertainty in those details.

As closure drift increases beyond what the nominal model can represent, a deterministic calibration would typically force closure parameters toward extreme values to match pressure drop. In contrast, the present framework exhibits a different behavior. The closure projection prevents coefficients from entering regions that would violate entropy dissipation, and the posterior instead increases in variance, particularly in the residual field and in process noise. This results in wider credible intervals that still cover the measurements [35]. While this may appear less informative, it is statistically more honest: the data indicate mismatch, and the model class cannot resolve it without violating constraints, so uncertainty must increase.

A key identifiability observation emerges when differential pressure sensors are removed, leaving only inlet and outlet pressure. In this case, global pressure drop is still observed, but spatial distribution of losses is not. The information matrix $I_T(\theta)$ becomes more ill conditioned, and the posterior over $z$ becomes correlated: increased interfacial drag and increased wall friction can both explain the same net drop. The posterior reflects this by spreading along a ridge in $(z_{\text{drag}}, z_{\text{fric}})$ space rather than collapsing to a point. This ridge is not a failure; it is a representation of non-identifiability. Moreover, when a differential pressure sensor is restored, the ridge collapses partially, and the posterior contracts in the direction that the sensor informs. This behavior provides a practical design insight: a small number of strategically placed differential measurements can dramatically improve closure identifiability for pressure-loss-related closures, whereas additional boundary pressure sensors may provide little incremental benefit [36].

The second configuration is a small network with three branches joining at a junction and feeding an outlet segment with thermal exchange. Junction loss coefficients and phase splitting behavior are uncertain and change with operating conditions. The truth model includes a junction loss that depends on a hidden variable representing mixing efficiency, which fluctuates with inlet composition. Sensors are placed only at the three inlets and at the outlet, with no direct measurements at the junction. This configuration challenges the inference model because junction closures influence downstream pressure and composition but are not directly observed. Inference includes node-associated closure variables in $z(t)$ that modulate junction loss and split fractions, with priors that allow modest drift.

The network experiments highlight how closure uncertainty propagates through topology. When inlet compositions vary, the downstream pressure response is sensitive to junction losses [37]. The posterior over the junction loss latent becomes informed indirectly through the outlet pressure. However, because outlet pressure is also influenced by downstream friction and interfacial drag, identifiability remains limited. The posterior again exhibits correlated uncertainty, indicating that data do not uniquely distinguish between junction loss increase and downstream friction increase. Adding a single pressure tap near the junction dramatically improves identifiability, but in many practical systems this is not available. The framework's value in such cases is that it quantifies the ambiguity rather than hiding it, enabling downstream decision-making or monitoring to incorporate uncertainty.

The third configuration focuses on thermal closure identification, which is often more difficult than momentum closure identification because temperature sensors are fewer and thermal dynamics can be slow. An annular segment with wall heating is simulated with a truth model in which the wall heat transfer coefficient depends on interfacial wetting state [38]. The inference model includes a closure latent that modulates wall heat partitioning between phases and a residual field that modulates local heat transfer. Measurements include inlet and outlet temperature and one mid-segment temperature averaged over time. Inference shows that thermal closures are only weakly identifiable unless transients are present. In quasi-steady operation, multiple heat transfer coefficients can yield the same outlet temperature because the thermal field adjusts slowly and because temperature integrates upstream behavior. When a transient is introduced by stepping wall heat flux, sensitivity increases and the posterior contracts. This again reinforces the view that closure identification benefits from excitation, but it also illustrates a limitation: if excitation is absent, the posterior must remain broad.

Posterior calibration is assessed not only for measured signals but also for unmeasured quantities of interest, such as predicted maximum void fraction or predicted minimum liquid holdup [39]. Because truth fields are available in simulation, one can evaluate whether credible intervals cover the truth at the appropriate rate. In the experiments, coverage is generally acceptable for integrated quantities like net pressure drop and outlet temperature. Coverage degrades for local quantities like cellwise void fraction where sensors provide limited information. This is expected, and the correct response is not to force sharper posterior intervals but to recognize that local fields are not identifiable. The framework's projection ensures that even when local fields are uncertain, they remain physically admissible, avoiding the common surrogate failure of producing negative volume fractions in regions not directly constrained by sensors.

Failure modes are informative. One failure mode occurs when the prior over the residual field $r$ is too flexible, for example if its correlation length is too short and its amplitude too large. In that case, the posterior can overfit measurement noise by creating spatially oscillatory closure corrections that match pressure signals but have no physical meaning [40]. The admissibility projection prevents the worst pathologies, but overfitting can still manifest as overly narrow posterior intervals that fail to cover truth under distribution shift. This failure is mitigated by stronger smoothness priors, by limiting the degrees of freedom of $r$ in the likelihood evaluation, and by using prior predictive checks to tune the prior amplitude.

Another failure mode arises when the admissibility constraints are overly strict or misaligned with the discretization. If the entropy constraint is too conservative, projection may distort the transition map significantly, causing the residual likelihood to misrepresent the dynamics and leading to biased inference. This highlights a practical need: the admissibility set should reflect sufficient conditions for stability without being so restrictive that it excludes physically plausible behavior. In practice, this is addressed by calibrating the entropy production bound $\Sigma$ to the discretization's numerical dissipation and by validating that projected updates remain close to unprojected updates in regimes where the model is trusted.

A third failure mode is confounding between sensor bias and closure drift [41]. If bias priors allow rapid drift, the posterior may attribute most mismatch to bias and leave closure near nominal, producing good measurement fit but poor predictive performance for unmeasured quantities. If bias priors are too tight, the posterior may attribute bias to closure, producing closure estimates that are not transferable. The framework navigates this by choosing bias prior timescales based on sensor specifications and by checking posterior correlations. When strong correlations persist, it indicates that the data do not separate these effects, and this uncertainty should be reported rather than ignored.

Distribution shift tests are also performed by changing fluid properties or boundary condition patterns outside the training envelope implied by prior tuning. Because the framework is Bayesian and not trained as a fixed predictor, robustness depends on whether the prior and transition model remain plausible under shift. Experiments show that when shift is moderate, posterior uncertainty increases but remains calibrated, with credible intervals widening appropriately. When shift is severe, for example when the truth includes physics absent from the model such as strong three-dimensional junction separation, the posterior can still remain admissible but becomes very broad [42]. This is again the correct behavior: the model cannot explain the data well, and the posterior reflects that by refusing to concentrate.

Across configurations, a consistent theme emerges: momentum-related closures tied to measurable pressure drops are more identifiable than thermal closures tied to sparse temperature sensors, and global closure modes are more identifiable than localized closure fields. The hybrid prior explicitly encodes these expectations, and the inference results align with them. The value of the approach is not that it always produces sharp closure estimates, but that it produces physically admissible posteriors whose uncertainty reflects the true limitations imposed by sparse measurements and model-form error.

## VII. Conclusion

This paper developed a stochastic closure identification framework for multiphase two-fluid network models that treats uncertain closure contributions as constrained random fields rather than as fixed parameters. The approach decomposes closure uncertainty into a low-dimensional latent process capturing regime-dependent drift and a spatially correlated residual field capturing localized mismatch, and it incorporates sensor bias as a separate latent process to mitigate confounding. A Bayesian state-space formulation combines a conservative residual likelihood with a measurement likelihood that accounts for time averaging and noise [43]. Crucially, admissibility is enforced through explicit projection in both closure and state spaces, imposing positivity, bounded volume fractions, and a discrete entropy inequality that prevents inference from exploiting unphysical energy injection to fit measurements.

The computational core is a differentiable, entropy-stable residual operator with projection that enables efficient evaluation of transition densities and measurement predictions across ensembles or gradient-based variational approximations. Identifiability is addressed explicitly through information measures that reveal which closure directions are informed by which sensors and when transient excitation increases sensitivity. Computational studies on annular and networked configurations demonstrate that the framework maintains physical admissibility, produces calibrated uncertainty for measured signals, and represents non-identifiability as posterior spread rather than as overconfident point estimates. The experiments also identify practical failure modes related to overly flexible residual priors, overly conservative admissibility constraints, and confounding between sensor bias and closure drift, emphasizing that prior design and constraint calibration are central to reliable inference.

The broader implication is that closure identification in multiphase systems should be treated as an uncertainty-aware inverse problem constrained by thermodynamics, not as a purely deterministic tuning task. When data are informative, the posterior contracts and yields transferable closure estimates; when data are insufficient, the posterior remains broad, providing a principled quantification of what cannot be learned from available sensors. This behavior is essential for downstream tasks that depend on multiphase predictions under uncertainty, including monitoring, forecasting, and safe decision-making in networks whose closures evolve with operating history [44].

## REFERENCES


[1] J. L. de Oliveira, "Modeling the dispersion phenomenon in batch transfer operations by the theory of structured mixture," *Brazilian Journal of Petroleum and Gas*, vol. 4, pp. 129–138, 12 2010.

[2] V. Kara, V. Yakhot, and K. L. Ekinci, "Generalized knudsen number for unsteady fluid flow.," *Physical review letters*, vol. 118, pp. 074505–074505, 2 2017.

[3] K. Manikonda, A. R. Hasan, C. E. Obi, R. Islam, A. K. Sleiti, M. W. Abdelrazeq, and M. A. Rahman, "Application of machine learning classification algorithms for two-phase gas-liquid flow regime identification," in *Abu Dhabi International Petroleum Exhibition and Conference*, p. D041S121R004, SPE, 2021.

[4] B. Sun, W. Shi, R. Zhang, S. Cheng, C. Zhang, S. Di, and N. Cui, "Transient behavior of vertical commingled well in vertical non-uniform boundary radii reservoir," *Energies*, vol. 13, pp. 2305–, 5 2020.

[5] R. Hu, J. Wan, Z. Yang, Y.-F. Chen, and T. K. Tokunaga, "Wettability and flow rate impacts on immiscible displacement: A theoretical model," *Geophysical Research Letters*, vol. 45, pp. 3077–3086, 4 2018.

[6] Y. Yang, D. McDermott, C. J. O. Reichhardt, and C. J. O. Reichhardt, "Dynamic phases, clustering, and chain formation for driven disk systems in the presence of quenched disorder.," *Physical review. E*, vol. 95, pp. 042902–042902, 4 2017.

[7] V. Mori, B. J. Smith, B. Suki, and J. H. T. Bates, "Linking physiological biomarkers of ventilator-induced lung injury to a rich-get-richer mechanism of injury progression.," *Annals of biomedical engineering*, vol. 47, pp. 638–645, 11 2018.

[8] J. Y. Lee, P. A. Kottke, and A. G. Fedorov, "Hydrodynamics of vortical gas jets coupled to point-like suction," *Physics of fluids (Woodbury, N.Y. : 1994)*, vol. 32, pp. 103602–, 10 2020.

[9] M. Mallach, M. Gevers, P. Gebhardt, and T. Musch, "Fast and precise soft-field electromagnetic tomography systems for multiphase flow imaging," *Energies*, vol. 11, pp. 1199–, 5 2018.

[10] K. Manikonda, R. Islam, C. E. Obi, A. R. Hasan, A. K. Sleiti, M. W. Abdelrazeq, I. G. Hassan, and M. A. Rahman, "Horizontal two-phase flow regime identification with machine learning classification models," in *International Petroleum Technology Conference*, p. D011S021R002, IPTC, 2022.

[11] Y. Camaraza-Medina, K. Mortensen-Carlson, P. Guha, Ángel Rubio-Gonzales, O. M. Cruz-Fonticiella, and O. García-Morales, "Suggested model for heat transfer calculation during fluid flow in single phase inside pipes (ii)," *International Journal of Heat and Technology*, vol. 37, pp. 257–266, 3 2019.

[12] K. Bybee, "Analysis of production data from hydraulically fractured horizontal wells," *Journal of Petroleum Technology*, vol. 60, pp. 78–80, 3 2008.

[13] J. Oosterhuis, S. Bühler, D. Wilcox, and T. H. van der Meer, "Jet pumps for thermoacoustic applications: design guidelines based on a numerical parameter study," *The Journal of the Acoustical Society of America*,

vol. 138, pp. 1991–2002, 10 2015.

[14] S. Bachman, "A geometric perspective on the modulation of potential energy release by a lateral potential vorticity gradient," *Fluids*, vol. 5, pp. 142–, 8 2020.

[15] N. Abbas, S. A. Wasimi, N. Al-Ansari, and N. Baby, "Recent trends and long-range forecasts of water resources of northeast iraq and climate change adaptation measures," *Water*, vol. 10, pp. 1–19, 11 2018.

[16] B. J. Akiode, V. Bhushan, R. Lee, P. Mukerji, and Z. Zaouali, "The value of surveillance and advanced logging applications for brownfield optimization," *SPE Offshore Europe Oil and Gas Conference and Exhibition*, 9 2007.

[17] C. S. Campbell, "Granular flows and gas fluidization," *International Journal of Chemical Reactor Engineering*, vol. 2, 7 2004.

[18] B. Lukitsch, R. Koller, P. Ecker, M. Elenkov, C. Janeczek, M. Pekovits, B. Haddadi, C. Jordan, M. Gfoehler, and M. Harasek, "Water as a blood model for determination of co2 removal performance of membrane oxygenators," *Membranes*, vol. 11, pp. 356–, 5 2021.

[19] S. Wang, S. Wang, K. Zhu, Y. Xiao, and Z. Lu, "Near quenching limit instabilities of concurrent flame spread over thin solid fuel," *Combustion Science and Technology*, vol. 188, pp. 451–471, 12 2015.

[20] M. Ramdin and R. Henkes, "Computational fluid dynamics modeling of benjamin and taylor bubbles in two-phase flow in pipes," *Journal of Fluids Engineering*, vol. 134, pp. 041303–, 4 2012.

[21] D. Biryukov and F. J. Kuchuk, "Pressure transient behavior of horizontal wells intersecting multiple hydraulic fractures in naturally fractured reservoirs," *Transport in Porous Media*, vol. 110, pp. 369–408, 9 2015.

[22] J. H. Vrijsen, A. R. Mazo, T. Junkers, and G. G. Qiao, "Accelerated polypeptide synthesis via n-carboxyanhydride ring opening polymerization in continuous flow.," *Macromolecular rapid communications*, vol. 41, pp. 2000071–, 7 2020.

[23] O. Taylan, M. A. Sattari, I. E. Essoussi, and E. Nazemi, "Frequency domain feature extraction investigation to increase the accuracy of an intelligent nondestructive system for volume fraction and regime determination of gas-water-oil three-phase flows," *Mathematics*, vol. 9, pp. 2091–, 8 2021.

[24] M. Wang, X. Dou, R. Ming, L. Weiqiang, W. Zhao, and C. Tan, "Semi-analytical rate decline solutions for a refractured horizontal well intercepted by multiple reorientation fractures with fracture face damage in an anisotropic tight reservoir," *Energies*, vol. 14, pp. 7482–, 11 2021.

[25] I. Andreolli, L. A. Bernardo, and R. H. Ruschel, "Petroleum exportation through pipelines: Important considerations," *Brazilian Journal of Petroleum and Gas*, vol. 9, pp. 55–71, 6 2015.

[26] F. Zhu, Z. Lu, and S. Wang, "Flame spread and extinction over a thick solid fuel in low-velocity opposed and concurrent flows," *Microgravity Science and Technology*, vol. 28, pp. 87–94, 11 2015.

[27] Y. Wang, H. Wang, S. Qi, S. Liu, Y. Zhao, and W. Yue, "Investigation of shale permeability evolution considering bivalued effective stress coefficients for co2 injection," *Geofluids*, vol. 2021, pp. 1132440–11, 10 2021.

[28] J. L. Moody, W. C. Keene, O. R. Cooper, K. J. Voss, R. Aryal, S. Eckhardt, B. N. Holben, J. R. Maben, M. A. Izaguirre, and J. N. Galloway, "Flow climatology for physicochemical properties of dichotomous aerosol over the western north atlantic ocean at bermuda," *Atmospheric Chemistry and Physics*, vol. 14, pp. 691–717, 1 2014.

[29] M. Yue, X. Huang, F. He, L. Yang, W. Zhu, and Z. Chen, "Analysis of the influence of different fracture network structures on the production of shale gas reservoirs," *Geofluids*, vol. 2020, pp. 1–11, 12 2020.

[30] N. Basha, A. Kovacevic, and S. Rane, "Analysis of oil-injected twin-screw compressor with multiphase flow models," *Designs*, vol. 3, pp. 54–, 12 2019.

[31] I. Belysheva, I. V. Mursenkova, and A. Chvyreva, "Experimental research of sliding surface distributed nanosecond discharge in supersonic air flow," *Journal of Physics: Conference Series*, vol. 516, pp. 012021–, 6 2014.

[32] "Numerical study of heat transfer enhancement for low-pressure flows in a square cavity with two fins attached to the hot wall using al2o3-air nanofluid," *Strojniški vestnik - Journal of Mechanical Engineering*, vol. 64, 1 2018.

[33] V. Hejtmánek and P. Čapek, "3d microstructure modeling of porous metal filters," *Metals*, vol. 2, pp. 344–352, 9 2012.

[34] Z. Yan, C. Cao, M. Xie, P. Jia, G. Liu, X. Wei, L. Cheng, and T. Zhiyong, "Pressure behavior analysis of permeability changes due to sand production in offshore loose sandstone reservoirs using boundary-element method," *Geofluids*, vol. 2021, pp. 1–10, 2 2021.

[35] M. Kellenberger, C. T. Johansen, G. Ciccarelli, and F. Zhang, "Dense particle cloud dispersion by a shock wave," *Shock Waves*, vol. 23, pp. 415–430, 3 2013.

[36] I. Pană, I. V. Gheţiu, I. G. Stan, F. Dinu, G. Brănoiu, and S. Suditu, "The use of hydraulic fracturing in stimulation of the oil and gas wells in romania," *Sustainability*, vol. 14, pp. 5614–5614, 5 2022.

[37] J. Franc, P. Horgue, R. Guibert, and G. Debenest, "Benchmark of different cfl conditions for impes," *Comptes Rendus. Mécanique*, vol. 344, pp. 715–724, 9 2016.

[38] H. Vogtenhuber, D. Pernsteiner, and R. Hofmann, "Experimental and numerical investigations on heat transfer of bare tubes in a bubbling fluidized bed with respect to better heat integration in temperature swing adsorption systems," *Energies*, vol. 12, pp. 2646–, 7 2019.

[39] S. Shadkam, F. Ludwig, M. T. H. van Vliet, A. Pastor, and P. Kabat, "Preserving the world second largest hypersaline lake under future irrigation and climate change.," *The Science of the total environment*, vol. 559, pp. 317–325, 4 2016.

[40] T. Persoons, E. V. den Bulck, and S. Fausto, "Study of pulsating flow in close-coupled catalyst manifolds using phase-locked hot-wire anemometry," *Experiments in Fluids*, vol. 36, pp. 217–232, 2 2004.

[41] M. M. Cui and J. Sauls, "Investigation on the oil supply system of a scroll compressor," *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, vol. 221, pp. 509–517, 6 2007.

[42] P. Aussillous, Y. Zhou, P. Ruyer, and P.-Y. Lagrée, "Discharge flow of granular media from silos with a lateral orifice and injection of air," *EPJ Web of Conferences*, vol. 140, pp. 03052–, 6 2017.

[43] S. Huang, Y. Wu, L. Cheng, H. Liu, Y. Xue, and G. Ding, "Apparent permeability model for shale gas reservoirs considering multiple transport mechanisms," *Geofluids*, vol. 2018, pp. 1–18, 6 2018.

[44] P. Berloff, S. A. Karabasov, J. T. Farrar, and I. Kamenkovich, "On latency of multiple zonal jets in the oceans," *Journal of Fluid Mechanics*, vol. 686, pp. 534–567, 9 2011.