International Journal of Applied Sciences, Research, and Innovation

OPEN ACCESS

FULL PAPER | IJASRI | OPEN

Perception-Aware Modeling and Vision-Based Identification of Two-Phase Flow During Gas Kick Migration

Brahim Ould Ahmed[1], **El Hacen Cheikh**[2], **Mohamed Yahya**[3]

[1]Institut Supérieur du Numérique, Zone Universitaire, Rue 42-178, Sebkha, Nouakchott 5035, Mauritania
[2]Université de Zouerate, Avenue de l' Indépendance, Quartier Kedya, Zouerate 1120, Mauritania
[3]Institut Supérieur des Sciences Appliquées et de Technologie de Kiffa, Route Nationale N1, Kiffa 6200, Mauritania

Abstract

Background conditions in drilling systems are increasingly shaped by the interaction of multiphase transport, uncertain downhole sensing, and limited direct observability of annular flow structure. Gas kick migration is therefore not only a hydraulic control problem but also an inference problem, because the quantities that matter operationally, such as gas holdup distribution, slip behavior, and transition to severe flow regimes, are only partially observed through pressure, return flow, and sparse surface measurements. This paper develops a perception-aware framework for gas kick analysis in vertical annuli that unifies mechanistic transport modeling with computer vision based flow regime identification. The central thesis is that regime recognition should be treated as an explicit state variable in well control and not as a descriptive afterthought. A physically structured model is formulated for transient gas migration under shut-in and circulating conditions, with closures that depend on local regime state, dissolution dynamics, and annular geometry. A complementary visual learning pipeline is then designed for regime classification from high-resolution annulus imagery, using temporal representation learning, calibration, and cross-domain adaptation to bridge laboratory observations and operational monitoring. The two components are fused through probabilistic state estimation so that image-derived regime beliefs update hydraulic state trajectories in real time. The resulting architecture supports earlier detection of hazardous regime transitions, more stable uncertainty quantification, and improved interpretation of surface indicators during kick handling. The paper argues that the most practical path toward robust digital well surveillance is not purely mechanistic or purely data-driven, but a coupled estimation architecture in which transport physics and visual evidence constrain one another continuously.

1. Introduction

Gas kicks remain one of the defining transient hazards in drilling because a small influx may evolve into a severe control problem before its structure is fully observable [17]. The principal difficulty is not merely that gas is compressible, buoyant, and composition dependent, but that the migration process reshapes the annular flow field as it rises. The same influx may manifest as an elongated bubble, dispersed packets, churn-like motion, or a mixed regime whose pressure signature does not admit a unique interpretation from conventional standpipe or casing measurements alone. In field operations, the operator reasons from indirect evidence while the flow system itself changes regime, density, slip, and frictional behavior across depth and time. That mismatch between hidden flow structure and decision-critical quantities has historically encouraged simplified models in which regime is implicit or fixed. Such simplifications are attractive for computation, yet they obscure the very mechanism by which pressure evolution becomes difficult to interpret. Once the state estimation problem is posed correctly, regime identification becomes central rather than peripheral.

A broader research opportunity therefore emerges at the intersection of drilling hydraulics, multiphase transport, machine vision, and state estimation. In that broader setting, gas-kick monitoring is best understood as a partially observed dynamical system in which two types of information coexist. The first type is mechanistic: conservation laws, hydrostatic balance, friction, slip, dissolution, and boundary conditions imposed by well control actions. The second type is perceptual: the visual texture, topology, and temporal organization of gas-liquid structures seen in a transparent annulus or other observational modality. The usual separation between these information types is increasingly difficult to justify. Mechanistic models without regime-aware observation models can fit pressure trends while misrepresenting local flow topology. Vision models without transport constraints can label frames convincingly while violating continuity, mass balance, or transition plausibility [14]. A technically serious monitoring framework must therefore couple them, not merely compare them.

This paper is organized around the thesis that gas-kick surveillance should be formulated as perception-aware multiphase state estimation. In that formulation, local flow regime is treated as a latent but inferable field that mediates between images, hydraulic closures, and control decisions. The concept is broader than the interpretation of any individual experiment. It implies a digital twin in which regime-conditioned closures evolve with pressure, gas fraction, solubility, and imposed liquid flow, while a visual learner supplies probabilis-

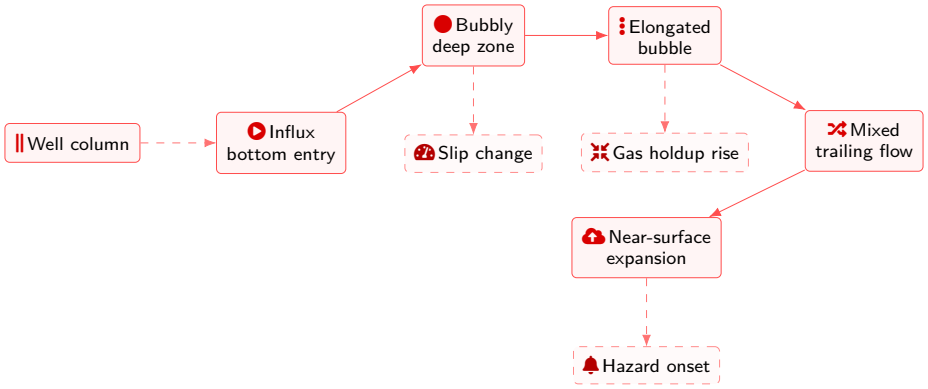


Figure 1. Conceptual migration path of a kick in a vertical annulus. A bottom influx evolves from relatively dispersed structures into elongated and mixed forms as it rises, while expansion toward the surface increases gas holdup and the likelihood of hazardous transient behavior. The diagram emphasizes morphology progression rather than a single fixed regime assumption.

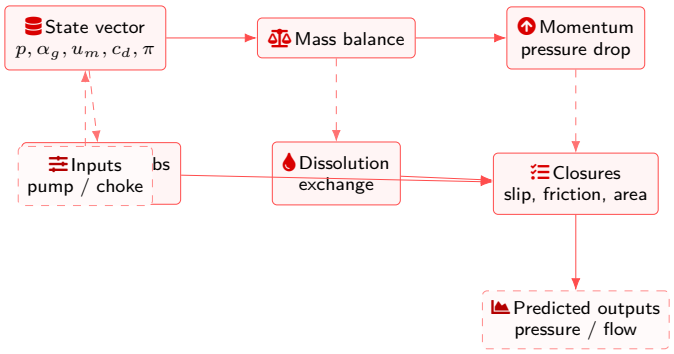


Figure 2. Mechanistic state-space structure. The transient annular model evolves pressure, gas holdup, mixture motion, dissolved gas, and regime probabilities. Closure selection is conditioned on the regime field so that slip, friction, and interfacial area adapt as morphology changes during migration.

tic evidence about local morphology and transition direction. The value of such a formulation lies in operational inference: it can reduce ambiguity when two pressure traces are compatible with different regime realizations, and it can temper overconfident image labels when transport dynamics render a transition unlikely.

The motivation for this perspective is visible in laboratory observations of kick migration. Obi et al. (2022a) experimentally investigated shut-in behavior in a 140 ft annulus and showed that the initial Taylor-bubble signature is comparatively insensitive to kick pressure and duration [32], whereas the downstream or trailing flow behavior changes materially with gas mass, injection history, and gas type, especially toward the surface where the evolving distribution becomes more complex. The significance of that observation for the present paper is methodological rather than bibliographic. It suggests that the part of the transient most likely to challenge control and diagnosis is not the first easily recognizable structure but the subsequent evolution, where flow morphology, slip, and apparent turbulence begin to diverge from simple template assumptions. A monitoring architecture that collapses the transient into one canonical regime is therefore structurally incomplete.

The same insight extends under circulation, but the problem becomes harder. Once a liquid background flow is imposed, the annulus no longer supports a migration pattern governed only by buoyancy and compressibility [26]. Additional momentum input, frictional losses, and altered liquid distribution modify both the transport velocity and the texture of the gas-liquid field. In practical terms, the operator is then trying to separate pressure changes due to control actions from pressure changes due to regime redistribution. This is a classic observability problem. Surface signals contain information, but not enough information unless they are interpreted through a model that admits evolving regime-dependent closures and, ideally, auxiliary evidence about the local phase arrangement.

For that reason, the paper develops an independent technical program with three tightly connected layers. The first layer is a mechanistic transient model for vertical annular kick migration under shut-in and circulating conditions, including free gas transport, dissolved gas exchange, pressure evolution, and regime-dependent slip and friction closures. The second layer is a computer vision model for two-phase regime identification in annular images, designed not merely as a classifier of isolated frames but as a temporal inference engine that respects transition continuity and uncertainty. The third layer is a probabilistic fusion architecture through which image-based regime beliefs update the hydraulic state estimate, allowing the mechanistic model and the perception model to discipline one another.

This framing deliberately avoids reducing the problem to a

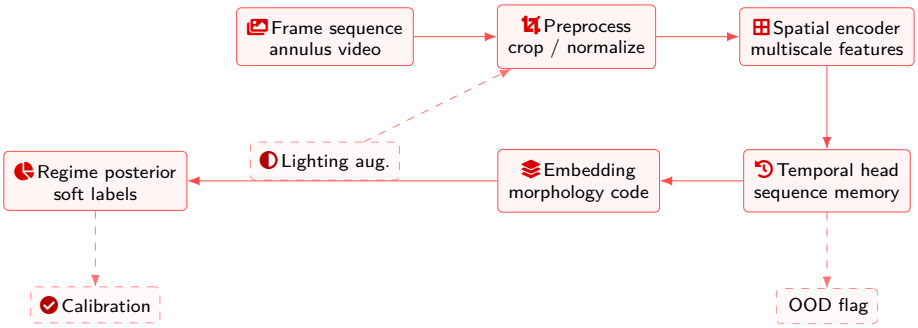


Figure 3. Computer-vision pipeline for annular flow recognition. A sequence model extracts multiscale spatial information and temporal persistence to produce morphology embeddings and calibrated regime posteriors. The additional out-of-distribution branch supports safe deployment by reducing trust in visual evidence when operating conditions drift away from the training domain.

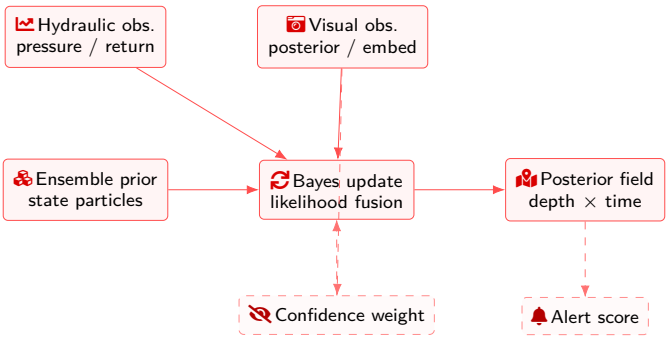


Figure 4. Perception–physics filtering loop. Hydraulic measurements and visual evidence are assimilated into an ensemble-based posterior over the hidden annular state. Confidence weighting allows image information to be trusted strongly in-distribution and downweighted when calibration or domain alarms indicate unreliable perceptual evidence.

review of particular experiments. The experiments matter because they reveal physically relevant transition behavior and provide data structures from which one can learn. But the paper′s objective is larger: to describe how annular two-phase flow recognition can become an operational variable in drilling surveillance and control. That objective requires discussing representation learning, calibration, domain shift, dissipation modeling, closure selection, and estimator design at a level suitable for integration into digital well control systems.

An additional reason the problem deserves this integrated treatment is that direct visual observation, once available in a laboratory setting, provides a richer target than conventional labels alone. Flow regimes in annuli are not merely nominal classes; they are spatiotemporal organizations of interfaces, bubbles, wakes, and liquid films [3]. As a result, a useful model must distinguish not only between classes such as bubbly, elongated-bubble, and churn-like behavior, but also between degrees of disorder, bubble clustering, intermittency, and the onset of regime mixing. Classical handcrafted image descriptors are not sufficient for this task because they struggle to represent both fine interface detail and longer-range temporal context. Modern deep sequence models are better suited, yet their outputs must still be constrained by fluid mechanics if they are to be trusted for safety-relevant inference.

The remainder of the paper proceeds from this premise. The next section establishes the physical setting and the design objectives for a perception-aware gas-kick model. The following section formulates the mechanistic transport problem and its regime-conditioned closures. A subsequent section develops computer vision models for annular regime identification using a vertical two-phase image dataset and explains how such models should be trained, calibrated, and evaluated for deployment in hydraulic diagnostics. The paper then presents a fusion framework for probabilistic state estimation and decision support, followed by a discussion of experimentalization, validation, transfer, and deployment. The overall goal is to show that visual flow information can be elevated from descriptive laboratory evidence to a formal component of gas-kick monitoring architecture.

2. Physical Setting and Design Objectives

The annular gas-kick problem sits at an awkward junction of scale interactions. At the largest scale, the wellbore or experimental annulus is a vertical conduit in which hydrostatic head, imposed circulation, and global boundary conditions dominate. At the intermediate scale, the gas phase redistributes through slip, coalescence, breakup, and possible dissolution or exsolution. At the smallest resolvable scale, the observable texture of the flow is determined by interfaces, wake structures, bubble size distributions, and the geometric asymmetry created by drill-pipe eccentricity or local wall effects [21]. In practical models these scales are rarely resolved simultaneously. Instead, large-scale solvers use closure relations for slip and friction that presuppose a regime, while vision systems tend to infer regime from local morphology

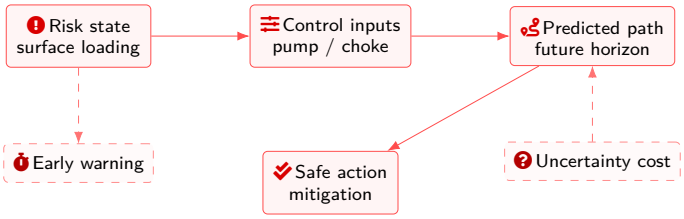


Figure 5. Risk-aware decision layer. The posterior state informs future trajectory prediction under candidate control actions, allowing the system to favor responses that not only suppress pressure excursions but also reduce uncertainty about severe regime transitions during kick handling.

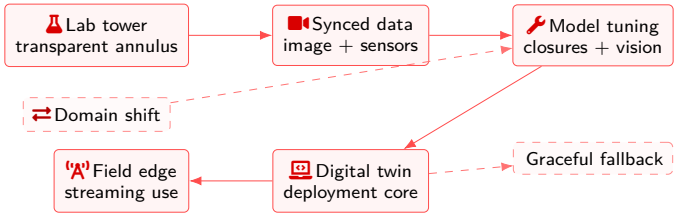


Figure 6. Experimentalization-to-deployment path. Transparent-annulus experiments create synchronized visual and hydraulic data for closure identification, vision training, and fusion validation. The resulting digital twin then transfers toward field use with explicit handling of domain shift and conservative fallback when visual evidence becomes unreliable.

without a formally coupled transport state. A perception-aware framework must make that scale hierarchy explicit so that information can be exchanged across it rationally.

Consider first the hydraulic variables that matter operationally. The pressure field controls formation safety margin, well control action, and surface alarm logic. The gas volume fraction field controls density reduction, mixture compressibility, and the magnitude of hydrostatic unloading. The liquid velocity field controls convective transport and the strength of momentum exchange under circulation. Dissolved gas concentration matters when carbon dioxide or other soluble components are present, because an apparent reduction in visible gas fraction at depth may be consistent with later exsolution near the surface. None of these variables is directly equivalent to the visible regime, yet each is influenced by it through constitutive behavior. A churn-dominated region and a bubbly region may support the same volumetric gas holdup at one location while producing different slip, pressure gradient, and temporal variability. The regime field is therefore not a cosmetic label; it is a closure selector and an observation mediator.

This point becomes more important in vertical annuli than in simple circular pipes. The annulus introduces curvature asymmetry, wall proximity effects, and in many cases eccentricity, all of which alter phase distribution and bubble migration paths. Local liquid films, bubble bypassing, and azimuthal nonuniformity become relevant even when a one-dimensional model is used at the system level [1]. The correct modeling response is not to abandon low-dimensional transport models, which remain indispensable for real-time use, but to acknowledge that their closures must absorb geometry-sensitive behavior. A vision system trained on annular imagery can help precisely because it sees morphological consequences of those geometry effects. If the visual evidence suggests a transition toward dispersed-bubble turbulence or churn-like intermittency, then the estimator should be allowed to update friction and slip in a regime-consistent direction even if scalar measurements are ambiguous.

The design objectives for such a system are more demanding than ordinary classification accuracy. The first objective is physical consistency. Any regime inference used in state estimation must be compatible with mass conservation, pressure continuity, and plausible transition rates. A classifier that changes abruptly between classes from one frame to the next without corresponding hydraulic support will destabilize the estimator. The second objective is uncertainty calibration. In safety-critical settings the cost of an overconfident wrong label is generally higher than the cost of an abstention or broad posterior. The third objective is temporal coherence. Flow regimes are persistent but not static; models should encode both inertia in the label process and the possibility of rapid transition under the right forcing. The fourth objective is domain robustness. Laboratory imagery obtained in controlled lighting and transparent equipment differs sharply from field camera feeds, borescope views, or surrogate sensing modalities [45]. The vision module must therefore expose uncertainty under domain shift rather than silently extrapolating.

A fifth objective concerns observability under limited instrumentation. Many field scenarios will not provide full-column imaging. Even in a laboratory tower, images may only be available at a few elevations. The estimator must then propagate information spatially and temporally, allowing localized regime evidence to influence the state elsewhere through the governing model. This is another reason to favor coupled estimation over standalone vision. The image is informative about local morphology, but the transport model determines what that local evidence implies for upstream and downstream states. Inference becomes a dynamic smoothing problem across depth and time.

A sixth objective is computational practicality. Real-time drilling surveillance cannot rely on prohibitively expensive simulations or overly deep models with unstable latency. The

Table 1. Core variables in the perception-aware gas-kick state representation.

Symbol	Meaning	Operational role	Type
$p(z,t)$	Pressure field	Governs safety margin, hydrostatics, and control response	Hydraulic
$\alpha_g(z,t)$	Gas volume fraction	Determines density reduction and compressibility	Hydraulic
$u_m(z,t)$	Mixture velocity	Controls transport and circulation-induced redistribution	Hydraulic
$c_d(z,t)$	Dissolved-gas concentration	Captures dissolution and later exsolution effects	Mass transfer
$r(z,t)$	Regime state	Selects slip, friction, and interfacial closures	Latent
$\pi_k(z,t)$	Regime probability	Represents uncertainty over admissible flow structures	Probabilistic

Table 2. Regime classes and their hydraulic implications in the annulus.

Regime	Visual signature	Expected closure tendency	Inference relevance
Bubbly	Small dispersed bubbles	Low-to-moderate slip, smoother friction response	Baseline transport state
Elongated-bubble	Coherent Taylor-like structures	Larger buoyancy-driven slip, organized interface motion	Early migration tracking
Dispersed-bubble turbulent	Dense, disordered bubble field	Elevated mixing, stronger momentum exchange	Circulation-sensitive behavior
Churn-like	Intermittent breakup and recirculation	Irregular slip, larger pressure fluctuation variance	Hazardous transition marker
Mixed transitional	Coexisting morphologies	Convex combination of neighboring closures	Soft classification region

mechanistic solver must be low order enough to run faster than data acquisition while retaining the essential nonlinearities of compressibility, slip, and dissolution. The vision model must compress spatial detail into temporally informed state variables without destroying edge features critical to interface recognition. The fusion scheme must update uncertainty without requiring a full high-dimensional posterior over every pixel and hydraulic variable simultaneously. Efficiency here is not a minor engineering detail; it determines whether perception-aware surveillance can ever leave the laboratory [38].

There is also a subtler design objective: the model should represent what is operationally actionable, not merely what is visible. A fine-grained taxonomy of dozens of visual motifs might improve image classification benchmarks, yet many of those distinctions may not correspond to materially different control consequences. The relevant representation must therefore be hierarchical. At the visual level, the model may learn rich embeddings of texture, interface elongation, bubble packing, and intermittency. At the hydraulic level, those embeddings should map to a smaller number of regime-conditioned closure families whose differences affect slip, friction, and pressure response. In this way, the system retains expressive perception while enforcing operational parsimony.

Another key consideration is how to represent the transition from an initially dominant elongated structure to the more heterogeneous trailing behavior that frequently governs diagnostic difficulty. In many transients the earliest stage is visually salient and comparatively easier to label. However, the inference problem becomes harder later, when the flow may contain mixed structures and the pressure signature reflects both accumulation and redistribution. A useful model must therefore avoid making the initial segment disproportionately influential. This can be achieved by emphasizing sequence models that preserve memory of transition history while continuing to attend to current local detail. It can also be aided by regime-conditioned priors in which early elongated structures increase the probability of certain subsequent states but do not determine them.

Finally, the design must accommodate soluble gas behavior [44]. When the gas phase exchanges mass with the liquid, the visible gas inventory no longer tracks the total gas inventory directly. A perception-only system could then underestimate risk during high-pressure segments where gas is partly dissolved, while a hydraulics-only system could misattribute later gas release if it lacks proper exsolution dynamics. Integrating a dissolved-gas state with visually inferred regime evidence offers a route to more stable interpretation. In practical terms, the model may treat visible regime as evidence about local free-gas configuration, while the transport state separately evolves dissolved concentration and pressure-dependent mass exchange.

These objectives motivate the remainder of the framework. The mechanistic component will be formulated to permit regime-conditioned closures, the computer vision component will be formulated to emit calibrated regime posteriors and feature embeddings, and the fusion component will be constructed to update the hidden hydraulic state under partial observation. The resulting architecture is not intended as a perfect simulator of all annular physics. Its purpose is narrower and more operationally valuable: to make gas-kick monitoring more observable by connecting what the flow

Table 3. Governing components of the mechanistic transport framework.

Model block	Principal quantity	Functional purpose	Regime dependence
Gas continuity	$\alpha_g \rho_g u_g$	Tracks free-gas migration and expansion	Indirect
Liquid continuity	$\alpha_l \rho_l u_l$	Preserves liquid balance and coupling to exchange	Indirect
Mixture momentum	$\rho_m u_m^2 + p$	Predicts pressure-gradient response and transport	Strong
Pressure evolution	$\beta_m^{-1} \partial_t p$	Resolves compressibility-driven dynamics	Moderate
Mass transfer	Γ_{ex}, $c^\star(p,T)$	Separates dissolved and free-gas inventories	Strong
Closure transport	π_k, λ_{jk}	Propagates regime uncertainty over depth and time	Explicit

Table 4. Design objectives for a perception-aware surveillance architecture.

Objective	Interpretation in this study	Deployment consequence
Physical consistency	Regime updates must respect conservation, continuity, and plausible transition rates	Prevents estimator instability
Uncertainty calibration	Posterior confidence should match empirical reliability	Reduces overconfident misdiagnosis
Temporal coherence	Regime evolution should be persistent but transition-capable	Improves sequential interpretability
Domain robustness	Visual model must expose uncertainty under lighting and sensor shift	Supports safe transfer beyond the lab
Computational practicality	Solver and vision stack must run within real-time constraints	Enables live surveillance use
Action relevance	Representation should preserve distinctions that affect control and risk	Keeps outputs operationally meaningful

looks like to how the flow is expected to evolve.

3. Mechanistic Model of Gas Kick Migration

A tractable model for gas-kick migration in a vertical annulus must balance realism with inferential usability. Full interface-resolved simulation is inappropriate for real-time monitoring, while overly reduced algebraic models conceal the dependence of pressure and slip on evolving morphology. The suitable compromise is a one-dimensional transient transport model with regime-conditioned closures. Such a model resolves depthwise advection, pressure evolution, gas fraction changes, and mass transfer while outsourcing unresolved topology to a latent regime field. The latent regime is then inferable from images and, indirectly, from pressure and flow observations [8].

The need for this formulation is reinforced by observations under circulating conditions. Obi et al. (2022b) reported that during drilling or circulation the annular flow structure becomes materially more complex than under shut-in, and that the post-Taylor-bubble behavior is strongly influenced by liquid flow rate, gas injection duration, and gas type, with carbon dioxide showing especially pronounced effects associated with solubility and near-surface release [33]. The modeling implication is that neither a fixed slip law nor a single two-phase friction correction is adequate over the full transient. Closures must adapt to regime and composition, and the estimator must be able to revise those closures as evidence accumulates.

Let the vertical coordinate be $z \in [0, L]$, increasing upward, and let A denote the annular cross-sectional area. The state variables are pressure $p(z,t)$, mixture velocity $u_m(z,t)$, gas volume fraction $\alpha_g(z,t)$, dissolved-gas concentration in the liquid $c_d(z,t)$, and a regime state $r(z,t)$ taking values in a small set of closure families. The liquid volume fraction is $\alpha_l = 1 - \alpha_g$. A low-order formulation can be derived from phase continuity together with a mixture momentum balance and a kinetic relation for dissolution or exsolution. Denoting gas and liquid densities by ρ_g and ρ_l, and gas and liquid superficial velocities by j_g and j_l, one useful state representation is

$$\begin{aligned} \partial_t(\alpha_g \rho_g A) + \partial_z(\alpha_g \rho_g u_g A) &= \Gamma_{ex} A + q_g^{inj}, \\ \partial_t(\alpha_l \rho_l A) + \partial_z(\alpha_l \rho_l u_l A) &= -\Gamma_{ex} A + q_l^{inj}, \\ \partial_t(\rho_m u_m A) + \partial_z(\rho_m u_m^2 A + pA) &= -A \partial_z p_f - \rho_m g A - \tau_w P_w \end{aligned} \tag{1}$$

where $\rho_m = \alpha_g \rho_g + \alpha_l \rho_l$, $u_m = (\alpha_g \rho_g u_g + \alpha_l \rho_l u_l)/\rho_m$, Γ_{ex} is the interphase mass exchange term, q_g^{inj} and q_l^{inj} denote source terms associated with influx or imposed flow, p_f is

Table 5. Visual-learning pipeline for annular flow regime recognition.

Stage	Input/output	Main function	Mode
Spatial encoding	Frames → multiscale features	Preserves interfaces, bubbles, films, and wall-adjacent structure	CNN/ViT
Temporal aggregation	Feature sequence → context vector	Captures persistence, intermittency, and transition direction	RNN/TCN/Attn.
Posterior head	Context vector → $q_\phi(r_t)$	Produces calibrated regime probabilities	Classification
Embedding head	Context vector → e_t	Encodes fine morphology for fusion and anomaly checks	Representation
Calibration layer	Raw scores → reliable probabilities	Aligns confidence with observed correctness	Reliability
OOD monitor	Embedding → s_t^{ood}	Detects domain shift and suppresses unsafe certainty	Safety

Table 6. Training strategy for temporal computer vision in two-phase annular flow.

Component	Purpose in training	Expected benefit
Self-supervised pretraining	Learns generic interface and motion structure before full supervision	Better sample efficiency
Sequence-level supervision	Uses temporal windows instead of isolated frames	Stronger transition discrimination
Soft or boundary-aware labels	Preserves ambiguity near mixed-regime segments	More faithful uncertainty estimates
Calibration loss	Penalizes confidence–accuracy mismatch	Safer posterior assimilation
Domain-adversarial adaptation	Reduces sensitivity to lab-specific appearance bias	Improved transfer robustness
Physics-compatible augmentation	Simulates blur, glare, jitter, and exposure changes without unphysical flow warps	Better deployment resilience

an effective frictional pressure contribution, τ_w is wall shear stress, and P_w is wetted perimeter. The phase velocities are connected to the mixture velocity through a slip relation,

$$u_g = u_m + (1-\alpha_g)u_s(r,p,\alpha_g,\mathrm{Re}_m,\chi), \qquad u_l = u_m - \alpha_g u_s(r,p,\alpha_g,\mathrm{Re}_m,\chi),$$

where u_s is a regime-dependent slip velocity and χ summarizes geometry and fluid-property parameters such as eccentricity, surface tension, and viscosity ratio.

A pressure equation can be obtained by combining compressibility and continuity. For practical estimation, it is often convenient to evolve pressure through a quasi-hyperbolic form in which mixture compressibility β_m^{-1} couples to divergence of volumetric flux:

$$\beta_m^{-1}\partial_t p + \partial_z(\alpha_g u_g + \alpha_l u_l) = s_p,$$
$$\beta_m^{-1} = \alpha_g\rho_g\frac{\partial \rho_g^{-1}}{\partial p} + \alpha_l\rho_l\frac{\partial \rho_l^{-1}}{\partial p} + \frac{\partial \alpha_g}{\partial p}\Big|_{m_g,m_l}. \tag{2}$$

where s_p absorbs source contributions induced by boundary control and local mass transfer. The gas density follows an equation of state, which in many practical conditions may be approximated by a compressibility-corrected real-gas relation [19]. The liquid density may be weakly pressure dependent, while dissolved gas changes the effective liquid properties and must therefore enter ρ_l or the exchange kinetics.

For soluble gas, especially carbon dioxide in water-like drilling fluids or aqueous phases, the free-gas and dissolved-gas inventories must be distinguished. A minimal representation evolves the dissolved concentration through convection and relaxation toward a pressure-dependent equilibrium $c^\star(p,T)$:

$$\partial_t(\alpha_l c_d A) + \partial_z(\alpha_l u_l c_d A) = k_m a_i\big(c^\star(p,T) - c_d\big)A,$$
$$\Gamma_{ex} = -M_g k_m a_i\big(c^\star(p,T) - c_d\big) \tag{3}$$

where k_m is a mass-transfer coefficient, a_i is interfacial area concentration, and M_g converts dissolved concentration flux to gas-phase mass exchange. The interfacial area concentration is itself regime dependent. Bubbly configurations expose a large area per gas volume, whereas elongated structures may exhibit smaller area concentration but stronger local recirculation. Churn-like conditions can increase exchange through interface renewal and breakup. Consequently, a_i is another closure mediated by $r(z,t)$.

The regime field r may be represented as a discrete la-

Table 7. Information sources in the coupled perception–physics estimator.

Source	Observation form	Contribution to posterior update	Frequency
Surface pressure	y_t^{hyd}	Constrains hydrostatic unloading and boundary response	Continuous
Return flow / pit gain	Flow-rate measurements	Informs influx magnitude and transport timing	Continuous
Image posterior	$\hat{q}_t$ at observed depths	Corrects local regime beliefs directly	Intermittent
Visual embedding	e_t	Adds morphology information beyond class probabilities	Intermittent
OOD score	s_t^{ood}	Downweights unreliable visual evidence	Event-driven
Control inputs	u_t	Condition prediction on pumps, choke, and circulation state	Continuous

Table 8. Comparison of inference strategies for gas-kick monitoring.

Strategy	Main strength	Main limitation	Best use case
Hydraulics-only	Enforces physical conservation and pressure dynamics	Regime state remains weakly observable	Sparse-vision settings
Vision-only	Detects local morphology and transition texture directly	Can violate transport plausibility under shift	Lab-centric diagnostics
Ensemble fusion	Scales to medium-dimensional nonlinear states	Approximate treatment of multimodality	Real-time field estimation
Particle-based fusion	Handles alternative influx/regime scenarios explicitly	Higher computational burden	Severe ambiguity cases
Hybrid scenario filter	Balances multimodality and efficiency	Requires careful scenario design	Online decision support

tent variable with classes such as bubbly, elongated-bubble, dispersed-bubble turbulent, churn-like, and mixed transitional. The purpose of this representation is not taxonomic purity but closure management. Each regime class selects families of constitutive laws for slip, friction multiplier, interfacial area, and signal-noise structure. For instance, a bubbly state may use a relatively small slip with weak intermittency, while an elongated-bubble state may use a larger buoyancy-driven slip and reduced effective friction. A churn-like state may exhibit irregular slip and larger pressure fluctuation variance [24]. Mixed transitional states can be represented as convex combinations of neighboring closure families.

To preserve continuity in the latent regime field, one can introduce a probabilistic transition model driven by local dimensionless groups and historical state. Let $\pi_k(z,t) = \mathrm{Pr}[r(z,t)=k]$. Then the regime probabilities evolve according to a controlled Markov transport relation,

$$\partial_t \pi_k + v_r \partial_z \pi_k = \sum_{j\neq k} \lambda_{jk}(\xi)\pi_j - \sum_{j\neq k} \lambda_{kj}(\xi)\pi_k,$$
$$\xi = \left(\alpha_g, u_m, p, \partial_z p, c_d, \mathrm{Re}_m, \mathrm{Fr}, \mathrm{Eo}, \chi\right) \quad (4)$$

where v_r is an effective regime-advection speed and λ_{jk} are transition rates conditioned on local hydraulic features ξ. In practice, these rates are not first-principles constants. They are learned or tuned from experimental data and constrained for plausibility. This probabilistic representation makes the model well suited for fusion with computer vision outputs, which naturally arrive as uncertain class probabilities rather than deterministic labels.

The friction and slip closures are the crucial points where regime awareness enters the hydraulic prediction. A generic regime-conditioned friction law may be written as

$$-\partial_z p_f = f_r(\alpha_g, \mathrm{Re}_m, \chi)\,\frac{\rho_m u_m |u_m|}{2D_h},$$

with D_h the hydraulic diameter and f_r chosen from a class-specific family. Likewise, the slip law may be parameterized as

$$u_s = U_r(\alpha_g, p, \rho_g/\rho_l, \mu_g/\mu_l, \sigma, \chi),$$

where σ is surface tension. Even if these expressions are empirical, the estimation architecture benefits from writing them explicitly because posterior uncertainty in r maps directly into posterior uncertainty in pressure and gas holdup [31].

Boundary conditions distinguish shut-in from circulating transients. Under shut-in, liquid throughflow is effectively zero at the boundaries, and pressure evolution is dominated by hydrostatic redistribution, gas expansion, and possible dissolution. Under circulation, a bottom or inlet liquid flow and a surface outlet condition must be imposed, often together with a control variable such as choke backpressure. The gov-

Table 9. Evaluation criteria for the proposed monitoring framework.

Criterion	What is assessed	Why it matters
Pressure posterior quality	Accuracy and calibration of predicted pressure trajectories	Reflects well-control relevance
Transition detection delay	Time lag in identifying hazardous regime changes	Measures early-warning capability
Gas-inventory estimation	Near-surface free-gas prediction under uncertainty	Connects directly to risk
Calibration error	Agreement between confidence and empirical correctness	Determines trustworthiness of labels
Robustness to degradation	Performance under sparse cameras, glare, delay, or noise	Tests graceful failure behavior
Downstream fusion gain	Improvement over standalone hydraulics or vision systems	Validates the integrated architecture

erning equations then respond not only to the evolving kick but also to operator actions. This is precisely why a regime-aware model is valuable under circulation: the same control action may produce different pressure responses depending on whether the gas is concentrated in an elongated structure or dispersed into a broad turbulent region.

For online use, the continuum model is discretized over depth into N cells. Let x_t denote the stacked state vector containing cellwise pressure, gas fraction, dissolved concentration, mixture velocity, and regime probabilities. A semi-implicit time discretization is advantageous because it stabilizes compressibility without erasing transient evolution. The discrete transition can be written compactly as

$$x_{t+\Delta t} = F_\theta(x_t, u_t) + w_t,$$

where u_t denotes control actions and boundary conditions, θ collects closure parameters, and w_t represents model discrepancy. Surface pressure, return flow, pit gain, and other sensors are then modeled as

$$y_t^{hyd} = H_\theta(x_t) + v_t.$$

The regime-conditioned closures make F_θ piecewise smooth but not globally linear. This motivates probabilistic filtering methods that can handle nonlinearities and multimodal uncertainty [23].

One of the most consequential modeling choices concerns the separation of deterministic dynamics from uncertainty. In ordinary simulation, parameter uncertainty is often treated as an afterthought. In surveillance, it must be built into the state-space description from the beginning. Sources of uncertainty include gas composition, influx depth, local eccentricity effects, friction bias, camera placement, and unmodeled regime mixing. A useful estimator should therefore propagate both state and closure uncertainty. Ensemble methods are particularly suitable because they can carry multiple realizations of the regime field and associated hydraulic trajectories while remaining computationally feasible. Particle methods are more expressive for multimodal posteriors but may be expensive at high state dimension. Hybrid schemes, such as ensemble Kalman filters with latent discrete regime probabilities or Rao-Blackwellized particle filters, provide a promising middle ground.

Another important aspect is the representation of mixed regimes. Real annular flows frequently violate crisp class boundaries, especially during the trailing phase after the first prominent structure passes a given depth. A hard classification scheme can force unrealistic closure jumps. It is therefore preferable to use soft regime probabilities and define effective closures by expectation or mixture:

$$\bar{f} = \sum_k \pi_k f_k, \qquad \bar{u}_s = \sum_k \pi_k U_k, \qquad \bar{a}_i = \sum_k \pi_k a_{i,k}.$$

This yields smoother hydraulic evolution and allows the vision model to express uncertainty naturally. It also reflects operational reality more faithfully: what matters is not whether a frame is purely churn or purely dispersed, but whether the posterior over plausible structures implies a high-risk change in slip, friction, or gas release [9].

The mechanistic model can also incorporate pressure-fluctuation statistics as auxiliary predicted observables. Churn-like or strongly intermittent regimes often induce larger variance and short-time asymmetry in local pressure or flow signals than more ordered bubbly states. If such statistics are available from sensors, they can augment the observation model and help identify regime without direct imaging. In this way, the estimator need not be vision dependent; rather, it is vision enhanced. The framework remains useful even when only sparse or intermittent images are available.

A final modeling concern is identifiability. Gas fraction, slip, and influx size can be partially confounded in pressure-only inference. For instance, a larger gas fraction with smaller slip may produce a pressure response similar to a smaller gas fraction with larger slip over a finite time window. By introducing a regime-informed observation model, one reduces this confounding because the image constrains the plausible slip family. Similarly, dissolved-gas dynamics can confound visible gas inventory. A model that separately tracks free and dissolved gas while using images to constrain visible morphology is therefore substantially more identifiable than one that treats all gas uniformly.

The mechanistic formulation described here is intentionally modular. One may refine the closure laws, include thermal effects, or extend the regime space without altering the basic structure [39]. What matters for the broader thesis of the paper is that hydraulic prediction and regime recognition are placed in the same state-space framework. This sets the stage for the computer vision module, which supplies the probabilis-

tic regime evidence needed to make the latent closure state observable.

4. Computer Vision Models for Two-Phase Flow Regime Identification

The computer vision problem in annular two-phase flow is not a generic image classification task. It is a spatiotemporal inference problem in which local appearance, interface geometry, motion, and persistence all matter, while labels may be ambiguous near transitions and the class boundaries themselves are partly operational rather than purely visual. A robust model must therefore combine high-resolution spatial encoding with temporal context, uncertainty estimation, and a training objective that reflects the ordinal or transitional relation among flow regimes. Such a model is especially valuable when the imagery is later used not as an end in itself but as evidence in hydraulic state estimation.

A relevant experimental basis for this task is provided by Manikonda et al. (2025), who presented a high-resolution vertical annulus image dataset from a 140 ft tower setup designed to capture gas-liquid flow regime behavior and transitions for water-air conditions [27], with an explicit focus on supporting machine vision and computational modeling. The importance of such a dataset within the present framework is not that it settles the regime-recognition problem, but that it enables the construction of visual models whose outputs can be fused with transport predictions. Because the images arise from annular rather than simple pipe geometry, they preserve the morphological cues most relevant to drilling-oriented multiphase diagnostics.

The first modeling decision concerns the target representation. A purely categorical label per frame is too crude. Regime evolution is gradual in many sequences, and frames near transition boundaries may contain mixed evidence. Moreover, some visual distinctions matter mainly because they correlate with closure changes in the hydraulic model. A practical target is therefore a calibrated posterior over a small set of operative regime classes together with a continuous embedding that captures finer morphology [42]. The posterior is used directly by the state estimator, while the embedding can support transition modeling, anomaly detection, and domain adaptation.

The raw imagery should be treated as a sequence rather than an unordered set of frames. Temporal structure matters because many classes are better distinguished by motion and persistence than by static texture alone. Elongated or Taylor-like structures show coherent upward motion and extended interfaces. Bubbly flow exhibits numerous small objects whose trajectories are less coherent. Churn-like patterns are characterized by intermittent large structures, breakup, and disorder. A high-performing model should therefore observe multiple consecutive frames and learn dynamics over a window long enough to capture these behaviors. This suggests a hybrid architecture in which a spatial encoder processes each frame and a temporal module aggregates information across time.

A suitable spatial encoder can be implemented using a convolutional backbone or a vision transformer, provided that interface detail is preserved. Convolutions are advantageous for local edge and texture recognition, particularly when bubbles and films occupy small fractions of the frame. Transformers offer stronger long-range context, which may help when large elongated structures span substantial image extents. In practice, a hybrid encoder is often effective: early convolutional stages capture local gradients and boundaries, while later self-attention layers integrate larger-scale morphology. Because annular imagery may contain fixed geometric boundaries and repeated backgrounds, the model should also be taught to separate static wall cues from dynamic phase structures [12]. This can be encouraged by temporal differencing channels or optical-flow-informed auxiliary tasks.

Let $I_{t-\tau:t}$ denote a sequence of frames. The vision model outputs regime logits $g_\phi(I_{t-\tau:t})$, regime posterior $q_\phi(r_t \mid I_{t-\tau:t})$, and embedding e_t. A generic formulation is

$$e_t = E_\phi(I_{t-\tau:t}), \qquad q_\phi(r_t = k \mid I_{t-\tau:t}) = \mathrm{softmax}_k\big(We_t + b\big).$$

However, this simple expression hides several requirements. First, E_ϕ should be multiscale so that both fine bubbles and large coherent interfaces are visible to the model. Second, the temporal receptive field should be long enough to disambiguate transient clutter from persistent structures. Third, the embedding should support calibration, meaning the posterior probabilities must align with empirical correctness rates. For deployment, a 0.9 posterior should correspond to something near 90% reliability, not simply a sharp score.

Supervised training alone is often insufficient because frame-level labels in flow data can be expensive, noisy, or ambiguous. A more robust approach begins with self-supervised or weakly supervised pretraining. Sequence contrastive learning can encourage embeddings of nearby temporal windows within the same regime segment to be similar, while embeddings across clearly different regimes are separated. Masked-frame reconstruction or temporal ordering tasks can further teach the encoder to represent interface continuity and motion. Once the encoder has acquired these generic visual priors, supervised fine-tuning on labeled regime segments becomes more sample efficient and less prone to overfitting to lighting or camera idiosyncrasies.

The loss function should reflect both classification and temporal structure [18]. One useful composite objective is

$$\begin{aligned}
\mathcal{L} &= \mathcal{L}_{cls} + \lambda_{tc}\mathcal{L}_{temp} + \lambda_{cal}\mathcal{L}_{cal} + \lambda_{dom}\mathcal{L}_{dom}, \\
\mathcal{L}_{cls} &= -\sum_t \sum_k y_{t,k} \log q_\phi(r_t = k \mid I_{t-\tau:t}), \\
\mathcal{L}_{temp} &= \sum_t d\big(q_\phi(r_t), \Pi_\psi(q_\phi(r_{t-1}), \hat{\xi}_t)\big), \\
\mathcal{L}_{cal} &= \sum_b \big|\mathrm{acc}(b) - \mathrm{conf}(b)\big|, \\
\mathcal{L}_{dom} &= \sum_t \ell_{adv}\big(D_\omega(e_t), s_t\big)
\end{aligned} \tag{5}$$

where $\mathcal{L}_{cls}$ is the supervised classification loss, $\mathcal{L}_{temp}$ encourages temporal consistency through a transition predictor Π_ψ, $\mathcal{L}_{cal}$ penalizes miscalibration through binwise accuracy-confidence mismatch, and $\mathcal{L}_{dom}$ is an adversarial domain-invariance term with domain label s_t. The transition predictor may optionally consume approximate hydraulic features

$\hat{\xi}_t$ so that temporal regularization is informed by fluid state rather than being purely visual.

A model trained this way yields more than labels. It learns a latent geometry of the regime manifold. In that geometry, bubbly states may occupy a region connected to dispersed-bubble turbulence, while elongated structures lie closer to certain mixed states and farther from high-disorder churn-like states. This structure is valuable because it allows one to treat classification errors asymmetrically. Confusing a mixed transitional state with churn may be less damaging than confusing a strongly elongated structure with quiescent bubbly flow. Loss shaping or cost-sensitive training can reflect these operational asymmetries without requiring a large number of classes.

Image preprocessing also deserves careful treatment. Naive resizing can destroy small-bubble detail, whereas aggressive contrast enhancement can introduce artifacts mistaken for interfaces. A robust pipeline should preserve native aspect ratio where possible, normalize illumination carefully, and include augmentation strategies that mimic likely deployment shifts, such as modest blur, noise, exposure variation, and reflections. Importantly, augmentations must not create physically impossible structures. Random warps that distort gravity-aligned trajectories can harm the learned representation more than they help. Physics-compatible augmentation is therefore preferable: temporal speed changes, realistic brightness variation, local occlusion, and camera jitter are reasonable; arbitrary geometric flips are usually not in vertical flow imagery because gravity provides a privileged direction.

Annotation strategy matters as much as model architecture [28]. Regime labels in two-phase flow are partly interpretive, and disagreement near transitions is unavoidable. Instead of forcing a single hard label in all cases, one can allow soft labels or confidence-weighted labels for ambiguous segments. This is compatible with the probabilistic regime field used in the hydraulic estimator. In fact, the estimator benefits when the vision module remains uncertain on mixed frames rather than overfitting to annotation conventions. Another useful tactic is segment-level labeling: human annotators identify intervals with dominant regime behavior and mark uncertain boundaries, after which the model learns with boundary-aware losses. This approach better reflects the temporal nature of the problem.

A central difficulty for deployment is domain shift between laboratory images and field conditions. The field may involve opaque fluids, poorer optics, different scales, partial views, or alternative sensing modalities such as limited borescope windows or fiber-optic surrogate images. A practical research agenda should therefore not assume direct transfer. Instead, the model should be designed for staged adaptation. The first stage learns morphology-rich embeddings from the transparent annulus dataset. The second stage performs domain adaptation using unlabeled target images or related synthetic renderings. The third stage calibrates uncertainty so that the model knows when it is outside its competence [34]. Techniques such as adversarial alignment, style transfer with content preservation, and feature-space regularization can help, but none replaces the need for explicit abstention logic.

For this reason, out-of-distribution detection should be treated as a first-class requirement. The vision model should report when a frame sequence does not resemble the training support, for example because of severe glare, new fluid properties, or a regime absent from the source dataset. Energy-based scores, deep ensembles, and distance-to-prototype metrics in embedding space can each supply such warnings. In the coupled estimator, an out-of-distribution signal does not cause failure; it simply lowers the weight placed on visual evidence and shifts reliance back toward mechanistic prediction and conventional sensors.

A perception-aware flow model can also benefit from auxiliary visual tasks beyond regime classification. Interface segmentation, bubble counting, interfacial area proxy estimation, and motion field estimation each provide information relevant to hydraulic closures. Even when the final estimator only consumes regime posteriors and embeddings, multi-task learning often improves the backbone by forcing it to attend to physically meaningful structures. For instance, a segmentation auxiliary loss can sharpen sensitivity to bubble boundaries, while a temporal motion loss can improve discrimination between persistent elongated structures and transient clutter. These tasks can be implemented without turning the system into a dense prediction pipeline at deployment time; their main value lies in representation learning.

It is also useful to connect the visual embedding directly to closure surrogates. Rather than predicting only a discrete regime class, the network can emit auxiliary estimates of effective slip indicator, interfacial complexity, or expected pressure-fluctuation variance. These quantities need not be treated as literal measurements. They can function as informative priors for the mechanistic model [11]. Such an approach blurs the line between classification and parameter inference in a productive way. The network remains supervised by regime labels, but its embedding is encouraged to organize itself around hydraulically relevant latent factors.

Temporal aggregation can be implemented with recurrent units, temporal convolutions, or transformer-style attention over frame sequences. Recurrent units are computationally economical and support streaming inference, which is useful for real-time monitoring. Temporal convolutions offer stable training and controlled receptive fields. Transformers can capture long-range dependencies and irregular intervals but may be heavier. A deployment-minded architecture might therefore combine a lightweight convolutional encoder with a gated recurrent temporal head. More computationally expensive teacher models can be used offline during training to distill richer temporal structure into the streaming model.

Evaluation protocols should go well beyond overall accuracy. Because the end use is state estimation and control support, the relevant metrics include calibration error, transition detection delay, confusion between hydraulically distinct classes, robustness across elevations, and resilience to synthetic domain perturbations. It is also valuable to evaluate whether the learned posteriors improve downstream hydraulic inference. A model that is slightly worse on standalone image accuracy but much better calibrated may be more valuable in the coupled system than a sharper but overconfident classifier. Thus, regime recognition should ultimately be judged by its contribution to state observability and risk estimation, not only by benchmark classification scores [22].

The computer vision module developed in this manner transforms images into physically interpretable evidence. It does not replace multiphase transport theory; it supplies the missing observational link between hidden regime structure and

closure selection. Once calibrated, uncertainty-aware regime posteriors and morphology embeddings are available, they can be assimilated into the mechanistic model through probabilistic filtering. That fusion is the next step in constructing a genuinely perception-aware gas-kick monitoring architecture.

5. Perception-Physics State Estimation and Control

Once the hydraulic state model and the visual regime-recognition model are available, the central challenge becomes fusion. The goal is not simply to average predictions. Rather, the estimator must determine how much trust to place in each information source at each time, given uncertainty, possible domain shift, and the dynamical constraints of multiphase transport. This is naturally formulated as Bayesian filtering or smoothing in a latent state-space model where the hydraulic variables and regime probabilities evolve jointly and observations arrive from both conventional sensors and the vision module.

Let x_t denote the hydraulic state and π_t the cellwise regime probability field. The mechanistic model defines a predictive transition

$$(x_{t+1}, \pi_{t+1}) \sim p_\theta(x_{t+1}, \pi_{t+1} \mid x_t, \pi_t, u_t).$$

Hydraulic sensors produce observations y_t^{hyd}, while the vision model produces either direct regime posteriors $\hat{q}_t$ at observed elevations or image embeddings e_t. A simple and effective choice is to treat the image posterior as a noisy observation of the latent regime probabilities. For an observed depth cell i, [29]

$$\hat{q}_{t,i} = \pi_{t,i} + \epsilon_{t,i},$$

with $\epsilon_{t,i}$ drawn from a calibrated distribution whose variance depends on model confidence and out-of-distribution score. In practice, the update may be implemented through a Dirichlet observation model or logistic-normal approximation rather than a Euclidean Gaussian, but the conceptual role is the same.

If the image model also emits embeddings e_t, these can be connected to hydraulic states through learned likelihood surrogates. For example, a decoder can map x_t, π_t to an expected embedding $\bar{e}_t$, trained on synchronized image-hydraulics data. Then one can write

$$e_t \sim \mathcal{N}(\bar{e}_t(x_t, \pi_t), \Sigma_e),$$

which allows the filter to use richer visual information than class probabilities alone. This is useful when regime boundaries are soft or when morphology contains signal about intensity within a regime, such as increased disorder or bubble crowding that may precede a class transition.

Several filtering strategies are compatible with this architecture. An ensemble Kalman filter is attractive because it scales reasonably to medium-dimensional hydraulic states and can assimilate nonlinear observations approximately. Each ensemble member carries a hydraulic trajectory and a regime realization or soft regime field. During prediction, members evolve according to the mechanistic model with perturbed closure parameters and source terms. During update, hydraulic sensor residuals and visual residuals jointly correct the ensemble. Because the regime field may be discrete or simplex-valued, one often updates unconstrained logits or latent continuous variables and then maps back to probabilities. This keeps the assimilation numerically convenient.

Particle filtering is more flexible when the posterior is strongly multimodal, as may happen when pressure-only data admit multiple plausible influx scenarios [7]. In such cases visual evidence can decisively reweight particles corresponding to different regime and slip histories. However, naive particle filtering suffers from degeneracy in high dimensions. A hybrid approach is often preferable: sample a moderate number of global scenario particles for uncertain variables such as influx size, composition, or closure family, while using an ensemble or Gaussian approximation for the continuous within-scenario state. This captures major multimodality without excessive cost.

An important practical issue is asynchronous and sparse visual sensing. Images may only arrive intermittently or at a subset of depths, whereas pressure and flow sensors may be continuous. The estimator should therefore treat visual observations as local corrections whose influence propagates through the model. If a camera at an intermediate elevation observes a transition toward churn-like behavior, the filter updates the local regime posterior there and then propagates the implied changes in slip and friction upward and downward according to the transport dynamics. This can sharpen predictions at depths not directly observed. Such propagation is one of the chief benefits of state-space coupling over standalone computer vision.

Weighting of visual evidence must be adaptive. When the image model is well calibrated and in distribution, it should exert substantial influence on regime state. When confidence is low or out-of-distribution risk is high, the update should soften [10]. This can be handled by inflating the visual observation covariance or by tempering the image likelihood. Formally, if s_t^{ood} is an out-of-distribution score, one can define an effective covariance

$$\Sigma_{q,t}^{eff} = \Sigma_{q,t}\big(1 + \kappa s_t^{ood}\big),$$

so that anomalous visual conditions reduce assimilation strength. The same logic applies when hydraulic conditions predict transitions absent from the source image dataset. The system should not force an overconfident visual label onto an implausible operating point; instead, it should widen uncertainty.

A coupled estimator also enables more informative risk metrics than pressure thresholds alone. Let R_t denote a risk functional depending on predicted free-gas inventory near the surface, probability of severe regime transition, expected pressure excursion, and estimator uncertainty. One can define

$$R_t = \mathbb{E}\big[\varphi(x_t, \pi_t) \mid \mathcal{D}_{1:t}\big] + \eta\, \mathrm{Var}\big[\varphi(x_t, \pi_t) \mid \mathcal{D}_{1:t}\big],$$

where $\mathcal{D}_{1:t}$ includes all data and φ maps state to a hazard-relevant scalar. The variance term is not cosmetic. In well control, uncertainty itself is a risk amplifier because it reduces confidence in operational margins. A perception-aware system can therefore flag elevated risk not only when the predicted mean pressure excursion is large, but also when model disagreement about regime evolution becomes pronounced.

Control support can be built on top of this risk-aware state

estimate. During circulation or managed pressure operations, the operator or supervisory controller chooses variables such as choke setting, pump schedule, or circulation rate. These actions influence both transport and observability. For example, increasing liquid throughput may improve transport of dispersed gas but also alter turbulence and visual ambiguity [4]. In a model predictive control formulation, the system can optimize future control actions against a cost that balances pressure safety, influx removal, and uncertainty reduction:

$$J = \sum_{\tau=t}^{t+H} \mathbb{E}\Big[\ell_p(p_\tau) + \ell_q(Q_\tau^{out}) + \ell_r(R_\tau) + \ell_u(u_\tau)\Big],$$
$$u_{t:t+H}^{\star} = \arg\min_{u_{t:t+H}} J \tag{6}$$

subject to the perception-aware state dynamics and operational constraints. Although a full control design is beyond the scope of the present paper, the state-estimation architecture already provides the ingredients needed for such an extension.

A related advantage of fused inference is interpretability. Purely data-driven alarms can be difficult to trust because they do not explain whether a warning arises from gas expansion, dissolution reversal, or a change in local morphology. Purely mechanistic forecasts can be difficult to trust when the closure state is opaque. In the coupled framework, an alarm can be decomposed into physically meaningful causes: rising probability of a high-slip regime at a given depth, increased expected near-surface exsolution, or growing divergence between image evidence and pressure-based prediction. Such decomposition is valuable for expert oversight and for post-event analysis.

The fusion architecture also provides a principled way to use laboratory visual data in field systems where direct imaging may be absent. During development, the image model helps identify and calibrate regime-conditioned closures in the mechanistic twin. Once transferred to the field, even sparse surrogate observations or pressure-fluctuation statistics can be interpreted through those better-calibrated closures. In other words, the role of computer vision is not limited to online classification in transparent test sections. It also strengthens the hydraulic model itself by making regime structure more measurable during the model-building phase [35].

From an implementation standpoint, differentiable programming offers additional possibilities. If the mechanistic solver is differentiable or approximately differentiable, one can train parts of the perception and fusion stack end-to-end using downstream estimation loss. For example, the visual embedding can be optimized not only for image classification accuracy but also for its utility in reducing state estimation error. This is attractive because it aligns the vision model with the actual operational objective. However, end-to-end training should be used carefully. Hard physical constraints and explicit calibration targets should remain in place so that the model does not discover visually convenient but hydraulically misleading shortcuts.

Finally, the coupled estimator should be designed for graceful degradation. Safety-critical systems cannot assume all sensors remain available. If images are lost, the system should revert to hydraulics-only inference with widened uncertainty. If pressure sensors are noisy, visual evidence can partially stabilize regime estimation. If both are weak, the system should state its uncertainty rather than offering brittle certainty. This philosophy of graceful degradation is one of the strongest arguments for perception-aware monitoring: by combining independent evidence channels under a coherent probabilistic framework, the system becomes more robust than either channel alone.

6. Experimentalization, Evaluation, and Deployment Considerations

The value of a perception-aware framework depends on whether it can be validated in a way that respects both multiphase physics and machine-learning realism. Validation should therefore not be reduced to matching one pressure trace or achieving one classification score [6]. Instead, the full system must be tested as an inference engine across controlled transients that vary gas type, injection history, circulation state, and observational quality. The experimental program needed for such validation can be built naturally around vertical annulus facilities with synchronized pressure, flow, and image acquisition.

A rigorous protocol begins with synchronized multimodal data collection. Images should be time-aligned with pressure measurements, injection records, and liquid flow conditions. The alignment tolerance must be small enough to preserve transient causality because regime transitions may appear quickly once the gas front reaches a camera elevation. Multiple elevations are valuable, not simply to increase data volume, but to reveal how the same transient changes morphology as it rises. This depth-resolved view is critical for learning transition models and for validating whether the hydraulic estimator propagates local visual information in a physically reasonable way.

The experimental matrix should include both shut-in and circulating cases. Shut-in data emphasize buoyancy, compressibility, and dissolution effects without imposed liquid momentum. Circulating data reveal how friction, background throughput, and injected liquid momentum alter the transition path and the observability of gas structures. A well-designed matrix also varies gas composition, especially where solubility can materially affect visible gas inventory and later release. Such variation is not merely a robustness exercise. It tests whether the model can separate appearance changes due to true regime transition from appearance changes induced by compositional mass transfer [43].

Evaluation of the mechanistic component should examine not only pressure prediction error but also inferred gas-fraction profiles, timing of front arrival, and posterior uncertainty. The correct standard is not deterministic point prediction, because the problem is structurally uncertain. A superior model is one that places the true trajectory inside a narrower calibrated posterior, not necessarily one that produces the single smallest root-mean-square pressure error on every run. Probabilistic scoring rules, such as interval coverage and negative log-likelihood on held-out runs, are therefore more informative than point metrics alone.

Evaluation of the vision component should similarly go beyond average classification accuracy. Calibration error, expected transition detection lag, uncertainty under domain

perturbation, and performance on ambiguous mixed regimes are essential. Sequence-level metrics are more meaningful than frame-level metrics because the operational use of the model depends on continuous inference. It is also valuable to stratify performance by elevation and by gas inventory severity. A model that performs well in low-gas bubbly cases but fails in high-severity mixed cases may still look strong on aggregate benchmarks if the data are imbalanced.

The most important evaluation, however, is downstream. One should compare three systems: hydraulics-only estimation, vision-only regime labeling, and fused perception-aware estimation. The comparison should focus on metrics that matter operationally, such as early warning of hazardous transition, uncertainty-calibrated near-surface gas inventory prediction, and robustness when one sensor modality is degraded. It would not be surprising if the fused system were only modestly better than the best standalone system on easy cases. Its principal value should appear on hard cases where observational ambiguity is highest [37]. Those are precisely the cases in which regime-aware fusion should reduce confounding between influx size, slip, and pressure response.

Experimental validation should also examine sensitivity to label noise and camera placement. In many facilities, one cannot instrument every desired depth, and some placements may see only a subset of the relevant morphology. The estimator should therefore be stress-tested under reduced camera coverage, delayed image availability, and partially occluded views. Because the fused system propagates local evidence through a transport model, it should degrade more gracefully than a standalone classifier. Demonstrating this property is important for practical credibility.

Another crucial issue is scale transfer. Laboratory annuli are invaluable because they provide transparent access and controlled transients, but field systems differ in length scale, fluid complexity, pressure level, and sensor quality. Transfer should therefore proceed through a hierarchy rather than a direct leap. First, the mechanistic closures are calibrated using the laboratory data with visual supervision. Second, synthetic data or intermediate-scale facilities can be used to adapt the visual encoder to new optics and flow characteristics. Third, in field settings the estimator can begin in a hydraulics-dominant mode, using visual cues only where confidence is high. This staged transfer acknowledges that the laboratory and the field are not identical while still allowing the laboratory to improve field observability [41].

Synthetic data can aid this process if used carefully. Computational fluid dynamics or interface-rendering tools can generate additional sequences with controlled variation in lighting, geometry, and regime severity. Yet synthetic images are only helpful if their morphology is credible. Low-fidelity synthetic bubbles with unrealistic breakup statistics may teach harmful shortcuts. The most promising use of synthetic data is therefore targeted augmentation: generating difficult or underrepresented transition patterns, or approximating optical artifacts expected in deployment, while anchoring the model primarily in real experimental imagery. Domain randomization can help the visual encoder learn invariances, but it should be constrained by fluid-mechanical plausibility.

Deployment considerations extend to computational architecture. In many drilling environments, edge hardware is limited and communication delays matter. A practical implementation may therefore split computation across tiers. A lightweight streaming vision model runs at the edge near the camera, producing embeddings, confidence scores, and provisional regime posteriors. A central estimator then fuses these outputs with hydraulic measurements and mechanistic prediction. This division preserves low-latency visual processing while allowing the more expensive state estimation to run on a more capable processor. It also aligns with data governance concerns, since raw video may be more expensive to store or transmit than compressed embeddings and posterior summaries.

Human factors should not be ignored [30]. Operators are more likely to trust a system that exposes its uncertainty and causal reasoning than one that simply announces a regime label. Interface design should therefore show evolving posterior bands for gas fraction and pressure, along with regime probabilities at observed depths and the confidence of the visual model. When the image stream is judged out of distribution, that status should be explicit. The operator should see not only the current best estimate but also why uncertainty has widened or narrowed. Such transparency is particularly important when perception and pressure evidence disagree.

Failure analysis is a central part of deployment readiness. Several failure modes are foreseeable. The vision model may mistake reflections or contamination for interfaces. The mechanistic model may underrepresent eccentricity effects or non-Newtonian rheology. Synchronization errors may corrupt the relation between images and pressure transients. Domain shift may invalidate calibration. The appropriate design response is not to expect zero failures but to identify them early and ensure that the fused system fails conservatively. Conservative failure means uncertainty grows, reliance shifts toward more trustworthy modalities, and alarms reflect ambiguity rather than false certainty [40].

There is also a methodological opportunity in using the estimator itself to improve data collection. If posterior uncertainty remains persistently high at certain depths or during certain transients, that indicates where additional cameras, annotations, or physical experiments would be most informative. Active experimental design can therefore be driven by the coupled model. Instead of collecting more of the same easy bubbly-flow images, one can target the regime boundaries and dissolved-gas conditions that currently dominate estimation error. This closes the loop between modeling and experimentation.

From a research standpoint, the integrated framework invites several extensions. Non-Newtonian drilling fluids could be incorporated through rheology-dependent closures and retrained visual embeddings. Multispectral or depth-aware imaging could enrich interface perception. Audio or vibration sensing could supplement visual cues in opaque systems. Differentiable simulators could tighten end-to-end training. Yet the basic experimental lesson remains stable: to make annular gas-kick monitoring more reliable, one must measure and model regime evolution as a dynamic latent state rather than treat it as a qualitative side observation.

The deployment path is therefore best understood as incremental but cumulative. First, laboratory image data improve visual regime recognition [16]. Second, visual regime recognition improves closure identification in mechanistic kick models. Third, those improved models support better state estimation and alarm logic even when field imagery is sparse. Fourth, wherever field imaging becomes available, the same

probabilistic fusion architecture can incorporate it directly. At each step, the technical core is unchanged: mechanistic transport and perception are coupled through an uncertainty-aware estimator.

7. Discussion

The integrated framework developed in this paper is motivated by a simple but consequential observation: in transient annular two-phase flow, what the flow looks like is inseparable from how the flow should be modeled. This is particularly true during gas-kick migration, where pressure behavior, slip, and hydrostatic unloading are mediated by the phase structure that forms in the annulus. Traditional workflows often acknowledge this qualitatively while still operationalizing the problem as if regime could be neglected or fixed. The present discussion returns to the broader implications of treating regime as a latent state in a perception-aware dynamical system.

A first implication concerns the meaning of model fidelity. In many engineering settings fidelity is equated with increasing the complexity of the governing equations. Yet for partially observed multiphase systems, fidelity also depends on whether the model′ s hidden variables are connected to measurable evidence. A closure-rich solver that cannot infer its closure state from observations may be less useful operationally than a somewhat simpler solver whose regime field is constrained by images and calibrated uncertainty. In this sense, perception-awareness is not an optional overlay but a route to effective model fidelity under observational limitation.

A second implication concerns the role of laboratory experiments [15]. Transparent annulus experiments are sometimes viewed as too idealized relative to field wells, and it is true that direct transfer requires caution. However, their value extends beyond direct scaling. They provide a setting in which hidden closure states become observable, making it possible to identify which visual structures correspond to meaningful changes in slip, friction, and dissolution behavior. This knowledge can be transferred into lower-observability environments through better-calibrated mechanistic models and probabilistic priors. The laboratory thus serves not merely as a benchmark arena but as an observability amplifier for model construction.

A third implication is methodological. The framework suggests that flow-regime datasets should be curated with downstream inference in mind. It is not enough to collect visually diverse images. The data should be synchronized with hydraulic variables, annotated with transition uncertainty, and structured so that models can learn not just class boundaries but the relationship between appearance and transport state. This favors multimodal datasets in which images, pressure, flow, and control actions are recorded together. Such datasets enable representation learning that is meaningful for estimation and control, rather than only for image recognition.

There is also a conceptual shift in how one interprets uncertainty. In many machine-learning applications uncertainty is treated as a confidence score attached to a prediction [36]. Here uncertainty is a structural quantity that propagates into risk and control. If the vision model is uncertain between a dispersed-bubble and a churn-like state, that uncertainty changes the predicted friction and slip, which changes the pressure forecast and possibly the recommended control action. The estimator must therefore preserve uncertainty through all layers of the architecture. Overconfident vision or overconfident hydraulics are equally problematic because both collapse ambiguity that is operationally real.

The framework also highlights the importance of mixed regimes and transition states. Engineering taxonomies often prefer clean class labels because they are convenient for correlation development. Yet the most diagnostically important moments may occur exactly when the flow does not fit a pure category. The trailing region after a prominent elongated structure, the onset of turbulence under circulation, or the near-surface release of previously dissolved gas can all produce mixed morphology. Soft probabilistic regime representations are therefore not merely mathematically elegant; they better reflect the operational reality of transitional multiphase flow.

From the viewpoint of machine vision, the application presents a distinctive challenge. Many high-performing visual models are rewarded for exploiting any predictive cue, including background artifacts. In annular flow imagery, such shortcuts are especially dangerous because experimental setups often contain fixed lighting patterns, wall features, or camera-specific artifacts that correlate with labels accidentally. A perception-aware pipeline counteracts this tendency by forcing the model to be useful for hydraulic estimation, not just static classification. If a visual feature improves frame-level accuracy but does not help predict plausible transport evolution, its value should be questioned [2]. This is one reason why downstream evaluation within the estimator is indispensable.

The discussion also reveals why direct end-to-end black-box prediction of well-control outcomes from images is not the right target. It might be possible, in a narrow experimental setting, to train a deep network that maps image sequences and a few sensor values directly to an alarm or a pressure forecast. But such a model would sacrifice physical interpretability, make uncertainty harder to trust, and likely fail under domain shift. By preserving a mechanistic transport core and using vision to infer the hidden regime state, the present framework keeps the causal structure intelligible. The image informs the regime; the regime selects closures; the closures shape hydraulic evolution. This layered causality is scientifically and operationally preferable to opaque shortcut learning.

Another point worth emphasizing is that the proposed approach is compatible with partial adoption. One does not need perfect computer vision to benefit from perception-aware modeling. Even intermittent visual evidence at a few elevations can improve closure identification and reduce uncertainty. Likewise, one does not need a perfect mechanistic model to benefit from fusion. The fusion process itself can reveal where the model is systematically biased, for example if visual evidence repeatedly suggests a transition earlier than the hydraulic model predicts. Such discrepancies can guide model refinement [5]. The architecture is therefore useful not only as a final monitoring solution but also as a scientific instrument for discovering where current transport assumptions are inadequate.

The framework further suggests a different way to think about digital twins in drilling. A digital twin is often presented as a high-fidelity simulation continuously updated by data. In multiphase drilling surveillance, that description is incomplete unless the twin includes latent structural states that connect measurements to closures. A twin that tracks pressure and flow but not regime will inevitably struggle with identifiability in transients. By contrast, a perception-aware twin includes the hidden morphological field, allows images or surrogate observations to update it, and represents uncertainty in the closure state. Such a twin is not necessarily more complicated in every equation, but it is more complete in epistemic structure.

One should also be candid about limitations. The proposed regime field is still a reduced representation of extremely rich flow topology. One-dimensional transport models cannot reproduce all azimuthal or local interface effects. Visual domain shift from transparent laboratory water-air systems to real drilling conditions may be severe. Solubility behavior, non-Newtonian rheology, and field instrumentation constraints add further complexity. The appropriate conclusion is not that the approach is overambitious, but that it should be advanced incrementally with careful calibration and conservative uncertainty handling. The architecture is strongest when it openly represents what it does not know [20].

Finally, the discussion points toward a broader research agenda in which fluid mechanics and machine perception are not separate disciplines exchanging outputs, but components of a unified inferential science. In systems where the most important state variables are hidden yet leave visible structural traces, the combination of mechanistic transport and learned perception is not a temporary trend. It is a necessary response to the limits of scalar sensing. Gas-kick migration in annuli is a particularly compelling example because its safety implications are immediate and its hidden structure is so central to its dynamics. The perception-aware framework developed here aims to place that insight on a technically operational footing.

8. Conclusion

This paper has argued that gas-kick migration in vertical annuli should be modeled and monitored as a perception-aware multiphase inference problem. The central premise is that flow regime is a latent state with operational significance because it mediates between observed morphology, closure selection, pressure evolution, and risk. Treating regime as a descriptive label added after simulation is therefore inadequate for modern monitoring architectures. Instead, regime should be represented explicitly within a coupled state-space framework in which mechanistic transport, visual evidence, and uncertainty are updated together.

To support that premise, the paper developed a unified technical framework with three components. The first component is a regime-conditioned mechanistic model for transient gas migration under shut-in and circulating conditions. It incorporates gas fraction transport, mixture momentum, pressure evolution, and dissolved-gas exchange, while allowing slip, friction, and interfacial area to depend on a probabilistic regime field. The second component is a computer vision pipeline for annular two-phase regime identification from image sequences [13]. That pipeline is designed around temporal representation learning, calibration, soft labeling, and domain robustness rather than static frame classification alone. The third component is a probabilistic fusion architecture in which image-derived regime evidence updates the hydraulic state estimate through filtering and uncertainty-aware assimilation.

The broader conclusion is that computer vision for two-phase flow becomes most valuable when it is not treated as an isolated benchmark task. Its real utility lies in making hidden closure states more observable, reducing confounding in hydraulic inference, and improving the reliability of risk indicators during transient gas migration. Likewise, mechanistic models become more operationally credible when their latent regime assumptions are exposed to visual evidence rather than remaining implicit. The resulting digital twin is not purely data-driven and not purely first-principles; it is a coupled estimator in which physical laws and perceptual measurements constrain one another continuously.

Several practical implications follow. Experimental programs should collect synchronized image and hydraulic data with transition-aware labels. Validation should focus on downstream state estimation and uncertainty calibration, not only on isolated classification metrics. Deployment should emphasize graceful degradation, out-of-distribution awareness, and staged transfer from laboratory facilities to field systems. Future extensions may incorporate richer rheology, additional sensing modalities, and more expressive differentiable solvers, but the essential architecture need not change.

Reliable surveillance of annular gas-kick transients requires acknowledging that the flow′ s hidden structure matters as much as its scalar measurements. A perception-aware framework provides a disciplined way to turn that acknowledgment into a computational system. By coupling multiphase transport modeling with computer vision based regime identification and probabilistic state estimation, one can improve observability, calibrate uncertainty more honestly, and support safer interpretation of gas-kick evolution in complex vertical annular flow [25].

References

[1] Review of applicable technology: solution mining of caverns in salt domes to serve as repositories for radioactive wastes, 6 1976.

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

[3] Development of a natural gas systems analysis model (gsam). annual report, january 1994–january 1995, 7 1994.

[4] Khalid S. Al-Mohanna, Suresh Jacob, Larry Ma, and Muhammad 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.

[5] J. Daniel 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.

[6] C.H. Atkinson. Western gas sands project status report, 10 1978.

[7] J.C. Austin. Direct utilization of geothermal energy resources in food processing. final report, may 17, 1978-may 31, 1982, 5 1982.

[8] F.C. Benedict, R.E. Criss, M.L. Davisson, G.F. Eaton, G.B. Hudson, J.M. Kenneally, T.P. Rose, and D.K. Smith. Hydrologic resources management program, fy 1998 progress report, 7 1999.

[9] W.J. Clark. Determination of importance evaluation for the subsurface exporatory studies facility, 6 1999.

[10] Ganesha R Darmawan and Aries Prasetyo. Drilling the undrillable; a review of indonesia onshore managed pressure drilling (mpd) operation experiences. *PETRO:Jurnal Ilmiah Teknik Perminyakan*, 10(4):211–217, 1 2022.

[11] Dennis Denney. Efficient conceptual design of an offshore gas-gathering network. *Journal of Petroleum Technology*, 60(12):91–93, 12 2008.

[12] Annaguly Deryaev. Economic justification of the need to drill wells with separate operation of several horizons at the same time. *International Science Journal of Management, Economics & Finance*, 1(3):127–132, 8 2022.

[13] K.B. Dollens, K.J. Harpole, E.G. Durrett, and J.S. Bles. Design and implementation of a cosub 2 flood utilizing advanced reservoir characterization and horizontal injection wells in a shallow shelf carbonate approaching waterflood depletion. annual report, july 1, 1996–june 30, 1997, 12 1997.

[14] B.. Dow, C.. Kamps, J.. Baker, S.. Likki, C.. Guzman, B.. Liezenberg, and J.. Guidry. Design and manufacture of an original equipment manufacturer deepwater managed pressure drilling integrated solution. In *IADC/SPE Drilling Conference and Exhibition*. SPE, 3 2018.

[15] Francisco Orlando Garzon, Jose Ricardo Amorocho, Moataz Al-Harbi, Nayef S. Al-Shammari, Azmi Al-Ruwaished, Mohammed Ayub, Wassim Kharrat, Vsevolod Bugrov, Jul Rune Gro Jacobsen, George Albert Brown, and Jose Vidal Noya. Stimulating khuff gas wells with smart fluid placement. In *SPE Deep Gas Conference and Exhibition*. SPE, 1 2010.

[16] Chinthaka P. Gooneratne, Bodong Li, and Moellendick Timothy E. Downhole applications of magnetic sensors. *Sensors (Basel, Switzerland)*, 17(10):2384–, 10 2017.

[17] Jon C. Helton, D.R. Anderson, and B.L. Baker. Computational implementation of a systems prioritization methodology for the waste isolation pilot plant: A preliminary example, 4 1996.

[18] Richard P. Holst. Preparing for lng: A montney gas field optimization workflow. In *SPE/CSUR Unconventional Resources Conference*. SPE, 10 2015.

[19] Natarajan Jeyaprakash, Che-Hua Yang, and Manickam Bhuvanesh Kumar. *Laser Machining*. IntechOpen, 7 2021.

[20] C.J. John, G. Maciasz, and B.J. Harder. Gulf coast geopressured-geothermal program summary report compilation. volume 2-a: Resource description, program history, wells tested, university and company based research, site restoration, 6 1998.

[21] G.D. Johnson. Program plan for evaluation and remediation of the generation and release of flammable gases in hanford site waste tanks, 8 1991.

[22] Ahmed Kamel. A technical review of radial jet drilling. *Journal of Petroleum and Gas Engineering*, 8(8):79–89, 10 2017.

[23] Maximillian J. Kieba and Christopher J. Ziolkowski. Differential soil impedance obstacle detection, 6 2004.

[24] Xiangwei Kong, Yuanhua Lin, Yijie Qiu, Hongjun Zhu, Long Dong, and Yanchao Chen. A new model for predicting dynamic surge pressure in gas and drilling mud two-phase flow during tripping operations. *Mathematical Problems in Engineering*, 2014(1):1–16, 1 2014.

[25] P.J. Lienau and H. Ross. Low-temperature resource assessment program. final report, 2 1996.

[26] Howard Steven Littell, James Webster Jessup, Wade Schoppa, M.R. Seay, and T.D. Coulon. Perdido startup: Flow assurance and subsea artificial lift performance. In *Offshore Technology Conference*. OTC, 5 2011.

[27] Kaushik Manikonda, Chinemerem Obi, Aarya Abhay Brahmane, Mohammad Azizur Rahman, and Abu Rashid Hasan. Vertical two-phase flow regimes in an annulus image dataset-texas a&m university. *Data in Brief*, 58:111245, 2025.

[28] Daniela Marum, Ansgar Cartellieri, Edisa Shahini, and Donata Scanavino. A novel and sustainable generation of advanced mud gas logging system for managed pressure drilling applications: An explorative well deployment from north germany. In *SPE Middle East Oil & Gas Show and Conference*. SPE, 12 2021.

[29] R.V. Matalucci, C. Esparza-Baca, and R.D. Jimenez. Characterization, monitoring, and sensor technology catalogue, 12 1995.

[30] Earl D. Mattson and Larry Hull. Documentation of inl′ s in situ oil shale retorting water usage system dynamics model, 12 2012.

[31] Talal Mokhalalati, Abdullah Al-Suwaidi, and Abdul El-Fatah Hendi. Managing onshore drilling wastes - abu dhabi experience. In *Abu Dhabi International Petroleum Exhibition and Conference*. SPE, 10 2000.

[32] Chinemerem E Obi, A Rashid Hasan, Luis Abril, Kaushik Manikonda, and Azizur Rahman. Experimental investigation of gas kick behavior during shut-in condition. In *ASTFE Digital Library*. Begel House Inc., 2022.

[33] Chinemerem E Obi, Kaushik Manikonda, Luis Abril, A Rashid Hasan, and Mohammad A Rahman. Dynamics of gas kick migration in the annulus while drilling/circulating. In *International Conference on Offshore Mechanics and Arctic Engineering*, volume 85956, page V010T11A089. American Society of Mechanical Engineers, 2022.

[34] Chuming Pang, Yongkui Shi, Xin Wang, Benzheng Li, and Hengjie Luan. Study on a high-efficiency mining technology system for gas outburst in coal seams—example of an h coal mine. *Minerals*, 12(7):795–795, 6 2022.

[35] Bhavin Patel, Todd Douglas Cooper, and Will Billings. The application of advanced gas extraction and analysis system complements early kick detection & control capabilities of managed pressure drilling system with added hse value. *SPE/IADC Drilling Conference*, 3 2013.

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

[37] Mehdi Mir Rajabi, Bjorn Rohde, Niall Maguire, Roger Sverre Stave, and Rolf Tapper. Successful implementations of top-hole managed pressure cementing and managed pressure drilling in the caspian sea. In *SPE/IADC Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*. SPE, 3 2012.

[38] Freddy Rojas Rodriguez, Andi Eka Prasetia, Adel Mettai, Lawrence Umar, Yap Yun Thiam, and M. Syazwan Murad. Hp/ht exploration well in offshore malaysia pushed automated mpd system to maximum utilization, identifying safest drilling operating window. *IADC/SPE Drilling Conference and Exhibition*, 3 2014.

[39] Tagir R. Salakhov, Vil U. Yamaliev, and Vladimir Dubinsky. An expert system approach to real-time diagnosis of drill bit condition and preventing its damage. In *SPE Oil and Gas India Conference and Exhibition*. SPE, 1 2010.

[40] Jerzy Sawicki and Tomasz Paczkowski. Effect of the hydrodynamic conditions of electrolyte flow on critical states in electrochemical machining. *EPJ Web of Conferences*, 92(92):02078–, 5 2015.

[41] Gunar Schroeder and Wieland Uffrecht. A new test rig for time-resolved pressure measurements in rotating cavities with pulsed inflow. In *Volume 7: Turbomachinery, Parts A, B, and C*, pages 2431–2440. ASMEDC, 10 2010.

[42] Milan Stanko and Michael Golan. Hexagonal reciprocating pump: advantages and weaknesses. *International Journal of Fluid Machinery and Systems*, 6(3):121–136, 9 2013.

[43] Zhongqiang Sun. Analysis on dust production mechanism and control measures in drilling and blasting method construction. *DEStech Transactions on Materials Science and Engineering*, (icmea), 4 2017.

[44] C.G. Wheat, H.W. Jannasch, M. Kastner, S. Hulme, J. Cowen, K.J. Edwards, B.N. Orcutt, and B. Glazer. *Fluid sampling from oceanic borehole observatories: design and methods for CORK activities (1990–2010)*. Integrated Ocean Drilling Program, 9 2011.

[45] Adam Wilson. Use of managed-pressure drilling requires adjustments to bridge gap to well control. *Journal of Petroleum Technology*, 70(01):59–60, 1 2018.