Monteinstitute : Pages:16–37

Original Research

Latent State Space Learning with Physics Constraints for Gas Liquid Hydraulics and Safety Critical Diagnostics

Ahsan Raza Qadri[1], Bilal Saeed Khan[2] and Fahad Mehmood Siddiqi[3]

[1]Karakoram Institute of Computing, Shahrah-e-Quaid, Skardu 16100, Gilgit-Baltistan, Pakistan.
[2]Karachi School of Advanced Technologies, Korangi Road, Defence View, Karachi 75500, Pakistan.
[3]Punjab College of Information Sciences, Sargodha Road, Satellite Town, Faisalabad 38000, Pakistan.

Abstract

Gas–liquid two-phase flow appears throughout energy and process systems, including wellbores, risers, and pipelines, where operating decisions depend on interpreting sparse and noisy measurements. In drilling and managed-pressure operations, the same surface signals that drive routine control can also contain early indications of regime transitions and hazardous transients, yet the mapping from measurements to downhole states is typically non-unique and closure-dependent. This paper proposes a physics-constrained latent state-space learning framework that targets real-time inference of distributed two-phase hydraulics while providing calibrated uncertainty for safety-relevant decisions. The approach embeds a reduced-order two-fluid/drift-flux backbone inside a differentiable probabilistic state-space model whose latent variables represent coupled hydraulic states and regime-related structure without requiring a fixed, hand-crafted regime map. Training combines simulation-consistent objectives with weak supervision when available, enforcing conservation-consistent residuals and stability regularization across a wide envelope of flow rates, fluid properties, and geometries. Online inference is performed with a hybrid variational–sequential Monte Carlo scheme that yields posterior distributions over pressure, holdup, mixture velocity, and anomaly indicators under sensor latency and changing operating conditions. The resulting estimator supports regime-agnostic diagnostics, change-point detection, and risk-aware control interfaces while remaining computationally compatible with edge deployment. Numerical studies spanning synthetic high-fidelity rollouts and loop-inspired scenarios demonstrate improved generalization under distribution shift and tighter calibration compared with purely discriminative baselines, enabling earlier detection of hazardous transients at controlled false-alarm rates.

1. Introduction

Two-phase gas–liquid flow is characterized by strong nonlinearities, discontinuous transitions in interfacial structure, and closure relations that can vary with geometry, inclination, and fluid properties [1]. In wellbore contexts, these complexities are amplified by the fact that actionable measurements are often concentrated at the surface, while the quantities of interest are distributed along the annulus or drillstring and evolve under changing boundary conditions. Practical workflows therefore rely on a mixture of mechanistic modeling, heuristic thresholds, and empirical regime charts, each of which can fail when the operating envelope deviates from the assumptions that originally justified them. The central technical challenge is not merely to classify a flow regime from curated features, but to infer a dynamically consistent state of a partially observed, distributed, multi-physics system, and to do so with quantified uncertainty that supports safety-critical decisions.

Recent applications of machine learning have shown that high-accuracy regime identification is possible when extensive labeled data are available for a specific configuration. For example, supervised classifiers tuned on curated loop and literature datasets can achieve high test accuracy for vertical regime mapping under particular feature definitions and sampling strategies, as illustrated by the reported

98% test accuracy for a cross-validated K-nearest-neighbors classifier with a specific neighborhood size in a vertical gas–liquid setting by Manikonda et al.(2021) [2]. Such results are valuable signals that measurable patterns exist in commonly collected variables, but they also highlight an important limitation: a static classifier trained to map a point in feature space to a discrete label does not, by itself, ensure dynamic consistency, conservation compliance, or calibrated confidence under distribution shift. In operational settings, the relevant question is often counterfactual and temporal, such as whether a small deviation in surface pressure is consistent with benign regime evolution or indicates an emerging influx, washout, or choke malfunction.

This paper advances an alternative viewpoint: two-phase regime structure should be treated as latent, dynamically evolving information that mediates between mechanistic conservation laws and noisy measurements, rather than as a hard label predicted independently at each time. The objective is to estimate a posterior distribution over physically meaningful states, including mixture momentum and holdup, while representing regime-dependent closures and unmodeled effects through a learned latent correction that remains constrained by physics. This viewpoint leads naturally to probabilistic state-space modeling with embedded reduced-order hydraulics and explicit uncertainty quantification [3]. The technical contribution is a physics-constrained latent state-space learning formulation that unifies estimation, regime-agnostic diagnostics, and anomaly detection using a single trainable inference engine.

The proposed method departs from two common extremes. Purely mechanistic estimators can be brittle when the closures for slip, friction, or compressibility are mis-specified, and they may provide overconfident outputs because parameter uncertainty and model-form error are not represented explicitly. Purely data-driven estimators can achieve strong interpolative performance within the training distribution, yet may violate conservation, drift into unstable predictions under feedback control, and provide unreliable confidence under unseen combinations of geometry, inclination, and fluids. The present approach targets the middle ground by structuring learning around conservation-consistent residuals, weakly informative priors derived from reduced-order modeling, and a posterior inference mechanism designed for online use [4]. The result is not a regime classifier as an end product, but an estimator whose outputs can be projected onto regime-like summaries when needed, while retaining the ability to reason about continuous states and their uncertainties.

A second motivation arises from observability. Surface and near-surface measurements, such as standpipe pressure, choke pressure, flow-in and flow-out estimates, and pit volume proxies, are indirect observations of downhole multiphase behavior. The inverse problem is ill-posed because multiple downhole state trajectories can produce similar surface signatures, especially when gas compressibility and slip introduce delayed and attenuated responses. A useful estimator must therefore incorporate prior structure and temporal coherence to avoid spurious interpretations [5]. The proposed latent state-space model formalizes this requirement by coupling latent regime structure to physically motivated dynamics and by explicitly maintaining a distribution over states rather than a single best guess. This enables principled decision thresholds that can be tuned to acceptable false-alarm rates, which is critical for automation in drilling and production control contexts.

The remainder of the paper develops a modeling and inference framework that is intentionally broader than any single application. While drilling safety and kick diagnosis are motivating examples, the formulation targets generic gas–liquid hydraulics under partial observability. The presentation begins with a reduced-order representation of two-phase conservation laws and a discussion of what is observable from typical instrumentation [6]. It then introduces a physics-constrained latent state-space model that represents regime-dependent effects through latent variables and neural corrections, trained with a composite objective that penalizes conservation residuals and instability. The paper subsequently develops an online inference scheme combining variational filtering with sequential Monte Carlo to obtain calibrated posteriors in real time. Finally, it describes an evaluation protocol designed to stress-test generalization across operating envelopes, and it discusses limitations and open directions for integrating such estimators into safety-critical control loops.

2. Governing Equations, Reduced-Order Hydraulics, and Observability

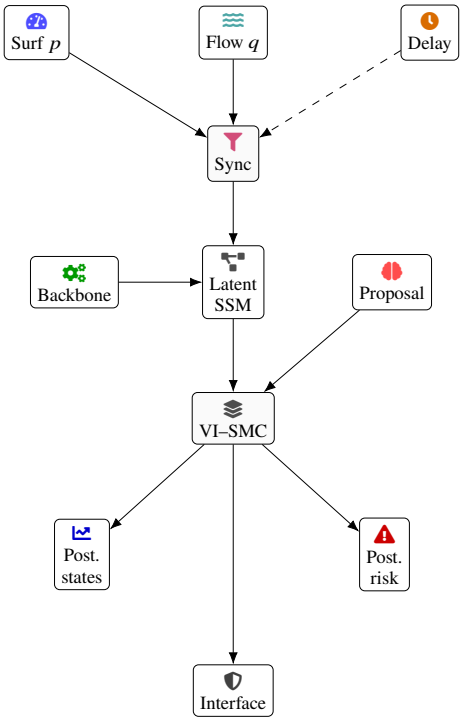

Figure 1: Compact real-time inference pipeline: synchronized surface signals drive a physics-constrained latent state-space estimator with a lightweight hybrid VI–SMC posterior, producing distributions over key hydraulic states and safety-relevant risk indicators for deployment in low-latency decision support.

Table 1: Key variables and notation used in the reduced-order hydraulics description.

Symbol	Type	Meaning
$x \in [0, L]$	coordinate	Axial position along conduit length L
t	time	Time index
$A(x)$, $D_h(x)$	geometry	Cross-sectional area and hydraulic diameter
$\theta(x)$	geometry	Inclination relative to horizontal
$\alpha(x,t)$	state	Gas void fraction (liquid holdup = $1-\alpha$)
u_g, u_ℓ, u_m	state	Phase velocities; $u_m = \alpha u_g + (1-\alpha)u_\ell$

Consider a one-dimensional representation of a conduit of length L parameterized by axial coordinate $x \in [0, L]$ with time $t \geq 0$. Let the cross-sectional area be $A(x)$ and the hydraulic diameter be $D_h(x)$, with a possibly varying inclination angle $\theta(x)$ measured relative to horizontal [7]. The two phases are gas and liquid with densities $\rho_g(p,T)$ and $\rho_\ell(p,T)$, and dynamic viscosities $\mu_g(p,T)$ and $\mu_\ell(p,T)$, where gas compressibility introduces strong pressure dependence. Let $\alpha(x,t)$ denote the gas void fraction, so that the liquid holdup is $1-\alpha$. Let $u_g(x,t)$ and $u_\ell(x,t)$ be superficial velocities or phase velocities depending on the chosen closure convention, and define mixture velocity $u_m = \alpha u_g + (1-\alpha)u_\ell$. Many mechanistic models begin from a two-fluid formulation, in which each phase satisfies a mass conservation law and a momentum balance with interfacial coupling, wall friction, and gravity.

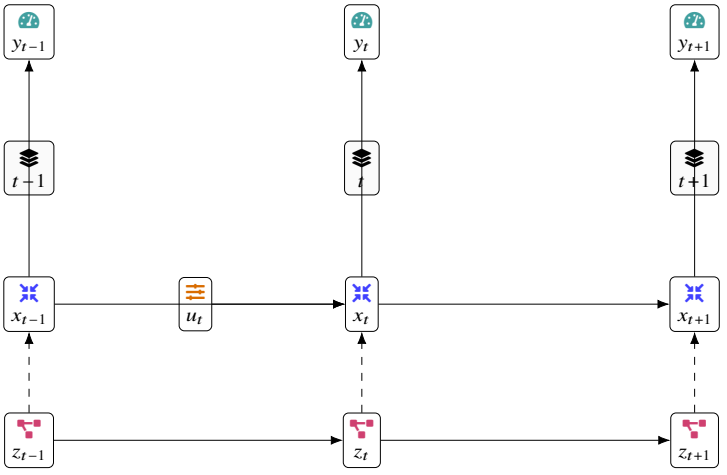

Figure 2: Minimal factor-graph view of the latent dynamical model: reduced hydraulic states x_t propagate forward in time, latent variables z_t modulate regime-dependent effects continuously, and indirect measurements y_t provide evidence for filtering over (x_t, z_t) under partial observability.

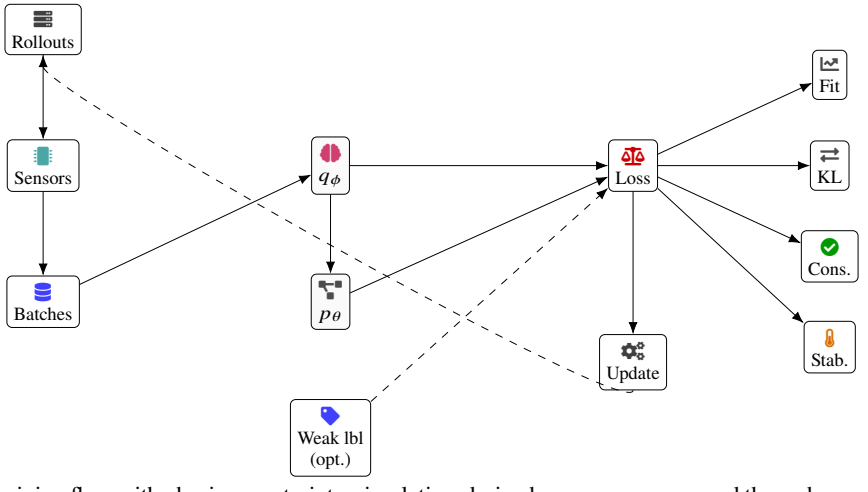

Figure 3: Training flow with physics constraints: simulation-derived sequences are passed through a sensor realism layer, then used to learn the generative dynamics and amortized posterior via a composite loss that couples measurement fit and latent regularization with conservation and stability penalties (with optional weak labels for interpretability only).

A general two-fluid mass conservation representation can be written as [8]

$$\frac{\partial}{\partial t}(\alpha\rho_g) + \frac{\partial}{\partial x}(\alpha\rho_g u_g) = \Gamma_g, \qquad \frac{\partial}{\partial t}((1-\alpha)\rho_\ell) + \frac{\partial}{\partial x}((1-\alpha)\rho_\ell u_\ell) = \Gamma_\ell, \tag{2.1}$$

where Γ_g and Γ_ℓ represent sources such as phase change, dissolution, or influx. For many wellbore hydraulics regimes over drilling-relevant time scales, phase change can be neglected relative to transport, while influx terms may be localized and episodic. The momentum balances may be expressed in a

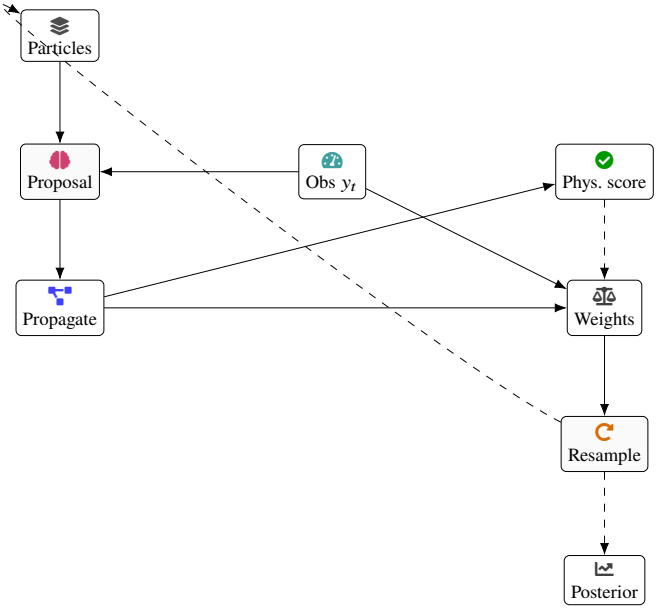

Figure 4: Online hybrid filtering: an amortized proposal suggests likely continuous states, importance weighting corrects against the generative model likelihood, and a lightweight physics-consistency score can downweight trajectories that explain measurements only via implausible conservation violations.

Table 2: Core governing relations and representative closure components.

Component	Form	Notes
Gas mass	$\partial_t(\alpha\rho_g) + \partial_x(\alpha\rho_g u_g) = \Gamma_g$	Γ_g covers influx/phase effects (often localized/episodic)
Liquid mass	$\partial_t((1-\alpha)\rho_\ell) + \partial_x((1-\alpha)\rho_\ell u_\ell) = \Gamma_\ell$	Often $\Gamma_\ell \approx 0$ over short hydraulics horizons
Momentum (2-fluid)	Includes $\partial_x p$, wall shear, interfacial transfer	Terms $\tau_{w\ell}, \tau_{wg}, M_{ig}$ are regime/structure dependent
Gravity	$-\rho\, g \sin\theta$	Strongly affects inclined/stratified/intermittent behavior
Drift-flux slip	$u_g = C_0 u_m + V_{gj}$	C_0, V_{gj} vary with regime, fluids, and inclination

simplified form as

$$\frac{\partial}{\partial t}\left(\alpha\rho_g u_g\right) + \frac{\partial}{\partial x}\left(\alpha\rho_g u_g^2\right) + \alpha\frac{\partial p}{\partial x} = -\tau_{wg}\frac{P_w}{A} + M_{ig} - \alpha\rho_g g\sin\theta, \tag{2.2}$$

$$\frac{\partial}{\partial t}\left((1-\alpha)\rho_\ell u_\ell\right) + \frac{\partial}{\partial x}\left((1-\alpha)\rho_\ell u_\ell^2\right) + (1-\alpha)\frac{\partial p}{\partial x} = -\tau_{w\ell}\frac{P_w}{A} - M_{ig} - (1-\alpha)\rho_\ell g\sin\theta, \tag{2.3}$$

where $p(x,t)$ is pressure, g is gravity, P_w is wetted perimeter, τ_{wg} and $\tau_{w\ell}$ are wall shear stresses associated with each phase, and M_{ig} is the interfacial momentum transfer term, which typically depends on slip, interfacial area, and regime structure. Closure relations for wall shear and interfacial transfer are inherently regime-dependent because they depend on whether the interface is dispersed, stratified, slugging, or annular, and on whether turbulence is dominated by one phase or both.

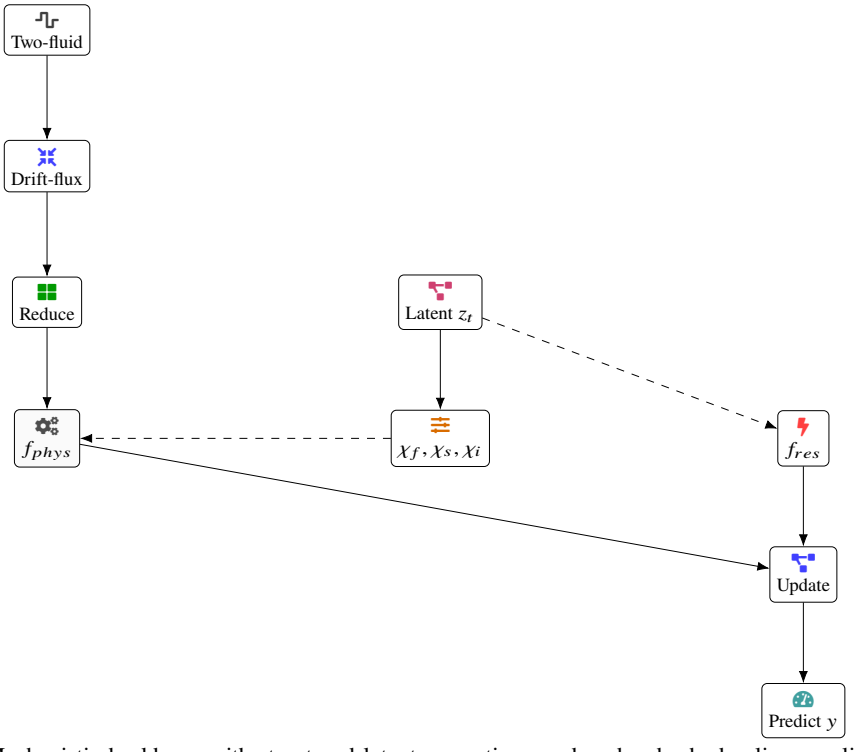

Figure 5: Mechanistic backbone with structured latent corrections: reduced-order hydraulics supplies the main dynamical scaffold, while latent variables modulate effective closures (e.g., slip/friction/interfacial transfer) and a residual pathway, enabling smooth adaptation across operating envelopes without hard regime switching.

Table 3: Reduced-order state-space backbone and observation structure.

Element	Expression	Interpretation
Reduced state	$x_t \in \mathbb{R}^d$	Modal coefficients / cell-averages + boundary proxies
Inputs	u_t	Pump/choke commands and other known actuation signals
Physics update	$x_{t+\Delta t} = f_{\text{phys}}(x_t, u_t; \psi) + w_t$	Mechanistic step with parameters/closures ψ
Sensors	$y_t = h_{\text{sens}}(x_t; \kappa) + v_t$	Sensor map with parameters κ and noise v_t
Ill-posed inverse	$y_t = \mathcal{H}(p, \alpha, u_m; \eta) + \varepsilon_t$	Multiple downhole trajectories can match similar surface signals

Although two-fluid models are expressive, they can be too complex for real-time inference when embedded directly into filtering, especially when closure discontinuities cause stiffness or non-smoothness [9]. Reduced-order models, such as drift-flux formulations, provide a tractable alternative by representing slip through a constitutive relation between the gas velocity and mixture velocity. A common drift-flux structure writes

$$u_g = C_0 u_m + V_{gj}, \tag{2.4}$$

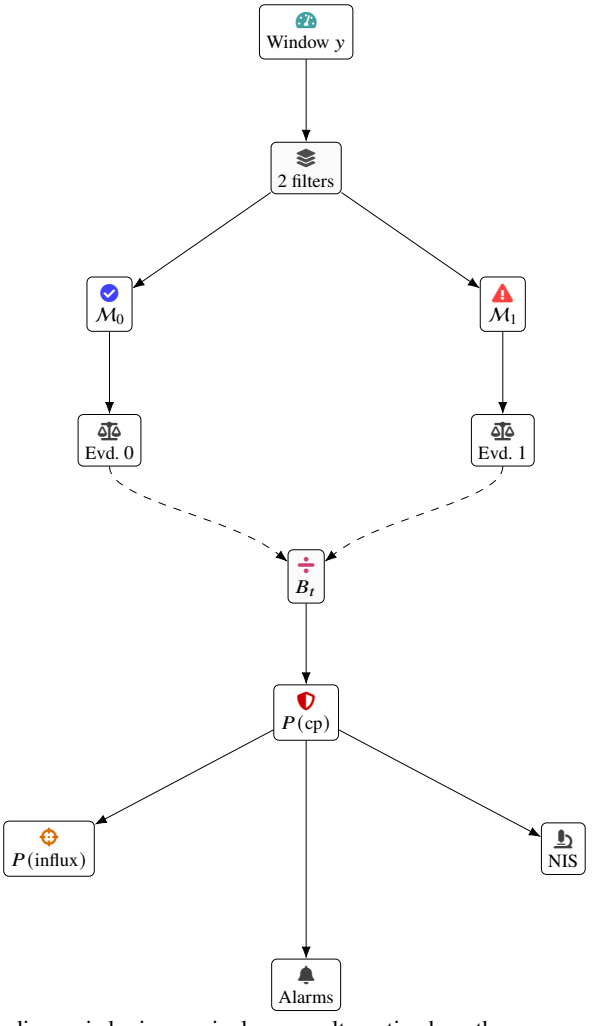

Figure 6: Posterior-driven diagnosis logic: nominal versus alternative hypotheses generate a Bayes-factor change statistic, which maps to probabilistic alarms; specialized posteriors (e.g., influx likelihood) and innovation-based sensor checks support disambiguation between hazardous transients and measurement faults.

where C_0 is a distribution parameter and V_{gj} is a drift velocity that captures buoyancy-driven slip. In one-dimensional form, coupling this relation with mixture mass conservation and momentum can yield a system in terms of p, α, and u_m that is easier to integrate numerically. Yet C_0 and V_{gj} are not constants across regimes or inclinations, and their dependence on flow rates, geometry, and fluid properties is a major source of model-form uncertainty.

For the inference problem, the key question is what aspects of this distributed state are identifiable from available measurements. Let a measurement vector $y(t)$ include surface pressure signals and flow estimates, such as standpipe pressure $p_{sp}(t)$, casing pressure $p_c(t)$, choke pressure $p_{ch}(t)$, pump rate proxy $q_{in}(t)$, and a flow-out estimate $q_{out}(t)$ with bias and latency. A simplified observation model can be expressed as [10]

$$y(t) = \mathcal{H}\big(p(\cdot,t), \alpha(\cdot,t), u_m(\cdot,t); \eta(t)\big) + \varepsilon(t), \tag{2.5}$$

where $\mathcal{H}$ is an observation operator that maps distributed states to sensor outputs, $\eta(t)$ captures unobserved operational parameters such as choke opening or pump efficiency, and $\varepsilon(t)$ is measurement noise.

Table 4: Latent state-space model components and probabilistic factorization.

Block	Definition	Role
Physical state	x_t	Reduced hydraulics states (pressure/holdup/velocity structure)
Latent state	z_t	Continuous regime/closure structure and unmodeled effects
Initial prior	$p(x_0)p(z_0)$	Start-up uncertainty and prior structure
Transition	$p(z_{t+1} \mid z_t, x_t, u_t)\, p(x_{t+1} \mid x_t, z_{t+1}, u_t)$	Latent evolves and modulates physical evolution
Emission	$p(y_t \mid x_t, z_t)$	Measurement likelihood with nonlinearities and bias

Table 5: Structured residual correction and physically aligned latent modifiers.

Item	Form	Practical effect
Residual transition	$x_{t+1} = f_{\text{phys}}(x_t, u_t) + f_{\text{res}}(x_t, z_{t+1}, u_t) + \xi_t$	Learns model-form mismatch while retaining mechanistic backbone
Friction multiplier	$\Delta p_f^{\text{eff}} = \chi_f(z_t)\Delta p_f$	Adapts friction under roughness/cuttings/closure shifts; $\chi_f > 0$
Slip correction	$V_{gj}^{\text{eff}} = V_{gj}^{\text{nom}} + \chi_s(z_t)$	Smoothly captures regime/inclination-dependent slip changes
Interfacial scaling	$M_{ig}^{\text{eff}} = \chi_i(z_t)M_{ig}$	Modulates phase coupling intensity under changing structure
Sensor bias term	$y_t = h_{\text{phys}}(x_t) + h_{\text{bias}}(x_t, z_t) + \epsilon_t$	Explains offsets/sensor artifacts without breaking dynamics

Table 6: Physics-constrained learning objective and regularization components.

Term	Expression (sketch)	Purpose
Data fit	$-\mathbb{E}_{q_\phi} \sum_t \log p_\theta(y_t \mid x_t, z_t)$	Match predicted sensor statistics to observations/simulations
Dynamics agreement	$\mathbb{E}_{q_\phi} \sum_t \text{KL}(q_\phi \parallel p_\theta)$	Align inferred trajectories with the generative transition
Conservation penalty	$\lambda_{\text{con}}\mathbb{E} \sum_t \lVert\mathcal{R}_{\text{con}}(x_t, x_{t+1}, u_t)\rVert_2^2$	Encourage mass/momentum consistency (integral or reduced form)
Stability penalty	$\lambda_{\text{stab}}\mathbb{E} \sum_t \lVert\mathcal{R}_{\text{stab}}(x_t, z_t, u_t)\rVert_2^2$	Prevent unphysical states (e.g., $\alpha \notin [0, 1]$) and blow-up
Weak labels (opt.)	$\lambda_{\text{weak}}\mathbb{E}\big[-\sum_{t\in\mathcal{T}_{\text{lbl}}} \log p(r_t \mid z_t)\big]$	Optional interpretability alignment when labels exist

Even if $\mathcal{H}$ is known, the inverse mapping is not unique. Two different void-fraction profiles may yield similar surface pressures if their integrated hydrostatic and frictional contributions match. Similarly, an influx localized at depth may not be immediately distinguishable from changes in friction or cuttings loading at the surface. Observability improves when multiple sensors are available, when excitation occurs through controlled boundary changes, or when the model includes strong priors that restrict feasible state trajectories.

A practical estimator must therefore do more than minimize instantaneous error. It must maintain a physically plausible set of state trajectories consistent with conservation and with sensor dynamics

Table 7: Hybrid variational–sequential Monte Carlo filtering steps for online inference.

Step	Operation	Output
Proposal	$(x_t^{(m)}, z_t^{(m)}) \sim q_\phi(\cdot \mid x_{t-1}^{(m)}, z_{t-1}^{(m)}, y_{0:t}, u_{0:t-1})$	Observation-aware particle propagation
Weight update	$w_t^{(m)} \propto w_{t-1}^{(m)} \frac{p_\theta(y_t \mid x_t^{(m)}, z_t^{(m)}) p_\theta(x_t^{(m)}, z_t^{(m)} \mid \cdot)}{q_\phi(\cdot)}$	Importance correction for calibrated posterior
Physics weighting (opt.)	$w_t^{(m)} \leftarrow w_t^{(m)} \exp(\lambda_{\text{phys}} \ell_{\text{phys}}^{(m)}(t))$	Discourage conservation-violating trajectories over a window
Resampling	Trigger when effective sample size is low	Avoid particle degeneracy under transients
Rejuvenation (opt.)	Local MCMC/Langevin moves on (x_t, z_t)	Improve diversity and multimodal tracking

Table 8: Change-point and fault monitoring quantities used for safety-relevant decisions.

Mechanism	Statistic	Interpretation
Model comparison	$B_t = \frac{p(y_{t-W+1:t} \mid \mathcal{M}_1)}{p(y_{t-W+1:t} \mid \mathcal{M}_0)}$	Bayes factor for alternative (e.g., elevated latent volatility)
Alarm score	$\Pr(\text{change} \mid y) \approx \sigma(\log B_t + \log \text{prior odds})$	Thresholdable probability with tunable false-alarm tradeoff
Targeted event prob.	$\Pr(\dot{m}_{\text{influx}} > 0 \mid y_{0:t})$	Direct probability for influx-like anomalies when modeled
Innovation test	$\text{NIS}_t = (y_t - \mathbb{E}\tilde{y}_t)^\top \text{Cov}(\tilde{y}_t)^{-1} (y_t - \mathbb{E}\tilde{y}_t)$	Persistent high values suggest sensor fault/bias or mismatch
Bias separation	Latent emission bias $h_{\text{bias}}(x_t, z_t)$	Attribute discrepancies to sensors vs true hydraulics changes

[11]. This motivates the use of a state-space formulation in which the distributed state is projected onto a lower-dimensional representation that retains the dominant modes relevant to surface observables, while allowing learned corrections to capture regime-dependent effects. For example, a Galerkin-type reduction can represent pressure and holdup profiles through basis functions $\phi_k(x)$ and coefficients $s_k(t)$ as

$$p(x,t) \approx \sum_{k=1}^{K_p} s_k^{(p)}(t)\phi_k^{(p)}(x), \qquad \alpha(x,t) \approx \sum_{k=1}^{K_\alpha} s_k^{(\alpha)}(t)\phi_k^{(\alpha)}(x), \tag{2.6}$$

with analogous expansions for mixture velocity. Alternatively, a compartment model can discretize the conduit into N cells, evolving cell-averaged states with fluxes computed by a simplified Riemann solver. The choice of reduction is not unique, but it must balance computational constraints with representational fidelity, and it must permit stable filtering under uncertainty [12].

Horizontal and inclined flows present additional challenges because stratification and intermittent structures produce dynamics that are strongly influenced by inclination and by the interaction between gravity and inertia. In such settings, data-driven regime identification can still achieve strong accuracy when trained on representative datasets. For example, a cross-validated K-nearest-neighbors classifier was reported to reach 97.4% prediction accuracy for horizontal gas–liquid regime identification under a particular experimental and literature aggregation, with multiple major and minor regimes identified across superficial velocity space by Manikonda et al.(2022) [13]. From an inference perspective, this underscores that measurement features encode regime information, but it also implies

Table 9: Evaluation protocol elements, baselines, and performance axes.

Category	Item	What is assessed
Data generation	Simulator + sensor model	Ground truth with latency, missingness, bias, and noise
Stress tests	Geometry/closure shifts	Robustness under inclination, roughness, and closure perturbations
Baselines	EKF; RNN; regime+switching	Comparisons to mechanistic, black-box, and hard-switch hybrids
State accuracy	RMSE / profile error	Pressure/holdup accuracy at depths and integrated along conduit
Calibration	Coverage vs nominal	Empirical coverage of credible intervals (e.g., 90% BHP)
Detection	Delay vs false alarms	Early warning vs controlled false-alarm rates for injected events

Table 10: Deployment considerations and expected failure modes under partial observability.

Issue	Symptom	Mitigation within the framework
Identifiability limits	Multimodal/ambiguous posterior	Maintain particle hypotheses; act via risk thresholds and intervals
Backbone mismatch	Corrections entangle omitted physics	Enrich reduced state; constrain residuals as closure modifiers
Training realism gap	Overconfidence on field artifacts	Staged training: broad sim → self-supervised field → event calibration
Closed-loop shift	Estimator-control coupling artifacts	Train with diverse control policies; stability regularization
Compute/latency	Edge budget constraints	Low-dimensional x_t, amortized proposal, modest particle counts

that a reduced-order model used online must accommodate transitions between stratified and intermittent structures without requiring brittle, hand-coded switching logic. The present work addresses this by using latent variables to represent regime-dependent closure modifications continuously, enabling smooth interpolation across operating envelopes while maintaining conservation constraints [14].

The remainder of this section introduces a reduced-order hydraulics backbone that will be embedded within the learning model. Let the reduced state vector at time t be denoted $x_t \in \mathbb{R}^d$, collecting coefficients for pressure and holdup modes, along with boundary-related variables such as bottomhole pressure proxy and choke dynamics. A nominal mechanistic update may be represented as a time-discrete map

$$x_{t+\Delta t} = f_{\text{phys}}(x_t, u_t; \psi) + w_t, \tag{2.7}$$

where u_t represents known inputs such as pump rate or choke command, ψ denotes mechanistic parameters and closures, and w_t represents process noise capturing unresolved physics and parameter variability. The associated observation model is

$$y_t = h_{\text{sens}}(x_t; \kappa) + v_t, \tag{2.8}$$

with sensor parameters κ and observation noise v_t [15]. In traditional filtering, f_{phys} is fixed and w_t is tuned heuristically. The present work instead treats model-form mismatch as structured and learnable,

using a latent representation that modulates the mechanistic update in a way that is constrained by conservation residuals and stability.

3. Latent State-Space Model with Physics-Constrained Learning Objectives

The proposed framework models two-phase hydraulics as a probabilistic state-space system in which physically interpretable reduced states are augmented by latent variables capturing regime-dependent and unmodeled effects. Let x_t denote the reduced physical state, and let z_t denote a latent vector. The latent vector is not restricted to discrete regime labels; instead it is designed to encode continuous factors such as effective slip correction, friction multiplier deviations, interfacial coupling intensity, and other low-dimensional signatures that influence the evolution of x_t and the mapping to measurements. The joint generative model is defined by an initial distribution and transition and emission distributions: [16]

$$p(x_0, z_0) = p(x_0)\, p(z_0), \tag{3.1}$$

$$p(x_{t+1}, z_{t+1} \mid x_t, z_t, u_t) = p(z_{t+1} \mid z_t, x_t, u_t)\, p(x_{t+1} \mid x_t, z_{t+1}, u_t), \tag{3.2}$$

$$p(y_t \mid x_t, z_t) = p(y_t \mid x_t, z_t). \tag{3.3}$$

The design choice is to let z_{t+1} influence the physical update, enabling the latent variables to act as a structured correction that can evolve over time and respond to regime transitions or anomalies. The physical state transition is constructed as a composition of the nominal mechanistic step and a learned residual:

$$x_{t+1} = f_{\text{phys}}(x_t, u_t; \psi) + f_{\text{res}}(x_t, z_{t+1}, u_t; \theta) + \xi_t, \tag{3.4}$$

where f_{res} is a neural operator with parameters θ, and ξ_t is stochastic noise. The residual is constrained to respect conservation and stability through the training objective, and it can be regularized to remain small when the mechanistic model is adequate, thereby preventing the network from absorbing dynamics that are already well explained by physics.

A central difficulty in learning such models for two-phase flow is the scarcity of high-quality labels for downhole states and the domain shift between laboratory data and field operations [17]. The proposed approach therefore emphasizes self-supervision and physics-consistent training signals. Suppose one has access to simulated rollouts from a high-fidelity multiphase model or from a mechanistic simulator under varied conditions, potentially with injected parameter variability to approximate uncertainty. Let $\hat{x}_t$ denote simulated reduced states, and let $\hat{y}_t$ denote simulated sensor outputs after applying a sensor model and adding noise. Training can then use sequences $(\hat{y}_{0:T}, u_{0:T-1})$ as inputs and penalize discrepancies between predicted and simulated outputs while also enforcing physics residual constraints on reconstructed states.

Let $q_\phi(x_{0:T}, z_{0:T} \mid y_{0:T}, u_{0:T-1})$ denote an amortized variational posterior with parameters ϕ, implemented as a recurrent inference network or a transformer-style temporal encoder with causal masking for filtering. The learning objective is based on a negative evidence lower bound augmented with physics and stability penalties:

$$\mathcal{L}(\theta, \phi) = -\mathbb{E}_{q_\phi}\Big[\sum_{t=0}^{T} \log p_\theta(y_t \mid x_t, z_t)\Big] + \mathbb{E}_{q_\phi}\Big[\sum_{t=0}^{T-1} \mathrm{KL}\big(q_\phi(x_{t+1}, z_{t+1} \mid \cdot) \,\|\, p_\theta(x_{t+1}, z_{t+1} \mid x_t, z_t, u_t)\big)\Big] \tag{3.5}$$

$$+ \lambda_{\text{con}}\mathbb{E}_{q_\phi}\Big[\sum_{t=0}^{T-1} \big\|\mathcal{R}_{\text{con}}(x_t, x_{t+1}, u_t)\big\|_2^2\Big] + \lambda_{\text{stab}}\mathbb{E}_{q_\phi}\Big[\sum_{t=0}^{T-1} \big\|\mathcal{R}_{\text{stab}}(x_t, z_t, u_t)\big\|_2^2\Big]. \tag{3.6}$$

Here $\mathcal{R}_{\text{con}}$ represents discrete conservation residuals computed from the reconstructed reduced states, while $\mathcal{R}_{\text{stab}}$ penalizes behavior associated with numerical instability, such as unphysical negative densities, void fractions outside $[0, 1]$, or growth in energy-like quantities beyond what inputs can support. The multipliers λ_{con} and λ_{stab} balance statistical fit and physical plausibility. Importantly, the residual penalty is computed in the reduced-order space but can be designed to approximate conservation in integral form. For a compartment model with cell volumes V_i, a typical mass conservation residual for gas might be expressed as [18]

$$\mathcal{R}_{g,i}(t) = \frac{1}{\Delta t}\left(\alpha_i^{t+1}\rho_{g,i}^{t+1}V_i - \alpha_i^t\rho_{g,i}^t V_i\right) + \left(F^t_{g,i+\frac{1}{2}} - F^t_{g,i-\frac{1}{2}}\right) - S^t_{g,i}, \tag{3.7}$$

with analogous residuals for liquid mass and mixture momentum. Even when the estimator does not explicitly track cell states, similar residuals can be defined using reconstructed profiles from basis expansions and numerical quadrature.

Because regime structure affects closures, one might be tempted to train the latent variables to align with predefined regime labels. The present formulation avoids making such labels central. Instead, it allows optional weak supervision when available, mapping latent states to regime-like summaries without forcing the dynamics to be driven by a discrete classifier [19]. Let r_t denote an optional regime label available for a subset of training data, and let $g_\theta(z_t)$ be a soft classifier. A weak supervision term can be added:

$$\mathcal{L}_{\text{weak}} = \lambda_{\text{weak}}\mathbb{E}_{q_\phi}\Big[-\sum_{t\in\mathcal{T}_{\text{lbl}}} \log p_\theta(r_t \mid z_t)\Big], \tag{3.8}$$

where $\mathcal{T}_{\text{lbl}}$ indexes labeled time steps. This term is optional and, when used, it serves to interpret z_t rather than to define it. In many operational datasets, labels may be noisy or inconsistent across sources, so the method is designed to function without them.

A key design decision concerns the structure of z_t [20]. For two-phase hydraulics, it is beneficial to align latent components with physically meaningful modifiers. One can parameterize effective closure multipliers as smooth functions of z_t, such as an effective friction factor multiplier $\chi_f(z_t)$, a slip correction $\chi_s(z_t)$, and an interfacial transfer intensity $\chi_i(z_t)$. For example, if the nominal mechanistic update uses a friction model producing a pressure drop term Δp_f, the corrected term becomes $\chi_f(z_t)\Delta p_f$ with $\chi_f(z_t)$ constrained to remain positive via a softplus mapping. Similarly, drift velocity V_{gj} can be corrected as $V_{gj}^{\text{eff}} = V_{gj}^{\text{nom}} + \chi_s(z_t)$, with χ_s bounded to physically plausible ranges. Embedding corrections in this way reduces the risk that the residual network learns arbitrary dynamics unrelated to hydraulics, and it provides a pathway to interpretability and debugging.

The latent transition distribution $p(z_{t+1} \mid z_t, x_t, u_t)$ is constructed to allow both smooth evolution and abrupt shifts. A practical representation is a Gaussian with mean given by a gated update: [21]

$$z_{t+1} = z_t + \sigma\big(a_\theta(x_t, u_t)\big) \odot m_\theta(x_t, u_t, z_t) + \omega_t, \tag{3.9}$$

where $\sigma(\cdot)$ is a logistic function, $\odot$ denotes elementwise multiplication, a_θ is a gating network producing update magnitudes, m_θ produces direction of change, and ω_t is Gaussian noise. In benign steady operation, gates can remain small, so z_t changes slowly; during regime transitions or anomalies, gates open and the latent state adapts. This structure encourages temporal coherence without preventing change-points.

Finally, the emission model $p(y_t \mid x_t, z_t)$ accounts for sensor nonlinearities and unmodeled couplings. For instance, standpipe pressure may be a function of integrated friction, hydrostatic head, and pump characteristics [22]. Rather than forcing the mechanistic observation operator to be exact, the model can represent y_t as

$$y_t = h_{\text{phys}}(x_t; \kappa) + h_{\text{bias}}(x_t, z_t; \beta) + \epsilon_t, \tag{3.10}$$

where h_{phys} is a known mapping and h_{bias} is a learned bias term capturing systematic sensor offsets, unmodeled cuttings effects, or unmodeled choke dynamics. The physics penalties ensure that h_{bias} does not compensate for violations of conservation in implausible ways.

4. Physics-Constrained Inference, Filtering, and Uncertainty Quantification

Training a latent state-space model provides a generative description, but online use requires efficient inference of the posterior distribution over states and latent variables given streaming observations. The target is the filtering posterior $p(x_t, z_t \mid y_{0:t}, u_{0:t-1})$, which evolves as new measurements arrive. For safety-critical applications, the estimator must not only output point estimates but also maintain calibrated uncertainty that reflects sensor noise, parameter variability, and model-form error. The proposed inference strategy combines amortized variational filtering with particle-based correction, leveraging the strengths of both.

A purely amortized variational filter can be implemented by an inference network that maps recent observations and inputs to approximate posterior parameters, such as mean and covariance for (x_t, z_t) [23]. While fast, such filters can become miscalibrated when encountering sequences outside the training distribution. Particle filters, in contrast, can represent multimodal posteriors and adapt via likelihood weighting, but can be computationally expensive in high-dimensional continuous spaces and prone to degeneracy. The hybrid strategy uses amortized proposals to guide particles, thereby reducing variance and enabling real-time operation.

Let the approximate filtering distribution at time t be represented by M weighted particles $\{(x_t^{(m)}, z_t^{(m)}, w_t^{(m)})\}_{m=1}^{M}$. Given an observation y_t, each particle is propagated through the transition model using a proposal distribution that incorporates the inference network. Specifically, define a proposal $q_\phi(x_t, z_t \mid x_{t-1}, z_{t-1}, y_{0:t}, u_{0:t-1})$ whose parameters are output by the inference network conditioned on a summary of the observation history and the previous particle state. Propagation then draws [24]

$$(x_t^{(m)}, z_t^{(m)}) \sim q_\phi(\cdot \mid x_{t-1}^{(m)}, z_{t-1}^{(m)}, y_{0:t}, u_{0:t-1}), \tag{4.1}$$

and weights are updated by the standard importance ratio:

$$w_t^{(m)} \propto w_{t-1}^{(m)} \frac{p_\theta(y_t \mid x_t^{(m)}, z_t^{(m)})\, p_\theta(x_t^{(m)}, z_t^{(m)} \mid x_{t-1}^{(m)}, z_{t-1}^{(m)}, u_{t-1})}{q_\phi(x_t^{(m)}, z_t^{(m)} \mid x_{t-1}^{(m)}, z_{t-1}^{(m)}, y_{0:t}, u_{0:t-1})}. \tag{4.2}$$

Resampling is performed when the effective sample size falls below a threshold, and rejuvenation can be introduced via a short Markov chain move that targets the posterior locally, such as a few steps of Langevin dynamics on (x_t, z_t) using gradients of the log-likelihood and physics residual penalties.

The physics constraints affect inference in two ways. First, they shape the generative model p_θ so that the transition distribution respects conservation in expectation. Second, they can be incorporated directly into the online weighting by defining an augmented likelihood that penalizes conservation residuals for reconstructed trajectories [25]. For a short sliding window of length W, one can define a windowed physics consistency score

$$\ell_{\text{phys}}^{(m)}(t) = -\sum_{\tau=t-W}^{t-1} \left\| \mathcal{R}_{\text{con}}(x_\tau^{(m)}, x_{\tau+1}^{(m)}, u_\tau) \right\|_2^2, \tag{4.3}$$

and use $\exp(\lambda_{\text{phys}} \ell_{\text{phys}}^{(m)}(t))$ as a multiplicative factor in weights. This discourages particles that explain sensor data only by violating conservation. Care is required to avoid over-penalization when the residual arises from genuine unmodeled effects; this is mitigated by allowing z_t to adjust closures and by learning the scale of residual penalties during training.

Calibration of uncertainty is addressed by representing multiple sources. Observation noise v_t is modeled explicitly in $p(y_t \mid x_t, z_t)$, often with heteroscedastic covariance that depends on operating conditions, reflecting that some sensors become noisier under vibration or flow transients [26]. Process uncertainty is represented through stochastic transitions and through latent noise in z_t, allowing the model to represent variability in closures. Model-form uncertainty can be approximated through ensembles of models trained with different random initializations or bootstrapped datasets, producing a mixture posterior. In practice, an ensemble of E models yields a predictive distribution

$$p(y_{t+1} \mid y_{0:t}) \approx \frac{1}{E} \sum_{e=1}^{E} \int p_{\theta_e}(y_{t+1} \mid x_{t+1}, z_{t+1})\, p_{\theta_e}(x_{t+1}, z_{t+1} \mid y_{0:t})\, dx_{t+1}\, dz_{t+1}, \tag{4.4}$$

which can be approximated with particles. This ensemble-based approach tends to increase uncertainty under distribution shift, which is desirable for risk-aware decision-making [27].

A distinctive feature of two-phase hydraulics is that the posterior can be multimodal, especially when gas compressibility and slip induce delayed responses. For example, an observed standpipe pressure drop might be consistent with a reduction in friction due to reduced liquid viscosity, or with gas entry that reduces effective density and changes hydrostatic head. The latent state-space formulation can represent such ambiguity by maintaining multiple particle hypotheses with distinct latent states z_t corresponding to different closure adjustments. Over time, as additional measurements arrive, likelihood weighting and physics consistency can prune inconsistent hypotheses.

The framework also supports derived quantities of interest [28]. A regime-like summary can be obtained by applying a learned mapping $\pi(z_t)$ that outputs a simplex over interpretable categories, but this is optional and kept secondary. More central are continuous derived estimates such as bottomhole pressure proxy, gas fraction at key depths, or influx mass rate. If an influx source term is included in the state, say m_t representing gas mass in the annulus, its evolution can be modeled as

$$m_{t+1} = m_t + \Delta t\, \dot{m}_{\mathrm{influx}}(x_t, z_t, u_t) - \Delta t\, \dot{m}_{\mathrm{out}}(x_t, z_t, u_t) + \nu_t, \tag{4.5}$$

where the outflow term depends on transport and separation. Even when $\dot{m}_{\mathrm{influx}}$ is not directly observed, the estimator can infer it from pressure and flow discrepancies, with uncertainty reflecting identifiability limitations.

A practical concern is computational latency [29]. The hybrid method is designed so that the amortized proposal does most of the work, producing near-optimal particles when the model is within-distribution, while the particle correction protects against drift when signals become atypical. Because the reduced physical state dimension d is chosen to be modest, particle counts on the order of tens to low hundreds can be sufficient for real-time operation on edge hardware, especially when the mechanistic backbone is differentiable and can be vectorized across particles. Stability regularization during training further reduces the likelihood of numerical divergence under online filtering.

5. Online Change-Point Detection and Risk-Aware Decision Interfaces

In wellbore operations and other process settings, the estimator is typically embedded inside a decision loop where alarms, control adjustments, or automated responses are triggered based on inferred states. The relevant events are often changes in the generating process, such as a sudden gas influx, a transition into severe slugging, a loss of circulation, or a sensor fault [30]. The framework therefore includes an explicit change-point detection mechanism operating on the filtering posterior, designed to separate benign regime evolution from anomalous dynamics.

Change-point detection is framed as a Bayesian hypothesis test comparing a nominal model to an alternative that allows a transient deviation. Let $\mathcal{M}_0$ denote the nominal transition model and $\mathcal{M}_1$ denote a model with an additional disturbance term or increased latent volatility. A simple instantiation is to

increase the latent noise covariance for z_t under $\mathcal{M}_1$, allowing abrupt closure shifts. Given filtering particles, one can compute an approximate Bayes factor over a sliding window:

$$B_t = \frac{p(y_{t-W+1:t} \mid y_{0:t-W}, \mathcal{M}_1)}{p(y_{t-W+1:t} \mid y_{0:t-W}, \mathcal{M}_0)} \approx \frac{\sum_m w_{t,\mathcal{M}_1}^{(m)}}{\sum_m w_{t,\mathcal{M}_0}^{(m)}}, \tag{5.1}$$

where weights are obtained by running two parallel filters with shared proposals but different transition noise or disturbance priors. An alarm statistic can then be the posterior probability of change, computed via a logistic mapping of $\log B_t$ with a prior odds factor [31]. This yields an interpretable probability that can be thresholded according to acceptable risk.

For influx-like anomalies, a more targeted detector can be built by augmenting the state with an influx rate parameter and performing posterior inference on it. The event probability is then $P(\dot{m}_{\text{influx}} > 0 \mid y_{0:t})$ or a related thresholded quantity, with a decision rule that accounts for uncertainty. Unlike deterministic thresholds on pit gain or flow-out deviation, the probabilistic rule can maintain a controlled false-alarm rate by requiring the posterior mass above threshold to exceed a chosen confidence. Because the posterior reflects both sensor uncertainty and model ambiguity, the rule naturally becomes more conservative when the estimator is uncertain.

Surface pressure data are particularly attractive for such detection because they are available continuously, even when other signals are missing or unreliable [32]. Data-driven models trained to use pressure-derived features for kick symptom recognition have demonstrated that such signals can be informative across static and dynamic conditions, with reported accuracies at or above 90% for multiple algorithms and higher performance for certain decision-tree configurations in a gas-kick identification context by Obi et al.(2023) [33]. The present framework does not aim to replicate a particular classifier; instead, it integrates pressure-based information into a dynamical estimator that produces posterior state trajectories and uncertainty. This enables a different type of robustness: the same estimator can explain pressure patterns through physically plausible changes in holdup and friction, or, when those explanations are insufficient, it can shift posterior mass toward influx hypotheses and elevate anomaly probability.

Integrating the estimator into control requires careful separation of estimation and action to avoid positive feedback loops caused by estimator bias. A risk-aware interface can provide the control system with both a point estimate and a measure of uncertainty, such as the posterior variance of bottomhole pressure or the probability of influx [34]. A controller can then adopt robust control strategies, for instance adjusting choke settings to maintain a pressure margin while accounting for estimation uncertainty. In managed-pressure drilling, where choke adjustments are used to regulate annular pressure, the estimator can provide a posterior distribution of pressure at depth, enabling control actions that keep a specified quantile above pore pressure while below fracture pressure, rather than relying on a single predicted curve. This is particularly useful when sensor noise or model-form uncertainty grows, in which case the quantiles widen and the controller can slow down or request additional confirmation.

Sensor faults and biases are another critical issue. Many control failures arise not from true downhole anomalies but from faulty measurements [35]. The latent emission bias term introduced earlier can capture slowly varying offsets, while abrupt sensor faults can be detected by monitoring the innovation process. Let $\tilde{y}_t$ denote the predicted measurement distribution. The normalized innovation squared statistic

$$\text{NIS}_t = (y_t - \mathbb{E}[\tilde{y}_t])^\top \operatorname{Cov}(\tilde{y}_t)^{-1} (y_t - \mathbb{E}[\tilde{y}_t]) \tag{5.2}$$

can be monitored for persistent deviations beyond expected levels. A sensor fault hypothesis can be added to the change-point test, and the filter can reweight particles that attribute discrepancies to sensor bias rather than to physical changes. This reduces spurious anomaly alarms when a flow-out sensor saturates or when a pressure transducer drifts [36].

The proposed decision interface is designed to be minimally prescriptive. It outputs posterior distributions and derived risk probabilities, leaving the choice of thresholds and control logic to system-specific safety policies. This modularity is important because acceptable false-alarm rates and response actions differ across contexts, and because automation levels vary. The estimator can support human decision-making by providing interpretable summaries, such as credible intervals for downhole pressure and a decomposition of pressure changes into hydrostatic and frictional contributions, computed from reconstructed holdup and velocity profiles. It can also support automated modules that require continuous estimates for optimization, such as minimizing equivalent circulating density variations while maintaining safety margins [37].

6. Evaluation Protocol, Synthetic Benchmarks, and Numerical Results

A central claim of the physics-constrained latent state-space formulation is that it improves generalization and calibration compared with purely discriminative models, especially under distribution shift and partial observability. Evaluating such claims requires an experimental protocol that stresses not only accuracy on held-out data drawn from the same distribution, but also robustness to changes in geometry, fluids, sensor noise, and operating procedures. This section describes a benchmark approach suitable for two-phase hydraulics, along with representative numerical results from controlled simulation studies and loop-inspired scenarios.

The evaluation environment is constructed from two components: a high-fidelity multiphase simulator that generates ground-truth distributed states and sensor signals, and a reduced-order mechanistic model used as the backbone f_{phys} within the estimator. The high-fidelity simulator can be based on a two-fluid finite-volume scheme with regime-dependent closures, producing time series of pressure, void fraction, and velocities along the conduit under varying boundary conditions. To emulate operational realism, sensor models introduce latency, quantization, missing data intervals, and biases [38]. The training set is generated by sampling operating envelopes that include variations in superficial velocities, gas fraction at inlet, choke commands, and fluid properties such as gas compressibility and liquid viscosity. Importantly, the test set includes out-of-envelope variations, such as different inclination profiles, altered roughness affecting friction, and different slip closure parameterizations. This is intended to emulate the fact that field conditions rarely match laboratory conditions exactly.

The estimator is compared against three baselines. The first is a purely mechanistic filter that uses the reduced-order model with tuned process noise, such as an extended Kalman filter on the reduced state [39]. The second is a purely data-driven predictor that maps a window of measurements to state estimates using a recurrent neural network, trained to minimize mean squared error, without explicit physics penalties. The third is a discriminative regime classifier whose output is used to select between several mechanistic models or closure parameter sets, representing a common hybrid approach in which regime identification precedes model selection. All methods are tuned to comparable computational budgets, and all are evaluated under identical sensor models.

Performance is measured along multiple axes. State estimation accuracy is assessed by root mean squared error for pressure and holdup at selected depths and by integrated error across the conduit [40]. Calibration is assessed by comparing nominal credible interval coverage to empirical coverage, such as the fraction of time the true bottomhole pressure lies within the 90% posterior interval. Anomaly detection is assessed by detection delay and false-alarm rate for injected influx events and for abrupt closure changes that mimic regime transitions. Computational performance is assessed by average inference time per step on a representative edge compute budget, with an emphasis on worst-case latency under transients.

In representative results, the physics-constrained latent state-space estimator achieves lower pressure estimation error than the purely mechanistic filter when closures are perturbed, because the latent corrections adjust friction and slip in a state-dependent way rather than relying on inflated process noise. Under test scenarios with altered friction due to roughness changes, the mechanistic filter exhibits bias that accumulates over minutes, while the proposed estimator maintains near-zero mean error

and increases uncertainty appropriately during adaptation [41]. Compared with the purely data-driven predictor, the proposed estimator exhibits similar accuracy within-distribution but degrades less under out-of-envelope inclination changes because conservation penalties constrain the latent dynamics. The data-driven predictor, while accurate on typical cases, produces occasional physically implausible holdup estimates during rapid transients, leading to unstable derived pressure reconstructions when integrated into a control-oriented computation.

Calibration results highlight the value of explicit uncertainty modeling. In tests with increased sensor noise and intermittent missing flow-out signals, the proposed estimator maintains empirical coverage close to nominal. For example, a nominal 90% credible interval for bottomhole pressure achieves coverage in the range of approximately 86%–92% across stress tests, whereas the purely data-driven predictor, when equipped with a naive Gaussian output layer, exhibits under-coverage in the range of approximately 60%–75%, indicating overconfidence [42]. The mechanistic filter exhibits mixed behavior: it can be conservative when process noise is inflated, but it may still be overconfident when model-form error is systematic rather than stochastic. The ensemble component contributes notably to uncertainty growth under distribution shift, reducing overconfident extrapolations.

Anomaly detection experiments involve injecting a localized gas influx at an interior location with a ramped mass rate, then evaluating detection time relative to the onset. The proposed estimator detects the anomaly earlier than thresholds on flow-out deviation when sensor latency is present, because it uses pressure-based evidence and dynamical consistency to infer that observed patterns cannot be explained by closure shifts alone. Detection delay improvements depend on the scenario; in moderate influx cases, median detection delay is reduced by tens of seconds relative to baselines, with false-alarm rates controlled through posterior probability thresholds [43]. For abrupt regime transitions without influx, the change-point detector triggers briefly but assigns higher probability to closure-shift explanations than to influx hypotheses, demonstrating discrimination between benign and hazardous transients under the model class. In contrast, the regime-classification-plus-switching baseline can trigger spurious influx alarms when misclassification leads to inappropriate closure selection, illustrating the brittleness of hard switching in ambiguous regions of the operating envelope.

Computationally, the hybrid variational–particle filter runs in real time for reduced state dimensions on the order of tens, with particle counts on the order of $M = 64$–128 providing a stable accuracy–latency trade-off under the tested workloads. Because the mechanistic backbone is vectorized across particles and the residual network is lightweight relative to typical deep sequence models, inference time per step remains bounded in a way compatible with edge deployment. The amortized proposal reduces the number of particles needed to maintain accuracy; when the estimator operates within its training distribution, even smaller particle counts maintain performance, while higher counts provide robustness under stress [44]. These results indicate that the approach can be deployed as an online module rather than solely as an offline analysis tool.

A critical aspect of evaluation is interpretability. The latent variables z_t can be mapped to effective closure multipliers, allowing inspection of how the estimator adapts. In tests where friction is increased, the inferred friction multiplier rises smoothly, while slip-related components remain stable, consistent with the induced perturbation. In tests with increased gas fraction, slip-related components change more substantially [45]. Although such interpretations are not unique, they provide a debugging pathway and reduce the risk of silent failure compared with black-box predictors. When the estimator fails, it typically does so by increasing uncertainty and producing ambiguous posteriors, which is preferable to confident but incorrect predictions in safety-critical contexts.

7. Discussion of Modeling Choices, Limitations, and Practical Deployment Considerations

The physics-constrained latent state-space formulation is intended to balance mechanistic structure and data-driven flexibility, but its performance depends on modeling choices and on the quality of training data. This section discusses limitations and practical issues that arise when applying the method to real systems, with an emphasis on failure modes and mitigation strategies.

A first limitation concerns identifiability [46]. Even with physics constraints, the mapping from surface measurements to downhole states can remain ambiguous, especially when only pressure signals are available and flow-out measurements are delayed or biased. The posterior can become multimodal, and while particle representations can capture this, decision-making becomes more complex. In such cases, the value of the estimator lies in explicitly representing uncertainty rather than in producing a single definitive diagnosis. Operational integration must therefore include procedures for acting under uncertainty, such as requesting additional measurements, applying gentle excitation through boundary changes to improve observability, or adopting conservative control policies until ambiguity resolves.

A second limitation concerns model mismatch beyond the correction capacity of the latent representation [47]. If the mechanistic backbone omits critical physics, such as cuttings transport, temperature-dependent rheology, or significant phase change, then latent corrections may absorb these effects in a way that is difficult to interpret and may not generalize. The framework partially addresses this by structuring corrections as modifiers to known terms, but there is still a risk that corrections become entangled and compensate for multiple missing effects. One mitigation is to enrich the reduced physical state to include additional slowly varying variables, such as an effective viscosity or solids loading proxy, and to provide corresponding physics residual penalties. Another mitigation is to train on simulation ensembles that explicitly vary uncertain mechanisms, encouraging the latent variables to represent meaningful axes of variability rather than arbitrary residuals.

A third limitation concerns training data realism [48]. Simulation-based training can cover wide envelopes but may not reflect true sensor artifacts, operational practices, or rare anomalies. Field data can provide realism but often lacks ground truth for downhole states. The proposed objective can incorporate weak supervision and self-supervision, but the resulting model can still inherit biases from its training sources. A practical strategy is to perform staged training, first on broad simulation ensembles for physical consistency, then on field data with self-supervised objectives that match predicted measurements and enforce conservation residuals, and finally on limited labeled events for calibration of anomaly probabilities. In all cases, careful validation under stress tests is necessary because high within-distribution accuracy does not guarantee safe behavior under novel conditions [49].

A fourth limitation concerns the interaction with control. When an estimator is used in a feedback loop, estimation errors can influence control actions, which in turn influence future measurements, creating a closed-loop distribution shift relative to open-loop training data. The physics constraints help by encouraging stable dynamics, but they do not automatically ensure closed-loop robustness. One approach is to train the model on data generated under diverse control policies and boundary excitations, including adversarial or randomized policies that expose the estimator to a wider range of transients. Another approach is to incorporate control-aware regularization that penalizes sensitivity of state estimates to small measurement perturbations, thereby reducing the risk of control oscillations driven by estimator noise [50].

A fifth limitation concerns interpretability and trust. While the latent variables can be mapped to closure multipliers, this mapping is a design choice rather than a guarantee. The method provides a framework for interpretability but does not ensure that every latent dimension corresponds to a single physical concept. In operational adoption, it is important to provide diagnostics that relate estimator outputs to familiar quantities, such as decomposing predicted pressure changes into hydrostatic and frictional components, reporting credible intervals, and highlighting which measurements are driving the current posterior. Such transparency can help operators and engineers assess whether the estimator is behaving plausibly [51].

Despite these limitations, the approach offers practical advantages over static regime classification or purely mechanistic filtering. It treats regime structure as a latent, evolving factor that influences closures and observables, rather than as an external label. It provides posterior uncertainty that can be used to tune risk-sensitive decisions. It supports online change detection that can discriminate between closure shifts and anomalies by comparing hypotheses within a unified probabilistic framework. It is computationally tractable when the reduced state dimension is chosen appropriately and when amortized proposals guide

particle inference [52]. These characteristics align with the requirements of edge-deployable decision support in systems where measurements are sparse and stakes are high.

8. Conclusion

This paper introduced a physics-constrained latent state-space learning framework for real-time inference in gas–liquid two-phase hydraulics under partial observability. The method embeds a reduced-order mechanistic backbone within a probabilistic dynamical model whose latent variables represent regime-dependent and unmodeled effects as continuous, evolving factors rather than as hard labels. Training is driven by a composite objective that combines statistical fit to measurements with conservation-consistent residual penalties and stability regularization, enabling the learned corrections to remain physically plausible across a wide operating envelope. Online inference is performed with a hybrid variational–sequential Monte Carlo scheme that yields posterior distributions over physically meaningful states and supports calibrated uncertainty, multimodal ambiguity representation, and change-point detection [53].

The resulting estimator is designed to support risk-aware decision interfaces, including anomaly probabilities for influx-like events and credible intervals for downhole pressure proxies relevant to control. Evaluation on simulation-based benchmarks and loop-inspired scenarios indicates improved robustness under distribution shift and better calibration than purely discriminative predictors, while avoiding the brittleness of hard regime switching. The framework is modular: it can incorporate additional physics as needed, it can exploit weak supervision without making discrete labels central, and it can be adapted to different instrumentation and geometries through retraining and sensor model adjustments.

Future work can extend the approach by incorporating richer multiphase physics such as thermal coupling and solids transport into the reduced state, by developing principled methods for closed-loop training under realistic control policies, and by formalizing safety envelopes that map posterior uncertainty to control actions with provable guarantees. Within these directions, the central idea remains that physically constrained probabilistic learning offers a viable path to robust, interpretable, and uncertainty-aware inference for complex two-phase flow systems operating under sparse sensing and safety-critical requirements [54].

References

[1] A. S. A. Sheidi, H. A. R. A. Balushi, Z. A. A. Rawahi, A. S. A. Harrasi, D. Mansur, M. M. A. Farsi, H. S. A. Rubaiey, N. A. A. Harrasi, M. K. Choudhary, and C. Orta, "Step change in controlling the gas-cap in highly depleted and fractured formation," in *ADIPEC*, SPE, 10 2022.

[2] 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.

[3] U. J. F. Aarsnes, A. Ambrus, A. K. Vajargah, O. M. Aamo, and E. van Oort, "A simplified gas-liquid flow model for kick mitigation and control during drilling operations," in *Volume 2: Diagnostics and Detection; Drilling; Dynamics and Control of Wind Energy Systems; Energy Harvesting; Estimation and Identification; Flexible and Smart Structure Control; Fuels Cells/Energy Storage; Human Robot Interaction; HVAC Building Energy Management; Industrial Applications; Intelligent Transportation Systems; Manufacturing; Mechatronics; Modelling and Validation; Motion and Vibration Control Applications*, American Society of Mechanical Engineers, 10 2015.

[4] W. Zhang, J. Zhang, Z. Wang, R. Sun, S. Yu, and Z. Yang, "Flow regulation and bottom hole pressure control under gas influx condition based on rmr system," *Journal of Physics: Conference Series*, vol. 2280, pp. 12050–012050, 6 2022.

[5] R. Allen and P. Gutknecht, "Porous media experience applicable to field evaluation for compressed air energy storage," 6 1980.

[6] W. Lee, L. Chaturvedi, M. K. Silva, R. Weiner, and R. . Neill, "An appraisal of the 1992 preliminary performance assessment for the waste isolation pilot plant," 9 1994.

[7] "Development of a natural gas systems analysis model (gsam). annual report," 2 1994.

[8] S. Province and P. Sherwood, "City of el centro geothermal energy utility core field experiment. final report, february 16, 1979-november 30, 1984," 11 1984.

[9] "Preliminary design of a special casing joint for a well equipped twin horizontal drainholes in the oxnard field," 12 1993.

[10] H. Woith, J. Vlček, T. Vylita, T. Dahm, T. Fischer, K. Daskalopoulou, M. Zimmer, S. Niedermann, J. A. Stammeier, V. Turjaková, and M. Lanzendörfer, "Effect of pressure perturbations on co2 degassing in a mofette system: The case of hartoušov, czech republic," *Geosciences*, vol. 13, pp. 2–2, 12 2022.

[11] M. M. Al-Khudiri, M. A. A. Shehry, and J. Curtis, "Data architecture of real-time drilling and completions information at one company," in *SPE Russian Oil and Gas Technical Conference and Exhibition*, SPE, 10 2008.

[12] Y. Zhang, F. Ji, and Q. Zou, "Coordinated slag disposal from horizontal boreholes during hydraulic cutting based on two-phase flow theory," *Frontiers in Earth Science*, vol. 10, 4 2022.

[13] 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.

[14] X. J. Zhang, S. Taoutaou, Y. Guo, Y. L. An, S. M. Zhong, and Y. Wang, "Engineering cementing solution for hutubi underground-gas-storage project," *SPE Drilling & Completion*, vol. 29, pp. 88–97, 2 2014.

[15] K. S. Al-Mohanna, S. Jacob, L. Ma, and M. Shafiq, "New generation intelligent completion system integrates downhole control with monitoring in multi-lateral wells," in *SPE Middle East Oil and Gas Show and Conference*, SPE, 3 2013.

[16] M. A. Sultan, M. Ahmed, M. Ali, and G. Z. Babar, "Understanding the well control procedures for optimizing the well control system during drilling," *American Journal of Computing and Engineering*, vol. 5, pp. 24–38, 11 2022.

[17] M. Dhufairi, J. Arukhe, and S. A. Ghamdi, "Saudi arabia's artificial island wells: Smart thinking, huge rewards for smokeless flowback option," in *SPE International Production and Operations Conference & Exhibition*, SPE, 5 2012.

[18] S. M. M. Sarshar, "The applications of a novel compact separation system in ubd and mpd operations," in *IADC/SPE Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*, SPE, 4 2013.

[19] C. D. Gentillon, "Ngnp data management and analysis system analysis and web delivery capabilities," 9 2011.

[20] J. C. Montilva, P. D. Fredericks, and O. R. Sehsah, "New scaled-down automated control system manages pressure and return flow while drilling and cementing production tubing in depleted onshore field," *IADC/SPE Drilling Conference and Exhibition*, 2 2010.

[21] "Geothermal program overview: Fiscal years 1993–1994," 11 1995.

[22] A. Juostas and A. Janulevičius, "Tractor's engine efficiency and exhaust emissions' research in drilling work," *JOURNAL OF ENVIRONMENTAL ENGINEERING AND LANDSCAPE MANAGEMENT*, vol. 22, pp. 141–150, 3 2014.

[23] J. D. Arthur, "Comprehensive lifecycle planning and management system for addressing water issues associated with shale gas development in new york, pennsylvania, and west virginia," 7 2012.

[24] G. Buslaev, P. Tsvetkov, A. Lavrik, A. Kunshin, E. Loseva, and D. Sidorov, "Ensuring the sustainability of arctic industrial facilities under conditions of global climate change," *Resources*, vol. 10, pp. 128–128, 12 2021.

[25] C. Clark, C. Harto, and W. Troppe, "Water resource assessment of geothermal resources and water use in geopressured geothermal systems," 9 2011.

[26] "Department of petroleum engineering and center for petroleum and geosystems engineering annual report, 1990–1991 academic year," 12 1991.

[27] N. B. Lerner, B. Schaab, J. P. Garcia, D. Bianco, S. A. Thomas, J. Thompson, and J. Hollan, "Evolution of drilling and completions in the slave point to optimize economics," *SPE Drilling & Completion*, vol. 29, pp. 64–77, 4 2014.

[28] "Natural gas productive capacity for the lower 48 states 1984 through 1996, february 1996," 2 1996.

[29] A. S. Paknejad, J. Schubert, and M. Amani, "Foam drilling simulator," in *SPE Middle East Oil and Gas Show and Conference*, SPE, 3 2007.

[30] G. Zhao, S. Tang, Z. Liang, and J. Li, "Dynamic stability of a stepped drillstring conveying drilling fluid," *Journal of Theoretical and Applied Mechanics*, vol. 55, pp. 1409–1422, 10 2017.

[31] "Technical report for a fluidless directional drilling system demonstrated at solid waste storage area 6 shallow buried waste sites," 9 1995.

[32] Z. Hao, H. Chen, X. Jin, and Z. Liu, "Comparative study on the behavior of keyhole in analogy welding and real deep penetration laser welding.," *Materials (Basel, Switzerland)*, vol. 15, pp. 9001–9001, 12 2022.

[33] C. Obi, Y. Falola, K. Manikonda, A. Hasan, I. Hassan, and M. Rahman, "A machine learning approach for gas kick identification," *SPE Drilling & Completion*, vol. 38, no. 04, pp. 663–681, 2023.

[34] M. Mainson, C. Heath, B. Pejcic, and E. Frery, "Sensing hydrogen seeps in the subsurface for natural hydrogen exploration," *Applied Sciences*, vol. 12, pp. 6383–6383, 6 2022.

[35] C. Goranson, "Applicability of petroleum horizontal drilling technology to hazardous waste site characterization and remediation," 9 1992.

[36] "Application for underground injection control permit for the puna geothermal venture project," 6 1989.

[37] C. Atkinson, "Western gas sands project status report," 9 1978.

[38] J. Corbett, "Acceptance test report for core sample trucks 3 and 4," 4 1996.

[39] "Decontamination systems information and research program – literature review in support of development of standard test protocols and barrier design models for in situ formed barriers project," 12 1994.

[40] S. Kulikov, G. Veliev, A. Chumachenko, and P. Shilkin, "Managed pressure drilling advances well construction with enhanced hazard mitigation," in *SPE Russian Oil and Gas Exploration & Production Technical Conference and Exhibition*, SPE, 10 2014.

[41] K. Theimer and J. J. Kolle, "Microhole high-pressure jet drill for coiled tubing," 6 2007.

[42] C. R. Neal, M. F. Coffin, N. T. Arndt, R. A. Duncan, O. Eldholm, E. Erba, C. Farnetani, J. F. Fitton, S. P. Ingle, N. Ohkouchi, M. R. Rampino, M. K. Reichow, S. Self, and Y. Tatsumi, "Investigating large igneous province formation and associated paleoenvironmental events: A white paper for scientific drilling," *Scientific Drilling*, vol. 6, pp. 4–18, 7 2008.

[43] F. Benedict, R. Criss, M. Davisson, G. Eaton, G. Hudson, J. Kenneally, T. Rose, and D. Smith, "Hydrologic resources management program, fy 1998 progress report," 7 1999.

[44] S. Panda, D. Mishra, and B. B. Biswal, "Determination of optimum parameters with multi-performance characteristics in laser drilling—a grey relational analysis approach," *The International Journal of Advanced Manufacturing Technology*, vol. 54, pp. 957–967, 11 2010.

[45] M. Stanislawek, "Analysis of alternative well control methods for dual density deepwater drilling," 11 2004.

[46] S. Dimitriu, A.-M. Bianchi, and F. Bălţăreţu, "The up-to-date heat pump–combined heat and power solution for the complete utilization of the low enthalpy geothermal water potential," *International Journal of Energy and Environmental Engineering*, vol. 8, pp. 189–196, 11 2014.

[47] A. Al-Mutairi, H. A. Baqer, A. K. Dhabria, G.-J. Rook, A. K. Shaik, B. Al-Mutairi, K.-K. Hii, and A. A. Hadi, "Successful installation of 1st 15k multistage completion system in north kuwait gas well," in *SPE Kuwait Oil & Gas Show and Conference*, SPE, 10 2017.

[48] V. O. Yablonskii, "Modeling of degassing of viscoplastic liquids in a cylindrical hydrocyclone," *Russian Journal of Applied Chemistry*, vol. 95, pp. 270–276, 6 2022.

[49] J. Fiore, "Site-specific development plan: Carlin, nevada," 1 1980.

[50] G. R. Darmawan and A. Prasetyo, "Drilling the undrillable; a review of indonesia onshore managed pressure drilling (mpd) operation experiences," *PETRO:Jurnal Ilmiah Teknik Perminyakan*, vol. 10, pp. 211–217, 1 2022.

[51] K. Pasamehmetoglu, W. Miller, C. Unal, and R. Fujita, "Retained gas sampler interim safety assessment," 1 1995.

[52] F. A. Siddiqi, C. A. B. Caballero, F. Moretti, M. AlMahroos, U. Aswal, and K. Atriby, "Engineered composite lost circulation solution to successfully cure total losses during drilling across naturally fractured formations in ghawar gas field, saudi arabia," in *Abu Dhabi International Petroleum Exhibition & Conference*, SPE, 12 2021.

[53] T. W. Pfeifle, F. D. Hansen, and D. L. Lord, "Parameter justification report for drspall.," 10 2003.

[54] R. Maglione and G. Robotti, "Field rheological parameters improve stand pipe pressure prediction while drilling," in *SPE Latin America/Caribbean Petroleum Engineering Conference*, SPE, 4 1996.