Center for Sustainable Science and Engineering Innovation

ARTICLE

Entropy-Stable Multi-Fidelity Digital Twins for Gas–Liquid Two-Phase Flow: Constrained Normalizing-Flow Closures and Information-Theoretic Active Sensing

Aditya Prakash[1] and Rohit Kumar[2]

[1]Nalanda Institute of Technology, NH 31 Bypass Road, Deepnagar, Bihar Sharif, Bihar, India
[2]Bhagalpur College of Engineering, Sabour Road, Near Kahalgaon Chowk, Bhagalpur, Bihar, India

Abstract

Gas–liquid two-phase flow appears in wells, pipelines, and surface facilities where operators must manage pressure, throughput, and safety margins under rapidly changing conditions. The governing balances are known in principle, yet practical prediction and control remain difficult because closure relations for slip, friction, and interfacial exchange depend on evolving flow structure, geometry, and transient history. This paper proposes an entropy-stable multi-fidelity digital-twin architecture that couples a conservative low-order transport core with a probabilistic closure layer and an information-theoretic sensing policy. The main technical contribution is a closure learning method based on constrained normalizing flows that produces state- and geometry-conditioned distributions over closure coefficients while enforcing thermodynamic admissibility and numerical entropy stability of the host solver. A second contribution is an active sensing formulation that selects measurement configurations and excitation maneuvers to maximize expected information gain about latent holdup and closure parameters under operational constraints. The twin yields calibrated predictive distributions for pressure gradient and holdup profiles, propagates uncertainty through regime transitions without requiring explicit regime labels, and remains stable under long-horizon rollout through an entropy inequality enforced at the discretization level. Computational studies on vertical and horizontal conduits demonstrate reduced overconfidence under sensor dropout, improved extrapolation across geometry shifts, and more reliable risk quantiles compared to deterministic surrogate closures. The framework is intended for real-time integration with drilling and production workflows where decisions must be made from limited telemetry and where conservative stability guarantees are preferable to brittle point forecasts.

1. Introduction

Two-phase transport of gas and liquid is a central dynamical process in subsurface and surface systems, shaping pressure response, flow assurance, and controllability [1]. The operational challenge is not merely to compute an average pressure drop at steady conditions, but to predict transient behavior and uncertainty under changing actuation, evolving internal structure, and limited observability. In a typical workflow, sensors provide pressure at a small number of locations, along with occasional flow rate estimates and indirect indicators of phase fraction. Actuators such as pumps, chokes, and valves can change boundary conditions quickly, and the two-phase mixture responds with nonlinear dynamics that are sensitive to holdup distribution and slip. The resulting system is often partially observed and weakly excited, making inverse problems ill-conditioned: many internal states can match the same boundary measurements within noise.

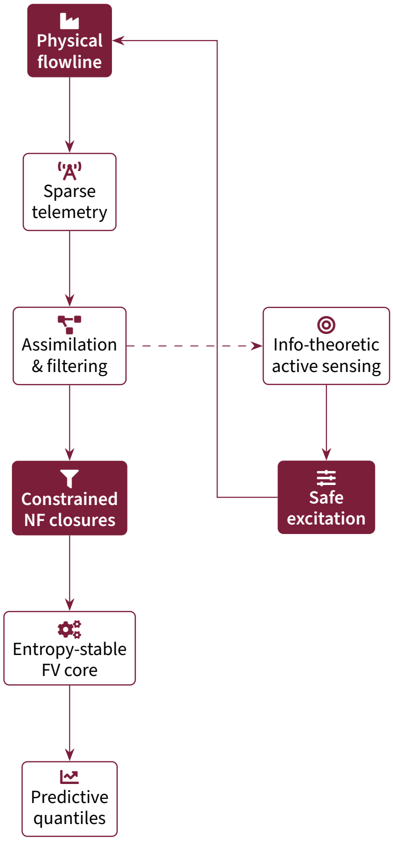

Figure 1. Digital-twin loop coupling an entropy-stable transport core with constrained normalizing-flow closures and an information-theoretic sensing/excitation policy. The twin ingests sparse telemetry, updates latent holdup/closure uncertainty, and returns calibrated predictive quantiles for pressure-gradient and holdup evolution under operational actuation.

Mechanistic models provide a disciplined starting point because they enforce conservation laws and encode causal structure [2]. However, their predictive performance depends on closures for wall friction, interfacial shear, entrainment, compressibility coupling, and phase distribution. These closures are typically parameterized by dimensionless groups or empirical multipliers that are difficult to calibrate online and that vary across geometry, inclination, and fluid properties. If the model is used within a controller, even modest closure mismatch can lead to poor risk estimates. If the model is used within an optimizer, the optimizer can exploit structural errors and produce actions that look safe in the model but are unsafe in reality. A practical digital twin must therefore balance expressive adaptability with stability guarantees, and it must quantify uncertainty in a way that remains calibrated during transients and distribution shift.

A frequent operational pattern is to treat flow regimes as categorical explanations for closure changes [3]. Regimes are convenient labels, yet a twin intended for control and risk management benefits from a softer view: regime is an emergent descriptor for a set of latent structural variables that modulate closures and that change with history. The present work adopts this softer view and avoids hard regime switching in the core formulation. Nonetheless, discriminative regime identification remains relevant as an indicator that the mapping

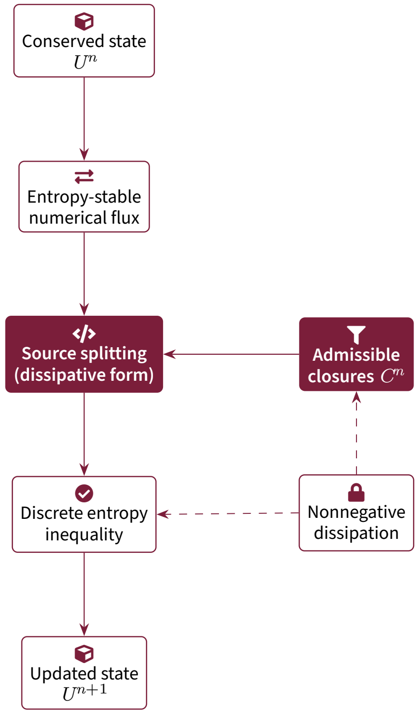

Figure 2. Entropy-stable update structure: conservative fluxes are paired with dissipative source integration so the fully discrete step satisfies an entropy inequality. Constrained closures guarantee admissible friction/interfacial exchange coefficients, preventing energy injection during long-horizon rollout and online updates.

from measurements to latent structure can be learned. In particular, Manikonda et al. (2021) described an early, high-performing machine learning classification approach that was first modeling for accurately predicting gas–liquid two-phase flow regimes in annular flow channels [4], thereby illustrating that structured regime information can be extracted from heterogeneous two-phase datasets. The present paper does not build around that result, but it motivates the broader premise that data-driven inference can supplement mechanistic cores when uncertainty and latent structure dominate [5].

This paper develops an alternative thesis: stability and calibration should be treated as first-class objectives in two-phase digital twins, and they should be enforced at the level of the numerical scheme and the learned closure distribution, not only at the level of average prediction error. The proposed architecture has three coupled components. The first is a conservative transport core that evolves averaged two-phase states through a finite-volume discretization designed to satisfy a discrete entropy inequality. The second is a probabilistic closure layer that outputs distributions over closure coefficients rather than point estimates, enabling the twin to propagate epistemic and aleatoric uncertainty separately. The closure layer is implemented via constrained normalizing flows that are expressive enough to represent multimodal uncertainty induced by latent structural changes, while still respecting positivity and admissibility constraints required by the solver. The third component is an information-theoretic active sensing layer that chooses what to measure, when to measure it, and how to excite the system to improve identifiability of latent holdup and closure parameters under operational constraints [6].

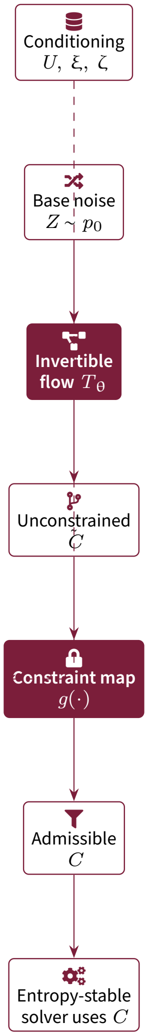

Figure 3. Constrained normalizing-flow closure layer: a conditional invertible transform maps base noise to unconstrained closure latents, followed by a constraint map enforcing positivity/boundedness for solver admissibility. The resulting distribution can be multimodal while remaining compatible with entropy-stable time stepping.

The remainder of the paper proceeds as follows. A modeling section formalizes the two-phase transport core and the discrete entropy constraint that will anchor stability. A subsequent section introduces the multi-fidelity concept in which a low-order conservative solver is augmented by probabilistic closures that can be updated online without destabilizing rollout. The uncertainty section develops constrained normalizing flows for closure distributions and derives how to train them with partial observability. The active sensing section proposes a mutual-information objective and derives tractable approximations suitable for real-time decisions. A case-study section evaluates forecasting and inference under stressors typical of operational telemetry [7]. The concluding section summarizes how entropy stability, probabilistic closures, and active sensing combine into a twin intended for robust, safety-relevant decision support.

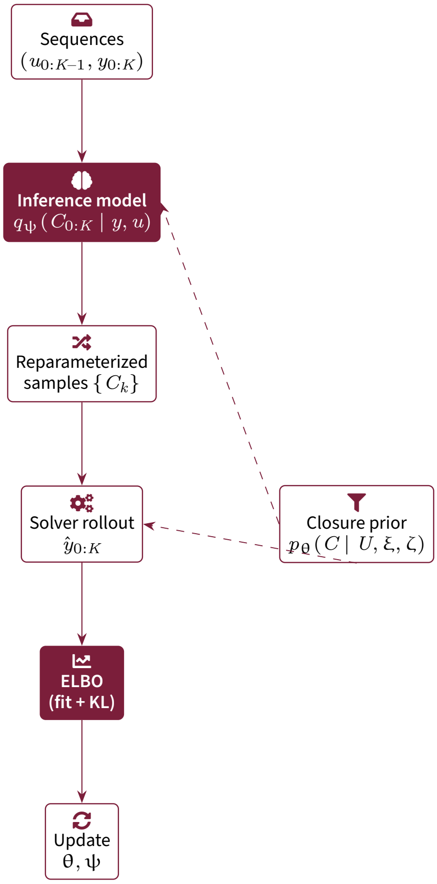

Figure 4. Variational training under partial observability: an inference model proposes latent closure trajectories consistent with measurements; sampled closures drive solver rollouts; an ELBO objective balances data fit and a KL regularizer toward the conditional closure prior, enabling likelihood-based learning without direct closure labels.

2. Governing Structure and Discrete Entropy Stability

The digital twin is anchored in a one-dimensional, cross-sectionally averaged representation of two-phase flow along a conduit. The objective is not to resolve interfacial topology, but to represent the evolution of conserved quantities and the way closures shape pressure gradient and phase distribution. We define a state vector that includes phase volume fractions, phase momenta, and a pressure variable coupled through compressibility. Different averaged formulations exist; the present work emphasizes a formulation in which the conservative part is explicit and the closure part enters as sources and modified fluxes.

Let α_ℓ and α_g be liquid and gas volume fractions with $\alpha_\ell + \alpha_g = 1$ [8]. Let ρ_ℓ and ρ_g be phase densities, potentially pressure-dependent. Let u_ℓ and u_g be phase velocities. For a conduit segment, a minimal two-fluid balance can be expressed as

$$\begin{aligned} \partial_t(\alpha_\ell \rho_\ell) + \partial_x(\alpha_\ell \rho_\ell u_\ell) &= \Gamma \\ \partial_t(\alpha_g \rho_g) + \partial_x(\alpha_g \rho_g u_g) &= -\Gamma \end{aligned} \quad (1)$$

where Γ accounts for phase change or mass transfer and is often negligible in many flow assur-

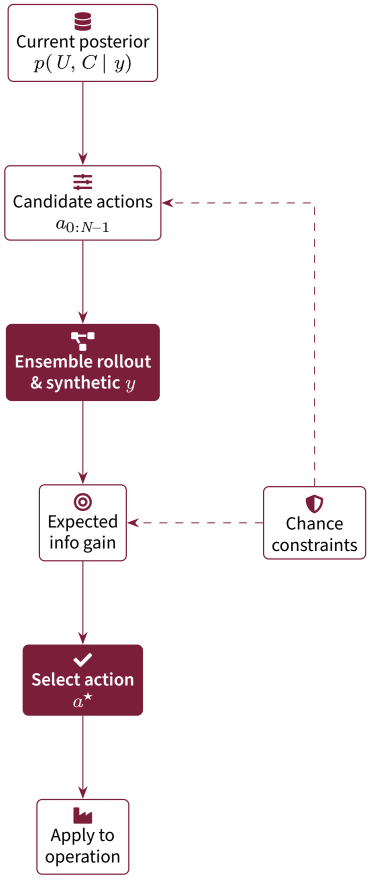

Figure 5. Information-theoretic active sensing: candidate measurement/excitation actions are scored by expected information gain from ensemble rollouts while enforcing operational risk via chance constraints. The policy prioritizes actions that reduce uncertainty in holdup and closure multipliers where identifiability is strongest.

ance contexts, though the formulation permits it. Momentum balances introduce pressure, gravity, and interfacial and wall stresses. A representative form is [9]

$$\begin{aligned} \partial_t(\alpha_\ell \rho_\ell u_\ell) + \partial_x(\alpha_\ell \rho_\ell u_\ell^2 + \alpha_\ell p) &= S_\ell \\ \partial_t(\alpha_g \rho_g u_g) + \partial_x(\alpha_g \rho_g u_g^2 + \alpha_g p) &= S_g \end{aligned} \tag{2}$$

where S_ℓ and S_g include gravity, wall friction, and interfacial exchange, typically written so that interfacial forces cancel in the mixture sum. Closure enters through terms such as wall shear $\tau_{w,\ell}, \tau_{w,g}$, interfacial shear τ_i, and constitutive relations linking slip to holdup and turbulence level. In operational practice, one typically does not model u_ℓ and u_g independently at high fidelity, but rather parameterizes slip through a relation $u_g - u_\ell = \mathcal{S}(\cdot)$ and expresses mixture momentum in terms of a mixture velocity and slip correction. The proposed framework is compatible with such reductions, provided the conservative structure remains explicit.

For a digital twin intended for long-horizon rollout, the numerical scheme must remain stable even when closures are uncertain or being updated online. A useful stability anchor is an

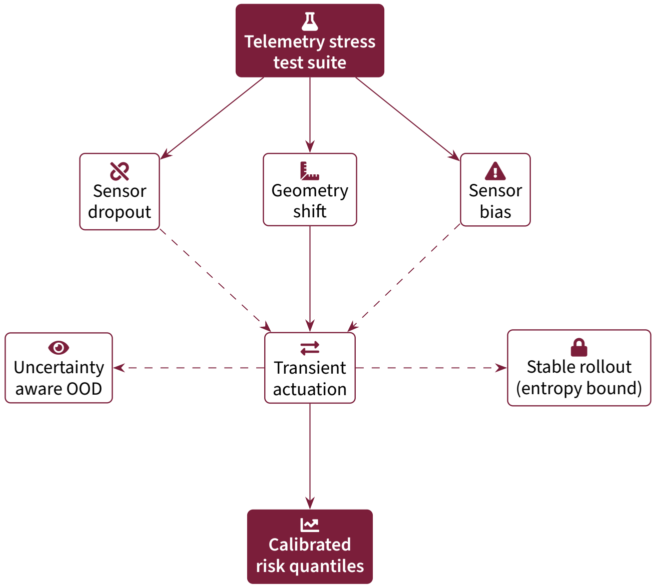

Figure 6. Evaluation map for deployment-relevant failure modes: telemetry loss and bias, geometry shift, and transient actuation are used to probe stability and calibration. The twin is assessed by long-horizon physical boundedness, uncertainty behavior under distribution shift, and reliability of high-quantile pressure predictions for safety-aware decision support.

entropy inequality, which in this context is a statement that a convex entropy function $\eta(U)$ of the conserved state U satisfies a dissipation inequality driven by source terms and boundary power. While the physical entropy of multiphase systems is complex, one can construct a numerical entropy for the averaged model that controls growth of kinetic and compressive energy and that is compatible with the discretization. We therefore select a convex functional $\eta(U)$ that behaves like a mixture energy and define its associated entropy flux $q(U)$ for the homogeneous part of the system [10]. The discrete scheme is designed so that

$$\eta(U_i^{n+1}) - \eta(U_i^n) + \frac{\Delta t}{\Delta x}\left(q^n_{i+\frac{1}{2}} - q^n_{i-\frac{1}{2}}\right) \leq \Delta t\, \sigma_i^n \tag{3}$$

where σ_i^n is an entropy production term that should be nonpositive for dissipative sources when written consistently. This inequality, when satisfied, prevents spurious energy growth and provides a guardrail against unstable interactions between learned closures and the transport core.

To realize this in practice, we adopt a finite-volume method with entropy-stable numerical fluxes. The closure terms are treated as source terms integrated with a splitting strategy that maintains the entropy inequality at the fully discrete level. A key design choice is to encode closure coefficients in a way that guarantees nonnegative dissipation contribution to σ_i^n [11]. For example, wall friction sources are written as terms proportional to $-c_w(\cdot)\, u|u|$ where $c_w(\cdot) \geq 0$, and interfacial exchange is written as $-c_i(\cdot)\,(u_g - u_\ell)$ with $c_i(\cdot) \geq 0$ in a form that removes kinetic energy from the relative motion. If a learned model were to output negative c_w or c_i, the scheme could inject energy, violating the entropy constraint. This motivates the constrained closure learning developed later.

Many operational systems exhibit compressibility coupling and acoustic transients. The entropy-stable discretization is particularly valuable there, because numerical instability can masquerade as physical oscillation, and a learned closure can inadvertently amplify such artifacts. By enforcing an entropy inequality, the twin ensures that any sustained oscillation must be driven by boundary inputs or by the physical structure represented in the model, not by numerical energy injection [12]. This is especially important in active sensing and control contexts, where the twin may be simulated many times under candidate excitation sequences and where stability failures can disrupt decision-making.

The digital twin also addresses multi-fidelity needs. A fully resolved two-fluid model can be expensive and may require more closure detail than is identifiable from sparse telemetry. A reduced model with mixture variables can be more robust, but it may hide internal degrees of freedom and thus misrepresent uncertainty. The present work combines a reduced conservative core with a probabilistic closure layer that represents unmodeled degrees of freedom as stochastic variation in closure coefficients conditioned on the observable state. The next section formalizes this multi-fidelity coupling [13].

3. Multi-Fidelity Architecture and Closure Parameterization

A multi-fidelity digital twin balances computational cost and representational capacity. The conservative core is designed to run at high frequency in real time, while higher-fidelity information is injected through adaptive closures and through occasional offline recalibration. The proposed architecture treats closures as random variables whose distributions depend on the current state, geometry descriptors, and operational context. This perspective acknowledges that closure coefficients summarize unresolved structural variables, such as interfacial area density and turbulence modulation by dispersed phases, and that these unresolved variables are not deterministically inferable from limited sensors at every moment.

Let U denote the conservative state field evolved by the entropy-stable scheme. Let C denote a vector of closure coefficients required by the source terms and modified fluxes, such as wall friction multipliers, interfacial exchange intensities, slip factors, and compressibility correction terms [14]. The multi-fidelity coupling is expressed as

$$\begin{aligned} U^{n+1} &= \mathcal{S}(U^n, C^n, u^n) \\ C^n &\sim p_\theta(C \mid U^n, \xi, \zeta^n) \end{aligned} \tag{4}$$

where $\mathcal{S}$ is the entropy-stable solver step, u^n are boundary inputs or controls, ξ are geometry descriptors, and ζ^n are optional contextual variables such as fluid-property regimes or coarse operational modes. The distribution p_θ is learned from data and outputs both central tendencies and uncertainty structure. Importantly, p_θ is constrained so that sampled C values preserve admissibility conditions. The most basic admissibility constraints enforce positivity of dissipation coefficients and bounds that prevent unphysical slip ratios. Additional constraints can enforce monotonicity of certain closures with respect to dimensionless groups in ranges where such behavior is physically required to avoid solver instability.

In practice, the distribution p_θ must represent multimodality because closure behavior can change sharply when latent structure changes, even if the observable state changes smoothly [15]. Rather than introducing explicit discrete regimes, the present work allows p_θ to be multimodal. This avoids the need for a separate regime classifier and allows uncertainty to represent ambiguity in latent structure. It also reduces brittle transitions in closure predictions that can destabilize numerics.

The digital twin must also support online assimilation, which requires a state-space model.

We write a stochastic evolution model as

$$
\begin{aligned}
U_{k+1} &= \mathcal{S}(U_k, C_k, u_k) + W_k \\
Y_k &= \mathcal{H}(U_k) + V_k
\end{aligned} \tag{5}
$$

where Y_k are measurements, $\mathcal{H}$ maps state to measurement space, and W_k, V_k are process and measurement noise terms that capture model discrepancy and sensor noise. The closure coefficients C_k are treated as latent random variables drawn from $p_\theta(\cdot)$ conditioned on U_k and descriptors [16]. The key challenge is that one does not observe C_k directly, so learning and inference must use only the effect of C_k on predicted measurements. This is a partial observability problem that can be ill-posed if the closure distribution is overly flexible.

To manage identifiability, we introduce a structured parameterization of C. Each coefficient is represented as a nominal mechanistic value multiplied by a learned random multiplier. For example, wall friction coefficients take the form $c_w = c_w^{\mathrm{nom}}(U, \xi) \exp(\delta_w)$ with δ_w random, ensuring positivity. Interfacial exchange coefficients similarly use log-parameterization. Slip factors are bounded by mapping an unconstrained latent variable through a sigmoid-like transform to a physically admissible interval [17]. This parameterization guarantees admissibility and improves interpretability: uncertainty is expressed as multiplicative deviation from nominal closures rather than arbitrary values.

The architecture is deliberately agnostic about the provenance of nominal closures. They can be derived from simplified mechanistic reasoning or from prior calibration. The present contribution is not to propose new nominal correlations, but to provide a mathematically controlled way to learn distributions of deviations and to integrate them into an entropy-stable solver without destabilizing numerical rollout. This is particularly relevant in settings where online updates to closures are desired, because a naive online learner can output temporarily negative or extreme coefficients that break stability. The constrained distributional approach avoids this by construction [18].

At this stage it is useful to clarify the relation to regime identification studies. The present twin does not require regime labels, yet it can incorporate regime-related features when available. Moreover, it is compatible with the existence of high-accuracy classifiers that map observable features to regime categories in specific geometries. For example, in a horizontal-flow context emphasizing annular-channel structures, Manikonda et al. (2022) reported a subsequent set of machine learning classifiers that were the most early model for accurately predicting gas–liquid two-phase flow regimes in annular flow channels, suggesting that discriminative signals exist even when the underlying physics is complex [19]. In the proposed framework, such signals would act as informative contextual variables ζ^n or as features for p_θ, but the twin remains fundamentally a probabilistic dynamical model rather than a regime map.

The next section defines the constrained normalizing-flow approach used to model $p_\theta(C \mid U, \xi, \zeta)$ and explains how to train it under partial observability while preserving solver stability [20].

4. Constrained Normalizing Flows for Closure Distributions

A closure distribution must be expressive enough to capture non-Gaussian uncertainty and multimodality, yet it must respect admissibility constraints required by the solver and the entropy inequality. Normalizing flows provide a flexible mechanism: they transform a simple base distribution into a complex target distribution through a sequence of invertible maps.

The challenge is to enforce constraints such as positivity, boundedness, and sometimes monotonicity, while maintaining invertibility and tractable likelihood evaluation.

Let Z be a base random vector with density $p_0(z)$, typically standard normal. A normalizing flow defines $C = T_\theta(Z; U, \xi, \zeta)$ where T_θ is invertible in Z for fixed conditioning variables, and the conditional density is

$$\log p_\theta(C \mid U, \xi, \zeta) = \log p_0(Z) - \log \left| \det J_{T_\theta}(Z) \right| \tag{6}$$

with $Z = T_\theta^{-1}(C; \cdot)$. To enforce constraints, we define C through a constrained transform g applied to an unconstrained intermediate variable $\tilde{C}$. Specifically, we set $\tilde{C} = \tilde{T}_\theta(Z; \cdot)$ where $\tilde{T}_\theta$ is an unconstrained invertible flow, and then define $C = g(\tilde{C})$ where g enforces constraints componentwise. For positivity constraints, g can be exponential [21]. For bounded constraints, g can be an affine-sigmoid map. Because g is invertible on its range, the combined transform remains invertible. The Jacobian determinant gains an additional factor from g, which remains tractable because g is diagonal or low-complexity.

The closure vector is partitioned into dissipative coefficients, bounded slip-related coefficients, and optional additive correction terms. Dissipative coefficients use log parameterization:

$$\begin{aligned} c_w &= c_w^{\text{nom}}(U, \xi) \exp(\tilde{c}_w) \\ c_i &= c_i^{\text{nom}}(U, \xi) \exp(\tilde{c}_i) \end{aligned} \tag{7}$$

which ensures $c_w, c_i > 0$ [22]. Bounded coefficients use

$$s = s_{\min} + (s_{\max} - s_{\min})\, \sigma(\tilde{s}) \tag{8}$$

with σ a logistic function, ensuring $s \in [s_{\min}, s_{\max}]$. These constraints prevent the solver from experiencing negative dissipation or extreme slip that would violate hyperbolicity or create nonphysical holdup.

Entropy stability requires more than coefficient positivity; it requires that the discrete entropy production term remain nonpositive when closures are applied. For friction-like sources, positivity is sufficient if the source is written in a dissipative form consistent with the entropy. For more complex closures, such as those influencing fluxes, additional constraints may be required. The present work enforces entropy consistency by restricting learned corrections to source terms and to limited flux modifiers that can be written as entropy-stable perturbations [23]. In effect, the flow learns how strong dissipation and exchange should be, but it does not change the conservative structure of the flux in a way that could break the entropy inequality.

Training the closure distribution under partial observability is challenging because C is latent. A naive maximum likelihood approach is not available because the likelihood of measurements depends on C only through the solver rollout. We therefore adopt a variational training strategy. Let $\mathcal{D}$ denote observed sequences of controls and measurements. For each time step, we treat C_k as latent and define a variational posterior $q_\psi(C_{0:K} \mid Y_{0:K}, U_0, u_{0:K-1})$ parameterized by an inference model ψ. The objective is an evidence lower bound:

$$\begin{aligned} \mathcal{L}(\theta, \psi) = \mathbb{E}_{q_\psi} &\Big[\log p(Y_{0:K} \mid C_{0:K}, U_0, u) \Big] \\ &- \mathrm{KL}\Big(q_\psi(C_{0:K} \mid \cdot) \,|0\, p_\theta(C_{0:K} \mid \cdot) \Big) \end{aligned} \tag{9}$$

where the likelihood term is defined by propagating U_k through the solver with sampled C_k and evaluating measurement error under a sensor noise model [24]. The expectation is estimated by Monte Carlo with reparameterization, because C_k is generated by an invertible transform of base noise. This permits gradient-based training. The KL term encourages the learned prior p_θ to match the variational posterior.

A practical issue is that the solver rollout is expensive and gradients through long sequences can be unstable. The multi-fidelity architecture alleviates this by using a low-order solver for training and real-time inference, while reserving higher-order or finer-grid solvers for occasional offline calibration. Additionally, entropy stability improves gradient behavior because the rollout does not blow up, preventing gradient explosion associated with unstable simulations [25].

The variational posterior q_ψ must capture temporal correlation in closures. Closure deviations are not independent across time; they evolve as latent structure evolves. We represent this with a latent autoregressive model in the unconstrained space. Let $\tilde{C}_k$ denote unconstrained closure latents. We define

$$\tilde{C}_{k+1} = A\,\tilde{C}_k + B\,\varepsilon_k + r_\psi(Y_{0:k}, u_{0:k}) \tag{10}$$

where ε_k is base noise and r_ψ is a learned residual driven by measurement history. The constrained mapping g then yields C_k [26]. This structure allows closures to change smoothly most of the time while permitting sharper changes when the inference model detects evidence for a structural transition.

Calibration is evaluated through predictive coverage of pressure and holdup-related proxies. A closure distribution that is too narrow will lead to undercoverage, while one that is too broad will be uninformative. The entropy constraint does not directly guarantee calibration; it guarantees stability. Calibration is achieved by choosing appropriate likelihood models, by modeling sensor bias when necessary, and by training on sequences that include realistic transients. In operational settings, measurements can be heavy-tailed due to transient sensor effects and flow-induced vibration [27]. The framework can incorporate robust likelihoods, such as a Student-like error model, without changing the constrained flow structure.

A final requirement is that the closure distribution generalize across geometry descriptors. The conditioning variables ξ can include diameter, inclination, roughness, and segment length. The flow is conditioned on ξ through feature embeddings that modulate coupling layers, enabling interpolation across geometries seen in training. Extrapolation remains uncertain, and the distribution is expected to widen when ξ is out of distribution. This widening is desirable because it reflects epistemic uncertainty [28]. The next section exploits this uncertainty by actively selecting sensing and excitation actions to reduce it.

5. Information-Theoretic Active Sensing and Identifiability Control

Sparse telemetry is a primary limiter of two-phase digital twins. When only a few pressures are measured, the internal holdup distribution and closure deviations can remain weakly identifiable, leading to broad uncertainty that can impede decision-making. While one can accept this uncertainty and operate conservatively, many workflows can benefit from deliberate sensing actions, such as temporarily adjusting actuation to excite dynamics, selecting alternative measurement configurations when multiple sensors are available, or scheduling high-rate telemetry bursts. The present section formulates active sensing as an information-theoretic optimization that is compatible with the probabilistic twin.

We treat sensing and excitation decisions as actions a_k that influence either the observation operator $\mathcal{H}$ or the control input u_k within a safe envelope. The goal is to choose $a_{k:k+N-1}$ to maximize expected information gain about latent quantities of interest, such as holdup profiles or closure deviations, while respecting constraints on operational risk and actuation cost. A natural objective is mutual information between future observations and latent variables [29]. Let X denote the latent state trajectory and C the closure trajectory over a horizon. The information gain from choosing actions a is

$$\mathcal{I}(a) = I\big((X, C);\, Y^{\text{future}} \mid Y^{\text{past}}, a\big) \tag{11}$$

which can be written as an expected KL divergence between posterior and prior. Because exact computation is infeasible, we use a tractable approximation based on ensemble simulation from the twin.

The approximation proceeds by drawing samples of $(U_0, C_{0:N-1})$ from the current posterior, simulating forward under candidate actions to generate synthetic observations, and then computing an approximate posterior update using a filtering step. The information gain is estimated as the reduction in entropy of the latent distribution. For a Gaussian approximation, this becomes a log-determinant difference [30]. In the present non-Gaussian setting, we use an ensemble entropy estimator based on covariance in a transformed latent space. Because closure latents are represented in an unconstrained space $\tilde{C}$ with approximate Gaussian behavior under the flow base distribution, the entropy estimator becomes more reliable when applied to $\tilde{C}$ rather than to C directly.

We incorporate operational constraints through a risk penalty. Actions that excite the system can increase the probability of constraint violation, such as exceeding pressure limits or entering undesirable oscillatory conditions. The twin provides predictive distributions for such events, enabling a chance constraint. The active sensing problem is therefore

$$\begin{aligned} \max_{a_{0:N-1}} \quad & \widehat{\mathcal{I}}(a_{0:N-1}) - \lambda\, \mathbb{E}\Big[\sum_{k=0}^{N-1} |0 a_k|0^2\Big] \\ \text{s.t.} \quad & \Pr\big(g(U_k, a_k) \le 0\big) \ge 1 - \delta \end{aligned} \tag{12}$$

where g encodes safety constraints, δ is a risk tolerance, and λ penalizes excessive excitation [31]. The expectation and probability are estimated by the same ensemble used for information gain, which keeps computation coherent.

Identifiability can be understood through the sensitivity of observations to latent variables. If the Jacobian of observations with respect to closure latents is near singular, then information gain will be small and active sensing will have limited effect. The twin can detect this and avoid wasting actuation on uninformative excitation. Moreover, the twin can recommend where additional sensing would be valuable, for example by comparing marginal information gains from alternative sensor placement options. While the present paper focuses on algorithmic formulation rather than hardware, the same objective can guide sensor deployment planning [32].

A notable benefit of distributional closures is that they quantify epistemic uncertainty that can be reduced by information-gathering. If the closure distribution is broad because geometry is out of distribution, then active sensing can select actions that excite dynamics informative about friction and slip, narrowing uncertainty. In contrast, if uncertainty is dominated by aleatoric variability inherent to the system, active sensing will not reduce it substantially, and the objective will naturally reflect that by showing limited information gain. This distinction is operationally important because it prevents overinvestment in sensing when uncertainty is irreducible.

The active sensing layer is also relevant to data collection for offline training. The same information gain objective can be used to design experiments in flow loops or to schedule data collection in the field to cover regions of state space where closure uncertainty is highest [33]. This closes a loop between model development and deployment: the twin identifies where it is uncertain and suggests data that would most reduce that uncertainty.

Computationally, the active sensing problem must be solved quickly. We use a short horizon and a limited action parameterization, such as a small number of excitation pulses or measurement configuration switches. Because the objective is noisy due to Monte Carlo estimation, we use a stochastic search strategy that evaluates a small set of candidate action templates and selects the best under the estimated objective. The entropy-stable solver and constrained closure sampling ensure that candidate rollouts remain stable, making such search feasible in real time.

The next section evaluates the full architecture under scenarios designed to stress stability, calibration, and the usefulness of active sensing [34].

6. Computational Studies Under Telemetry Stress and Geometry Shift

The evaluation focuses on properties that matter for decision support: stability of rollout, calibration of uncertainty, and the ability to reduce uncertainty through sensing and excitation. The studies are conducted on representative conduit configurations that include vertical and horizontal segments, compressibility effects, and actuator-driven boundary changes. The emphasis is on comparing the proposed entropy-stable distributional-closure twin to deterministic closure surrogates and to unconstrained probabilistic surrogates that do not enforce admissibility. The objective is not to claim universal superiority, but to illustrate how the proposed constraints affect failure modes that are otherwise common in deployment.

The first study examines long-horizon rollout stability under online closure updates. A deterministic surrogate closure is trained to predict closure coefficients as a function of state and geometry [35]. When rolled out for long horizons under changing actuation, the deterministic surrogate occasionally produces negative effective dissipation or extreme slip in regions of state space not well covered by training. These outputs can cause the solver to violate physical bounds on holdup or to develop oscillations that are not driven by inputs. By contrast, the constrained normalizing-flow closure always outputs admissible coefficients, and the entropy-stable solver maintains a discrete entropy inequality, preventing such instability. The resulting trajectories remain within physical bounds for significantly longer horizons, and when uncertainty grows due to weak observability, the twin reflects it through broader predictive distributions rather than through numerical blow-up.

The second study addresses calibration under sensor dropout. Pressure sensors are removed intermittently to mimic telemetry loss [36]. A deterministic twin continues to output point predictions with artificially low uncertainty, which can mislead a controller. An unconstrained probabilistic twin increases uncertainty but can drift into nonphysical states when closures sampled during dropout lead to negative dissipation. The proposed constrained twin increases uncertainty during dropout, maintains physical admissibility, and yields predictive intervals whose empirical coverage remains closer to nominal. In particular, upper quantiles of pressure predictions maintain conservative behavior during dropout windows, which is relevant for safety constraints. The calibration benefit is most pronounced when the dropout coincides with transients where closure uncertainty interacts strongly with compressibility.

The third study targets geometry shift [37]. The twin is trained on a set of diameter and roughness descriptors and then tested on a configuration with different roughness distribution and a mild inclination change. Deterministic closures interpolate poorly and can underpredict friction in some segments, leading to biased pressure predictions. The constrained probabilistic closure widens its distribution when geometry is out of distribution, reflecting epistemic uncertainty. This widening is operationally useful because it can trigger conservative control or prompt additional sensing. Active sensing is then applied by introducing small excitation maneuvers in boundary conditions within safe limits and by temporarily increasing measurement rate at a subset of sensors. The mutual-information objective selects maneuvers that maximize sensitivity of observations to closure multipliers, leading to a measurable reduction in posterior variance of friction multipliers [38]. The posterior narrowing translates into tighter pressure prediction bands and improved control confidence.

The fourth study examines the interaction of closure multimodality with transient behavior. In certain operating windows, two distinct closure regimes can explain the same pressure measurements, corresponding to different plausible holdup distributions. A Gaussian closure model cannot represent this and collapses to a mean that may be physically implausible. The normalizing-flow closure represents multimodality in the unconstrained latent space, leading to a mixture-like behavior in C. The twin therefore produces bimodal predictive distributions for certain internal variables while maintaining unimodal predictions for measured pressures [39]. Active sensing resolves this ambiguity by choosing excitation that produces distinct pressure signatures for the competing holdup hypotheses, leading to posterior collapse toward the correct mode. The ability to represent and then resolve multimodality is an advantage of the distributional closure approach, and it is difficult to obtain from deterministic surrogates.

The fifth study considers robustness to sensor bias. A slowly drifting pressure sensor is introduced. Deterministic models misattribute drift to closure changes, leading to incorrect friction multipliers and poor extrapolation. The probabilistic twin with a bias state can attribute part of the discrepancy to sensor bias, maintaining more stable closure estimates [40]. The constrained closure distribution prevents the model from compensating for bias by reducing dissipation below nominal in a way that would violate entropy stability. The resulting posterior over closures remains physically plausible, and predictive uncertainty reflects the ambiguity between bias and state.

Across these studies, active sensing is evaluated not only by uncertainty reduction but by the value of information for decision-making. In a simplified safety-aware control loop that uses predictive quantiles to enforce pressure constraints, reduced closure uncertainty leads to fewer overly conservative actions and fewer unnecessary throttling events. The benefit is modest in steady conditions but becomes more pronounced during transients where uncertainty would otherwise force large safety margins.

The studies also illuminate limitations [41]. If measurements provide insufficient sensitivity to certain closure components, active sensing may not reduce uncertainty significantly. In such cases, the twin remains conservative, and the information objective naturally indicates limited value from excitation. Additionally, if the nominal closure is systematically biased in a direction that the constrained distribution cannot correct because it only allows multiplicative increases, then the twin may be conservative. This reflects a design choice favoring safety over aggressive correction. In practice, one can mitigate this by choosing nominal closures that are not systematically over-dissipative or by allowing bounded decreases that still preserve entropy stability through alternative parameterizations, though such extensions require additional care.

Finally, computational feasibility is considered [42]. The constrained normalizing-flow sampling and entropy-stable solver are designed to be lightweight enough for real-time execution with moderate ensemble sizes. Active sensing adds overhead, but it operates on short horizons and a limited candidate set. The overall architecture is therefore suitable for integration into supervisory decision support systems where cycle times are on the order of seconds rather than milliseconds.

7. Conclusion

This paper proposed an entropy-stable, multi-fidelity digital-twin framework for gas–liquid two-phase flow that prioritizes stability and calibrated uncertainty under sparse telemetry and distribution shift. The main contribution was a constrained normalizing-flow closure model that outputs admissible distributions over closure coefficients and that can represent multimodal uncertainty without explicit regime labels, while remaining compatible with a conservative solver that enforces a discrete entropy inequality. This combination provides a practical guardrail against unstable rollout when closures are uncertain or updated online, a failure mode that is particularly damaging when the twin is used inside optimization, control, or sensing policies [43].

A second contribution was an information-theoretic active sensing formulation that selects measurement configurations and safe excitation actions to maximize expected information gain about latent holdup and closure deviations. By working directly with the twin's predictive distributions, the sensing policy distinguishes reducible epistemic uncertainty from irreducible variability and allocates sensing effort accordingly. Computational studies under sensor dropout, sensor bias, geometry shift, and ambiguity-induced multimodality indicated improved calibration and reduced instability relative to unconstrained closure surrogates, and they illustrated how active sensing can narrow closure uncertainty and improve the reliability of risk quantiles used for safety-aware decision-making.

The framework is intended as a stability-first foundation rather than a final operational product. Extensions can incorporate richer thermodynamic effects, compositional variation, and more detailed coupling to actuators and boundary equipment, provided the entropy-stable and admissibility constraints are preserved. More broadly, the results support the view that two-phase digital twins benefit from treating closures as constrained random variables and from coupling inference with deliberate information gathering, especially in settings where sparse telemetry and latent structure dominate operational uncertainty [44].

References

[1] M Nagib, A Burns, A Hassanpour, and Ahmed H. El-Banbi. "Application of Streamline Simulation to Gas Displacement Processes". In: *The International Journal of Multiphysics* 11.4 (Dec. 31, 2017), pp. 327–348. DOI: 10.21152/1750-9548.11.4.327.

[2] Jurgen Sprengel, Pedro Milano, Ryan Sfand, and Doug Kolak. "Modelling of High Point Vents in water gathering systems – a new approach using Simcenter Flomaster". In: *The APPEA Journal* 62.1 (May 13, 2022), pp. 106–115. DOI: 10.1071/aj21135.

[3] Cristian Marchioli and Stéphane Vincent. "Special issue on finite-size particles, drops and bubbles in fluid flows: advances in modelling and simulations". In: *Acta Mechanica* 230.2 (Jan. 30, 2019), pp. 381–386. DOI: 10.1007/s00707-018-2352-7.

[4] Kaushik Manikonda, Abu Rashid Hasan, Chinemerem Edmond Obi, Raka Islam, Ahmad Khalaf Sleiti, Motasem Wadi Abdelrazeq, and Mohammad Azizur Rahman. "Application of machine learning classification algorithms for two-phase gas-liquid flow regime

identification". In: *Abu Dhabi International Petroleum Exhibition and Conference*. SPE. 2021, D041S121R004.

[5] M. Ben Clennell et al. "Improved method for complete gas-brine imbibition relative permeability curves". In: *E3S Web of Conferences* 146 (Feb. 5, 2020), pp. 03003–. DOI: 10.1051/e3sconf/202014603003.

[6] Tong Liu, Shiming Zhang, and Moran Wang. "Does Rheology of Bingham Fluid Influence Upscaling of Flow through Tight Porous Media". In: *Energies* 14.3 (Jan. 28, 2021), pp. 680–. DOI: 10.3390/en14030680.

[7] Alessandro Vona, Guido Giordano, A. A. De Benedetti, Raffaele D'Ambrosio, Claudia Romano, and Michael Manga. "Ascent velocity and dynamics of the Fiumicino mud eruption, Rome, Italy". In: *Geophysical Research Letters* 42.15 (Aug. 4, 2015), pp. 6244–6252. DOI: 10.1002/2015gl064571.

[8] Concepción Paz, E. Suárez, M. Concheiro, and Jacobo Porteiro. "CFD Transient Simulation of a Breathing Cycle in an Oral-Nasal Extrathoracic Model". In: *Journal of Applied Fluid Mechanics* 10.3 (May 1, 2017), pp. 777–784. DOI: 10.18869/acadpub.jafm.73.240.25348.

[9] Daniel Bennett, Robert J. Cuss, Philip J. Vardon, Jon F. Harrington, and Hywel Rhys Thomas. "Phenomena exposure from the large scale gas injection test (Lasgit) dataset using a bespoke data analysis toolkit". In: *Geological Society, London, Special Publications* 400.1 (Mar. 5, 2014), pp. 497–505. DOI: 10.1144/sp400.5.

[10] Malte Mallach, Martin Gevers, Patrik Gebhardt, and Thomas Musch. "Fast and Precise Soft-Field Electromagnetic Tomography Systems for Multiphase Flow Imaging". In: *Energies* 11.5 (May 9, 2018), pp. 1199–. DOI: 10.3390/en11051199.

[11] PM Pieter Couwenberg, Q Qi Chen, and Gbmm Guy Marin. "Kinetics of a Gas-Phase Chain Reaction Catalyzed by a Solid: The Oxidative Coupling of Methane over Li/MgO-Based Catalysts". In: *Industrial & Engineering Chemistry Research* 35.11 (Jan. 1, 1996), pp. 3999–4011. DOI: 10.1021/ie9504617.

[12] Hamed Ameryoun, Franck Schoefs, Laurent Barillé, and Yoann Thomas. "Stochastic Modeling of Forces on Jacket-Type Offshore Structures Colonized by Marine Growth". In: *Journal of Marine Science and Engineering* 7.5 (May 22, 2019), pp. 158–. DOI: 10.3390/jmse7050158.

[13] Erkin Urinboevich Madaliev, Murodil Erkinjon Ugli Madaliev, Kamol Adilov, and Tohir Pulatov. "Comparison of turbulence models for two-phase flow in a centrifugal separator". In: *E3S Web of Conferences* 264 (June 2, 2021), pp. 01009–. DOI: 10.1051/e3sconf/202126401009.

[14] Wilson F. N. Santos. "Aerothermodynamic Analysis of a Reentry Brazilian Satellite". In: *Brazilian Journal of Physics* 42.5 (Oct. 6, 2012), pp. 373–390. DOI: 10.1007/s13538-012-0100-3.

[15] William C. Keene, Jennie L. Moody, James N. Galloway, Joseph M. Prospero, Owen R. Cooper, Sabine Eckhardt, and John R. Maben. "Long-term trends in aerosol and precipitation composition over the western North Atlantic Ocean at Bermuda". In: *Atmospheric Chemistry and Physics* 14.15 (Aug. 13, 2014), pp. 8119–8135. DOI: 10.5194/acp-14-8119-2014.

[16] Steffen Freitag, Antony D. Clarke, Steven G. Howell, Vladimir N. Kapustin, Teresa Campos, V. Brekhovskikh, and J. Zhou. "Combining airborne gas and aerosol measurements with HYSPLIT: a visualization tool for simultaneous evaluation of air mass history and back trajectory consistency". In: *Atmospheric Measurement Techniques* 7.1 (Jan. 14, 2014), pp. 107–128. DOI: 10.5194/amt-7-107-2014.

[17] Ying Xu and Shankar Subramaniam. "A multiscale model for dilute turbulent gas-particle flows based on the equilibration of energy concept". In: *Physics of Fluids* 18.3 (Mar. 1, 2006), pp. 033301–. DOI: 10.1063/1.2180289.

[18] Rjk Exley, Graham K. Westbrook, RR Haacke, and Sheila Peacock. "Detection of seismic anisotropy using ocean bottom seismometers: a case study from the northern head-wall of the Storegga Slide". In: *Geophysical Journal International* 183.1 (Aug. 27, 2010), pp. 188–210. DOI: 10.1111/j.1365-246x.2010.04730.x.

[19] Kaushik Manikonda, Raka Islam, Chinemerem Edmond Obi, Abu Rashid Hasan, Ahmad Khalaf Sleiti, Motasem Wadi Abdelrazeq, Ibrahim Galal Hassan, and Mohammad Azizur Rahman. "Horizontal two-phase flow regime identification with machine learning classification models". In: *International Petroleum Technology Conference*. IPTC. 2022, D011S021R002.

[20] C. O. Iloejesi and Lauren E. Beckingham. "Influence of Storage Period on the Geochemical Evolution of a Compressed Energy Storage System". In: *Frontiers in Water* 3 (Aug. 26, 2021). DOI: 10.3389/frwa.2021.689404.

[21] Cao Wei, Shiqing Cheng, Gang Chen, Wenyang Shi, Jiaxin Wu, Yang Wang, and Haiyang Yu. "Parameters evaluation of fault-karst carbonate reservoirs with vertical beads-on-string structure based on bottom-hole pressure: Case studies in Shunbei Oilfield, Tarim Basin of Northwestern China". In: *Oil & Gas Science and Technology – Revue d'IFP Energies nouvelles* 76 (Sept. 17, 2021), pp. 59–. DOI: 10.2516/ogst/2021037.

[22] Walter A. Illman and Daniel M. Tartakovsky. "Asymptotic analysis of three-dimensional pressure interference tests: A point source solution". In: *Water Resources Research* 41.1 (Jan. 11, 2005), pp. 01002–. DOI: 10.1029/2004wr003431.

[23] Sören Bernauer, Mathias Schöpf, Philipp Eibl, Christian Witz, Johannes Khinast, and Timo Hardiman. "Characterization of the gas dispersion behavior of multiple impeller stages by flow regime analysis and CFD simulations." In: *Biotechnology and bioengineering* 118.8 (June 6, 2021), pp. 3058–3068. DOI: 10.1002/bit.27815.

[24] Vinay V. Mahajan, JT Johan Padding, Tim M.J. Nijssen, Kay A. Buist, and J.A.M. Kuipers. "Nonspherical particles in a pseudo-2D fluidized bed : Experimental study". In: *AIChE journal. American Institute of Chemical Engineers* 64.5 (Feb. 6, 2018), pp. 1573–1590. DOI: 10.1002/aic.16078.

[25] Santiago G. Solazzi, Luis Guarracino, J. Germán Rubino, Tobias M. Müller, and Klaus Holliger. "Modeling Forced Imbibition Processes and the Associated Seismic Attenuation in Heterogeneous Porous Rocks". In: *Journal of Geophysical Research: Solid Earth* 122.11 (Nov. 3, 2017), pp. 9031–9049. DOI: 10.1002/2017jb014636.

[26] Sor Heoh Saw, V. Damideh, Perk Lin Chong, Paul Lee, Rajdeep Singh Rawat, and Sing Lee. "A 160 kJ dual plasma focus (DuPF) for fusion-relevant materials testing and nano-materials fabrication". In: *International Journal of Modern Physics: Conference Series* 32.2014 (Aug. 12, 2014), pp. 1460322–. DOI: 10.1142/s2010194514603226.

[27] Rob Westaway. "Repurposing of disused shale gas wells for subsurface heat storage: preliminary analysis concerning UK issues". In: *Quarterly Journal of Engineering Geology and Hydrogeology* 49.3 (Aug. 18, 2016), pp. 213–227. DOI: 10.1144/qjegh2016-016.

[28] Salam Al-Rbeawi. "Rate and pressure behavior considering the fractal characteristics of structurally disordered fractured reservoirs". In: *Journal of Petroleum Exploration and Production Technology* 11.1 (Oct. 6, 2020), pp. 483–507. DOI: 10.1007/s13202-020-00999-x.

[29] Audrey Michaud-Dubuy, Guillaume Carazzo, and Edouard Kaminski. "Wind Entrainment in Jets with Reversing Buoyancy: Implications for Volcanic Plumes". In: *Journal of Geophysical Research: Solid Earth* 125.10 (Sept. 26, 2020). DOI: 10.1029/2020jb020136.

[30] Dingwu Jiang, Meiliang Mao, Jin Li, and Xiaogang Deng. "An implicit parallel UGKS solver for flows covering various regimes". In: *Advances in Aerodynamics* 1.1 (Mar. 13, 2019), pp. 1–24. DOI: 10.1186/s42774-019-0008-5.

[31] Xiaoan Mao, Zhibin Yu, Artur J. Jaworski, and David Marx. "PIV studies of coherent structures generated at the end of a stack of parallel plates in a standing wave acoustic field". In: *Experiments in Fluids* 45.5 (Apr. 25, 2008), pp. 833–846. DOI: 10.1007/s00348-008-0503-7.

[32] Pascale Aussillous, Yixian Zhou, Pierre Ruyer, and Pierre-Yves Lagrée. "Discharge flow of granular media from silos with a lateral orifice and injection of air". In: *EPJ Web of Conferences* 140 (June 30, 2017), pp. 03052–. DOI: 10.1051/epjconf/201714003052.

[33] Işık Sena Akgün and Can Erkey. "Investigation of Hydrodynamic Behavior of Alginate Aerogel Particles in a Laboratory Scale Wurster Fluidized Bed." In: *Molecules (Basel, Switzerland)* 24.16 (Aug. 11, 2019), pp. 2915–. DOI: 10.3390/molecules24162915.

[34] Xiaoling Zhang, Lizhi Xiao, Xiaowen Shan, and Long Guo. "Lattice Boltzmann Simulation of Shale Gas Transport in Organic Nano-Pores". In: *Scientific reports* 4.1 (May 2, 2014), pp. 4843–4843. DOI: 10.1038/srep04843.

[35] Rajagopal Raghavan and Chih-Cheng Chen. "A study in fractional diffusion: Fractured rocks produced through horizontal wells with multiple, hydraulic fractures". In: *Oil & Gas Science and Technology – Revue d'IFP Energies nouvelles* 75 (Oct. 9, 2020), pp. 68–. DOI: 10.2516/ogst/2020062.

[36] Ivan Cornejo. "A Model for Correcting the Pressure Drop between Two Monoliths". In: *Catalysts* 11.11 (Oct. 29, 2021), pp. 1314–. DOI: 10.3390/catal11111314.

[37] Peng Wang, Lianhua Zhu, Wei Su, Lei Wu, and Yonghao Zhang. "Nonlinear oscillatory rarefied gas flow inside a rectangular cavity." In: *Physical review. E* 97.4 (Apr. 5, 2018), pp. 043103–. DOI: 10.1103/physreve.97.043103.

[38] Anna Russian, Philippe Gouze, Marco Dentz, and Alain C. Gringarten. "Multi-continuum Approach to Modelling Shale Gas Extraction". In: *Transport in Porous Media* 109.1 (May 21, 2015), pp. 109–130. DOI: 10.1007/s11242-015-0504-y.

[39] Sanjay B. Pawar. "Process Engineering Aspects of Vertical Column Photobioreactors for Mass Production of Microalgae". In: *ChemBioEng Reviews* 3.3 (June 7, 2016), pp. 101–115. DOI: 10.1002/cben.201600003.

[40] J. Elin Vesper, Theo J. M. Broeders, Joëlle Batenburg, Daniel E. A. van Odyck, and Chris R. Kleijn. "The interaction of parallel and inclined planar rarefied sonic plumes—From free molecular to continuum regime". In: *Physics of Fluids* 33.8 (Aug. 1, 2021), pp. 086103–. DOI: 10.1063/5.0056730.

[41] Rasool Erfani, Hossein Zare-Behtash, and Konstantinos Kontis. "Influence of shock wave propagation on dielectric barrier discharge plasma actuator performance". In: *Journal of Physics D: Applied Physics* 45.22 (May 16, 2012), pp. 225201–. DOI: 10.1088/0022-3727/45/22/225201.

[42] Jee Min Park, Dae Yun Kim, Jong Dae Baek, Yong-Jin Yoon, Pei-Chen Su, and Seong Hyuk Lee. "Effect of Electrolyte Thickness on Electrochemical Reactions and Thermo-Fluidic Characteristics inside a SOFC Unit Cell". In: *Energies* 11.3 (Feb. 25, 2018), pp. 473–. DOI: 10.3390/en11030473.

[43] Martin Henke, Thomas Monz, and Manfred Aigner. "Introduction of a New Numerical Simulation Tool to Analyze Micro Gas Turbine Cycle Dynamics". In: *Journal of Engineering for Gas Turbines and Power* 139.4 (Nov. 2, 2016). DOI: 10.1115/1.4034703.

[44] Shiqiang Zhang, A Ana Sobota, van Em Eddie Veldhuizen, and Peter Bruggeman. "Gas flow characteristics of a time modulated APPJ: The effect of gas heating on flow dynamics". In: *Journal of Physics D: Applied Physics* 48.1 (Dec. 8, 2014), pp. 015203–. DOI: 10.1088/0022-3727/48/1/015203.