Invariant Structural Causal Modeling for Counterfactual Prediction and Intervention Planning in Multiphase Thermal-Hydraulic Networks

Suman Acharya[1], **Rohit Shrestha**[2], **and Dipesh Poudel**[3]

[1]***Purbanchal University, Biratnagar 56613, Morang, Nepal***
[2]***Nepal Sanskrit University, Beljhundi 22400, Dang, Nepal***
[3]***Lumbini Buddhist University, Lumbini 32900, Rupandehi, Nepal***

Abstract

Multiphase thermal-hydraulic networks are governed by conservation laws but are operated and monitored through sparse, heterogeneous measurements. In gas–liquid systems, regime-dependent interfacial organization and equipment interactions induce distribution shifts across operating campaigns, so predictors that fit one condition can fail under another even when the underlying physics is unchanged. This paper develops an invariant structural causal modeling framework for counterfactual prediction and intervention planning in networked two-fluid systems. The core idea is to treat uncertain closures, sensor transformations, and unobserved disturbances as causal variables that mediate between control actions, boundary conditions, and measured signals, and to identify a representation whose residual mechanisms remain invariant across environments defined by operating policies and facility states. The approach couples a constrained mechanistic backbone with a data-driven causal layer that separates stable causal effects from spurious correlations created by changes in excitation, telemetry filtering, and latent regime mixtures. Counterfactual operators are derived by intervening on inputs and policies within the causal model while preserving invariants, enabling robust prediction of pressure and holdup outcomes under planned actuation changes. Identification is performed using a joint objective that combines likelihood, conservation-consistent residuals, and an invariance penalty enforcing independence between environment labels and structural noise. A risk-aware planning formulation then selects interventions that are robust to remaining causal ambiguity. Computational studies on annular segments and branched networks demonstrate improved out-of-environment calibration and more stable intervention ranking compared to non-invariant baselines, particularly under sensor drift and closure drift scenarios common in practice.

1 Introduction

Multiphase thermal-hydraulic networks arise in energy transport, process operations, and infrastructure systems where gas and liquid phases share conduits, junctions, and equipment interfaces [1]. Despite the maturity of conservation-law modeling, practical prediction remains difficult because the observable signals are sparse and filtered, the effective closures change with interfacial structure and surface conditions, and the operating policies that generate data vary from campaign to campaign. When a facility changes its control

policy, a predictive model trained on historic telemetry can shift from being broadly reliable to being systematically biased, even if the governing physics has not changed. This recurring phenomenon is a form of distribution shift that is operationally unavoidable, because control actions are adjusted for safety, throughput, and maintenance constraints. A technical framework that is explicitly built to remain stable under policy and environment changes is therefore central to reliable forecasting and to planning interventions.

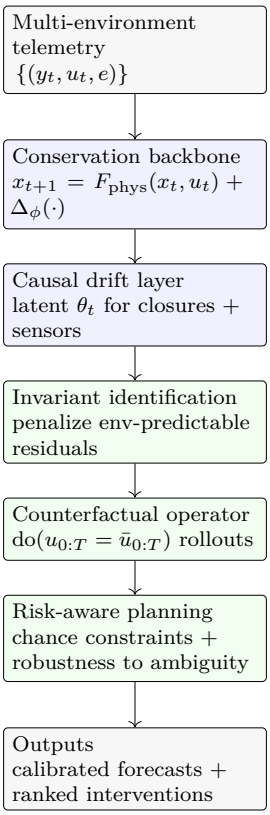

Figure 1. End-to-end workflow: multi-environment network telemetry is fused with a conservation-law backbone and a latent causal drift layer. Invariant identification discourages environment-specific residual patterns, enabling counterfactual rollouts under planned actuation and risk-aware selection of robust interventions.

A typical modeling pipeline for networked multiphase flow starts with a mechanistic two-fluid model or a mixture-level approximation and then tunes closures and parameters to match measurements. This yields a single calibrated model per facility configuration. Such calibration can be effective for a narrow operating envelope, but it implicitly assumes that the statistical relationships between measured signals and latent state are stable across time. In practice, these relationships can change due to sensor drift, changes in telemetry filtering, small geometry modifications, deposition and erosion, and shifts in upstream mixing that alter interfacial organization [2]. The resulting mismatch is rarely random; it is structured and correlated with operating policy. For example, if a facility reduces inlet gas fraction during certain seasons, then the measurement distribution shifts and the mapping from pressure signals to holdup changes, confounding a purely predictive learner. A model that does not separate causal mechanisms from spurious correlations can overfit the policy-induced distribution and fail when the policy changes.

The need for structured separation is suggested indirectly by recent machine-learning results in annular settings, where measured features have been shown to contain enough information to infer interfacial organization categories at high accuracy. Manikonda et al.

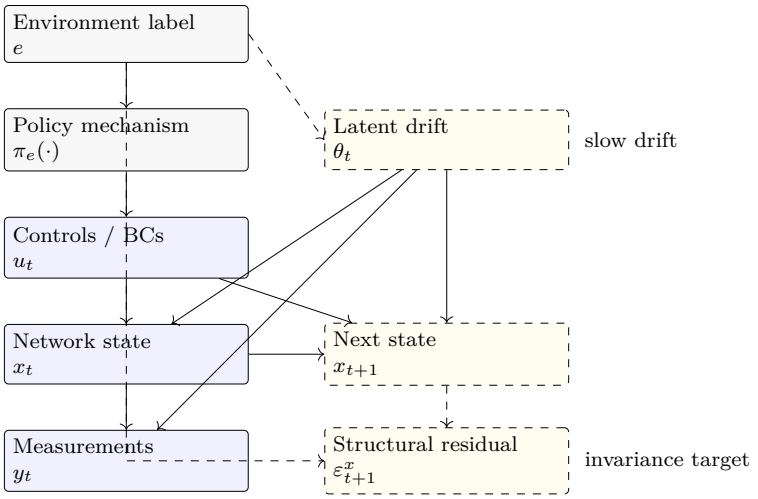

Figure 2. Time-slice structural causal view. The environment label indexes operating campaigns and induces distribution shifts through policy π_e and drift in latent variables θ_t (closures and sensor transformations). Stable mechanisms map (x_t, u_t, θ_t) to x_{t+1} and (x_t, θ_t) to y_t, while invariance is expressed by residuals that remain non-informative about e once conditioned on their causal parents.

(2021) introduced an early set of high-performing machine-learning classifiers for recognizing gas–liquid flow regimes in annular channels, demonstrating that practical signals can encode regime-dependent structure in a way that can be learned with low latency [3]. Such results emphasize that the data carry information about latent multiphase organization, but they also highlight a limitation of purely discriminative mappings: a classifier can be accurate under the data distribution induced by a particular experimental protocol while remaining fragile under shifts in excitation, sensor characteristics, or facility state. For intervention planning, robustness under such shifts is at least as important as in-distribution accuracy.

This paper proposes a structural causal modeling approach to address this robustness requirement. The central thesis is that multiphase network telemetry should be interpreted as arising from a causal data-generating process in which control actions and boundary conditions influence latent physical state, which then influences measurements through sensor transforms, while unobserved disturbances and closure variability mediate these relationships [4]. In this view, a change in operating policy changes the distribution of inputs, and therefore changes the distribution of observed telemetry, but it does not necessarily change the causal mechanisms that map state to measurements or inputs to state. A predictive model that explicitly targets invariance of structural mechanisms across environments can therefore generalize better under policy changes, because it is discouraged from exploiting correlations that are specific to a particular policy distribution. The key is to identify which mechanisms are stable and to represent uncertainty when the data cannot uniquely identify causal structure.

The proposed framework has three components. The first component is a constrained mechanistic backbone that enforces conservation structure and basic admissibility, providing a physically grounded representation of mass, momentum, and energy transport along network edges and through junctions. The backbone is not treated as a perfect model; instead it provides a structured set of state variables and residual operators that restrict the hypothesis space. The second component is a causal latent layer that represents closure variability, sensor transformations, and unobserved disturbances as explicit variables in a structural causal model, enabling separation between stable causal effects

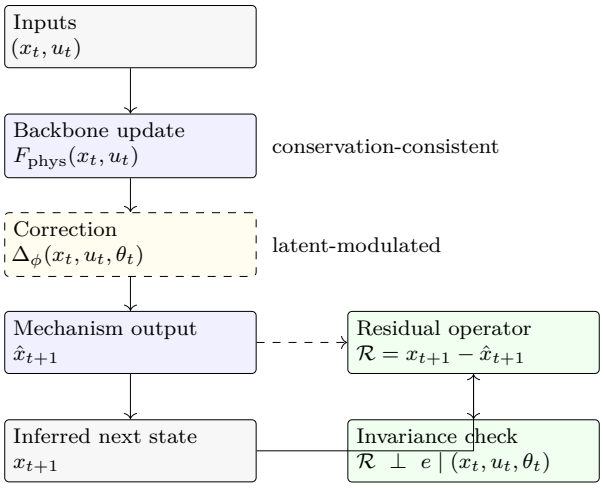

Figure 3. Mechanistic decomposition used to define structural residuals. A fixed conservation-law update F_{phys} is augmented by a constrained correction Δ_ϕ modulated by latent drift θ_t. Residuals $\mathcal{R}$ are computed against inferred x_{t+1} and are treated as structural noise; enforcing their invariance across environments regularizes the learned correction and supports counterfactual validity.

and environment-dependent confounding. The third component is a counterfactual operator and planning formulation: interventions are represented as do-operations on control actions and, when appropriate, on policy variables, and counterfactual outcomes are computed by propagating interventions through the invariant mechanisms while integrating over remaining uncertainty.

This emphasis on counterfactual reasoning distinguishes the contribution from standard prediction and from standard calibration [5]. Intervention planning requires ranking candidate actions by their expected effect on outcomes and constraints. If a model has learned spurious correlations, it can rank interventions incorrectly, even if it predicts observed telemetry well under the current policy. Causal invariance provides a systematic constraint that mitigates this failure by enforcing that structural noise terms, interpreted as exogenous disturbances, remain statistically independent of environment labels once conditioned on their causal parents. This independence is a measurable criterion that can be enforced during identification.

The present paper does not rely on regime labels as the core representation. Nevertheless, results showing that annular interfacial categories are learnable in varied annular-channel conditions reinforce that the latent structure is not arbitrary. Manikonda et al. (2022) reported similarly strong regime-recognition performance using modern classification models across annular-channel conditions, indicating that annular interfacial organization is expressed consistently in accessible signals [6]. In the causal framework developed here, such consistency motivates a latent representation that can encode regime mixtures without requiring discrete switching logic, and it motivates enforcing invariance across environments so that the inferred latent structure remains meaningful when operating policies change.

The paper is organized to build from causal formulation to identification and then to planning [7]. The next section formulates a structural causal model for multiphase networks coupled to a mechanistic backbone, defining environments, interventions, and counterfactual queries. The following section develops an invariance-based identification objective that combines likelihood, conservation-consistent residual penalties, and independence constraints between structural noise and environment labels. The next section studies identifiability and interventional design, clarifying when causal effects of interest

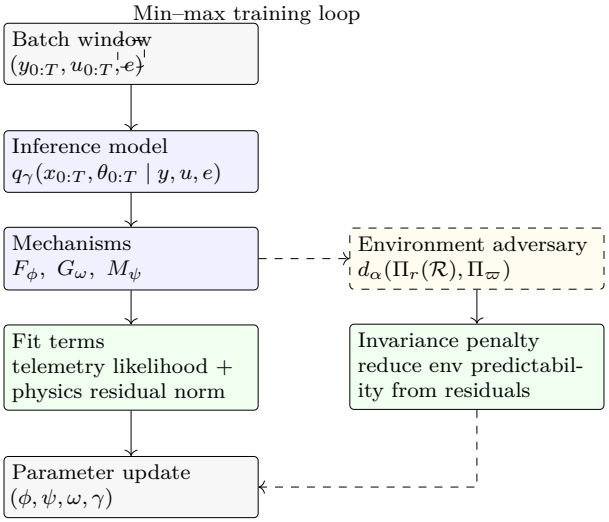

Figure 4. Invariant learning loop with latent inference and an environment adversary. Latent trajectories $(x_{0:T}, \theta_{0:T})$ are inferred from sparse telemetry and actions, mechanisms produce predictions and residuals, fit terms enforce data consistency and conservation consistency, and an adversary discourages residual patterns that reveal the environment label beyond causal parents.

can be estimated from observational data and when planned excitation is needed. The subsequent section introduces computational algorithms that implement the identification and counterfactual planning under admissibility constraints. A computational study section then evaluates robustness, calibration, and intervention ranking stability under synthetic policy shifts, sensor drift, and closure drift scenarios. The conclusion summarizes implications and limitations, focusing on how causal invariance can be used as a technical tool in multiphase network operations.

2 Structural Causal Formulation for Multiphase Networks

A structural causal model provides a formal representation of how observed and latent variables are generated from exogenous disturbances through structural equations. For multiphase networks, a direct causal representation in terms of full spatiotemporal fields is impractical, but a causal model can be built around a reduced state representation that is consistent with conservation laws and that aligns with available measurements and control inputs [8]. The aim is not to claim that the causal graph is fully known, but to impose a structured factorization that makes the sources of distribution shift explicit and that enables counterfactual reasoning under interventions.

Consider a network represented as a directed graph with edges corresponding to one-dimensional conduit segments and nodes corresponding to junctions and boundary manifolds. Let x_t denote a discretized state at time index t that includes conservative variables for gas and liquid phases along each edge, augmented by node variables required for coupling. Let u_t denote control actions and boundary inputs at time t, including inlet mass fluxes, valve openings, pump settings, outlet pressures, and possibly heat-flux controls. Let y_t denote the measurement vector at time t, including pressures, differential pressures, temperatures, and other telemetry channels. The data are collected under an operating policy that determines how u_t depends on past measurements and operational context. Let e denote an environment label that indexes an operating campaign, a facility condition, or a policy regime. The environment label is not assumed to be a cause of physics directly; rather it is a proxy for changes in the distribution of exogenous variables and in

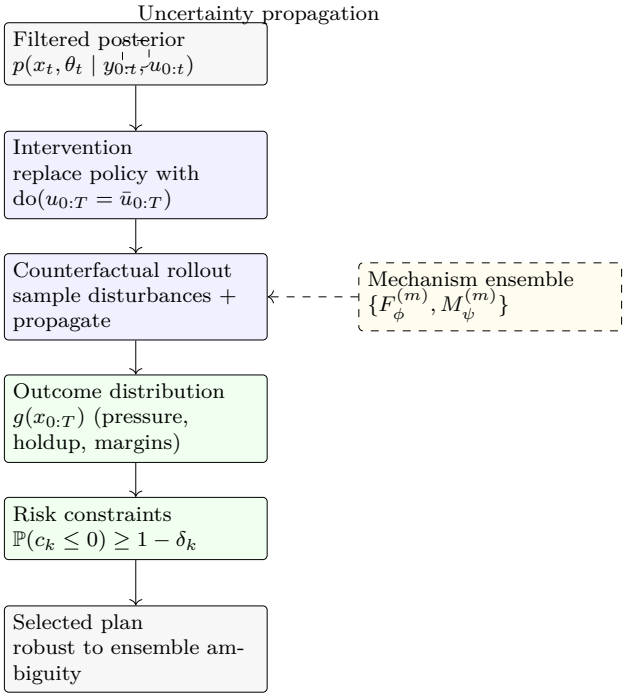

Figure 5. Counterfactual prediction and intervention planning. A posterior over the current physical state and drift variables is combined with a do-intervention that fixes the planned control trajectory. Rollouts through the invariant mechanisms yield an outcome distribution used to enforce chance constraints and to choose interventions that remain robust across an ensemble of plausible mechanisms.

the distribution of inputs induced by policy and facility state.

The mechanistic backbone is represented as a discretized operator that maps (x_t, u_t) to a predicted next state [9]. Because closures are uncertain and because unmodeled disturbances exist, the backbone is treated as an imperfect structural equation. A convenient representation is

$$x_{t+1} = F(x_t, u_t, c_t) \quad + \quad w_t, \tag{2.1}$$

where F is induced by an entropy-consistent discretization of two-fluid balance laws and junction coupling, c_t denotes closure and equipment parameters that influence interfacial exchange and losses, and w_t denotes exogenous process disturbances. The closure variable c_t is itself modeled causally as a function of slower facility state, operating conditions, and exogenous drift. A minimal structural equation for closure is

$$c_{t+1} = G(c_t, z_t, u_t) \quad + \quad \xi_t, \tag{2.2}$$

where z_t is a latent facility-state variable representing effects such as surface condition, deposition, and sensor calibration context, and ξ_t is exogenous closure noise. The latent variable z_t evolves slowly and can be driven by cumulative operating stress and maintenance actions:

$$z_{t+1} = H(z_t, u_t) \quad + \quad \upsilon_t. \tag{2.3}$$

This hierarchy makes explicit that a change in environment can correspond to a change in the distribution of z_t and therefore in c_t, even if the instantaneous physics mapping F

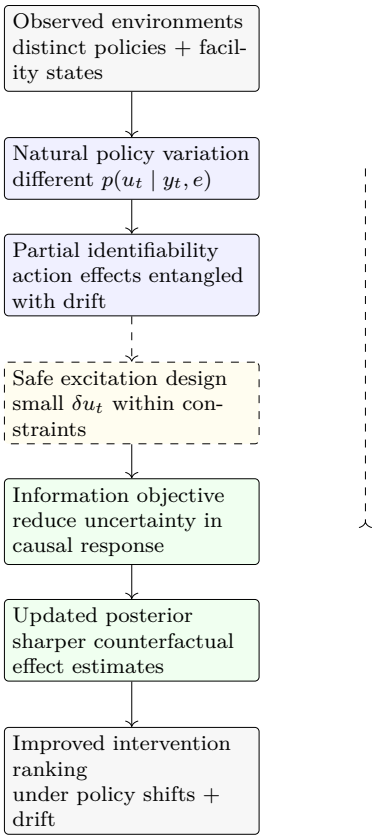

Figure 6. Identifiability pathway under policy-induced confounding. Multiple environments can provide quasi-experimental variation via differing policies, but remaining ambiguity in action effects motivates deliberate, constraint-aware excitation. Designing δu_t to maximize information about causal response parameters improves posterior concentration and stabilizes intervention ranking across future environment shifts.

remains unchanged.

Measurements are generated from the physical state through a sensor transform that includes calibration bias, filtering, and possibly nonlinear measurement mappings [10]. A structural equation for telemetry is

$$y_t = M(x_t, s_t) \quad + \quad \eta_t, \tag{2.4}$$

where s_t denotes sensor calibration and telemetry processing states, and η_t denotes measurement noise. The sensor state s_t can drift:

$$s_{t+1} = Q(s_t, z_t) \quad + \quad \rho_t, \tag{2.5}$$

reflecting the practical observation that sensor drift is correlated with facility condition and can covary with closure drift. The environment label e influences the distribution of (z_0, s_0) and can also influence the policy that generates u_t, but e is not treated as a direct parent of x_t when conditioned on u_t, c_t, and exogenous noises. This distinction is essential for invariance: the causal mechanisms F, G, M, and Q are assumed stable across environments, while the distribution of exogenous variables and the policy-induced distribution of u_t can change.

To make the causal structure explicit, one can describe a graph in which exogenous noises $(w_t, \xi_t, \upsilon_t, \eta_t, \rho_t)$ are mutually independent and influence endogenous variables through the structural equations. The environment label e influences the initial distributions of

Symbol	Description	Role	Domain
x_t	Discretized conservative state on network edges and nodes	Endogenous state	Mass, momentum, energy fields
u_t	Control actions and boundary inputs (valves, pumps, inlets, outlets, heat flux)	Intervention / policy variable	Actuation space
y_t	Telemetry (pressures, Δp, temperatures, auxiliary signals)	Observed output	Sensor space
e	Environment label indexing policy, facility state, or campaign	Context index	Finite set $\mathcal{E}$
θ_t	Low-dimensional latent drift vector (closures, sensor, facility)	Confounder / mediator	$\mathbb{R}^{d_\theta}$
F_{phys}	Entropy-consistent two-fluid network update operator	Mechanistic backbone	State transition map

Table 1. Core variables and operators in the structural causal model for multiphase thermal-hydraulic networks.

(z_0, s_0, c_0) and can influence the policy mapping that generates u_t. A policy is represented as [11]

$$u_t = \pi_e\big(h_t, \kappa_t\big), \tag{2.6}$$

where h_t denotes an information state available to the operator or controller, such as past telemetry and planned schedules, and κ_t denotes exogenous policy noise. Different environments correspond to different policies π_e or to different distributions of κ_t. This representation captures the operational reality that data are not generated by random excitation but by feedback policies designed to meet constraints, a feature that creates selection bias if not modeled.

Counterfactual reasoning requires defining interventions. An intervention corresponds to setting a variable to a value independently of its usual parents. In this framework, an intervention on control is represented as replacing the policy equation with $u_t = \bar{u}_t$ for a planned trajectory $\bar{u}_{0:T}$. Under such an intervention, the causal mechanisms F, G, M, and Q remain the same, but the distribution of u_t changes. The counterfactual outcome of interest might be a pressure constraint margin or a holdup-related functional over a horizon. Let $g(x_{0:T})$ denote such an outcome. The counterfactual query is the distribution of g under the do-intervention: [12]

$$p\big(g(x_{0:T}) \mid \mathrm{do}(u_{0:T} = \bar{u}_{0:T}),\, y_{0:t}\big), \tag{2.7}$$

conditioned on observed telemetry up to time t. This distribution integrates over posterior uncertainty in current state and latent variables and over future exogenous disturbances.

The role of invariance is to ensure that the estimated causal mechanisms generalize across environments. Formally, invariance is expressed by requiring that the structural noise terms be independent of environment labels when conditioned on their parents. For example, for the state transition equation, define the structural noise as

$$\varepsilon^x_{t+1} = x_{t+1} - F(x_t, u_t, c_t). \tag{2.8}$$

If F is correctly specified as a mechanism stable across environments, then ε^x_{t+1} should

Mechanism	Inputs	Output	Interpretation
$F(x_t, u_t, \theta_t)$	State, control, drift	x_{t+1}	Network two-fluid dynamics with closure and equipment variability
$G(\theta_t, u_t)$	Drift state, control	θ_{t+1}	Evolution of closures and sensor/facility condition
$M(x_t, \theta_t)$	State, drift	y_t	Sensor and telemetry map with calibration and filtering effects
$Q(s_t, z_t)$	Sensor and facility states	s_{t+1}	Long-term calibration drift coupled to facility condition
$\pi_e(h_t, \kappa_t)$	Information state, policy noise	u_t	Environment-specific operating policy generating actions

Table 2. Selected structural mechanisms and their roles in the causal data-generating process.

have the same conditional distribution across environments, and in particular should be independent of e given (x_t, u_t, c_t). Similar invariance conditions apply to closure noise and sensor noise. In practice, exact independence is unattainable due to model mismatch, but enforcing approximate invariance discourages the learner from encoding environment-specific correlations into F.

The above causal model is high dimensional due to x_t being a discretized field. The framework therefore uses a structured reduction [13]. Rather than representing c_t and z_t as full fields, they are represented as a low-dimensional latent vector that modulates closures and sensor transforms. This is consistent with the observation that many closure drifts and sensor drifts exhibit low-dimensional structure, such as a global friction scaling change and a small number of calibration offsets. Let θ_t denote a latent vector capturing these effects, so that $c_t = c(\theta_t)$ and $s_t = s(\theta_t)$. The causal model becomes

$$\begin{aligned} x_{t+1} &= F(x_t, u_t, \theta_t) + w_t, \\ \theta_{t+1} &= G(\theta_t, u_t) + \xi_t, \\ y_t &= M(x_t, \theta_t) + \eta_t. \end{aligned} \tag{2.9}$$

This reduced causal structure makes identification feasible while preserving the essential confounding pathways: environment affects u_t through policy and affects θ_t through facility state, and both u_t and θ_t influence x_t and y_t.

The causal representation also clarifies a key difficulty: observational data under a policy may not identify the causal effect of changing u_t if u_t is strongly correlated with unobserved θ_t due to feedback and selection. For example, if operators open a valve more when they believe the system is accumulating liquid, then u_t is correlated with latent holdup

Environment type	Dominant source of shift	Typical examples	Modeling impact
Policy regime	Change in feedback or scheduling logic for u_t	Aggressive vs. conservative outlet pressure control	Alters input excitation patterns
Sensor regime	Calibration and filtering changes	Transducer replacement, software filter update	Changes mapping from state to telemetry
Facility configuration	Geometry and equipment modifications	Added bypass line, valve replacement	Modifies network topology and losses
Upstream mixing	Interfacial organization at inlets	Different gas fraction programs or separators	Affects closure behavior and slip
Seasonal operation	Load profiles, ambient conditions	Seasonal throughput and temperature shifts	Shifts operating envelopes and stress

Table 3. Representative environment labels and associated sources of distribution shift in the causal formulation.

state and with closure drift conditions, creating confounding. A causal model provides a language to express and address this confounding, either through explicit modeling of θ_t and policy variables or through designed interventions that break correlations. The next section develops an identification objective that enforces invariance across environments as a practical surrogate for stable causal mechanism learning under such confounding [14].

3 Invariant Mechanism Identification with Conservation-Consistent Residuals

Identification seeks to estimate the structural mechanisms and latent variables of the causal model from data collected across environments. The distinguishing feature of the proposed approach is that it uses invariance across environments as an explicit constraint, combined with conservation-consistent residual penalties from the mechanistic backbone. The goal is to learn a mechanism representation that is predictive within environments and remains stable under changes in environment label distribution, thereby supporting counterfactual reasoning under interventions that shift the input distribution.

Let the set of environments be indexed by $e \in \mathcal{E}$. For each environment, data consist of sequences $\{(y_t, u_t)\}_{t=0}^{T_e}$, and possibly boundary condition metadata and planned schedules. The latent state x_t and latent drift variables θ_t are unobserved. The identification problem is thus a latent-variable learning problem. A purely likelihood-based approach would maximize the joint likelihood of observed telemetry under a parameterized model of transitions and measurements. Such approaches can overfit environment-specific correlations if the model class is flexible, especially when the latent variables are underconstrained by sparse measurements. Invariance provides an additional constraint: the distribution of structural noise should not depend on environment once conditioned on the causal parents.

A practical mechanism representation is a parameterized transition operator F_ϕ and measurement operator M_ψ, with latent drift dynamics G_ω [15]. The parameters are

Drift component	Physical meaning	Temporal scale
Friction scaling mode	Slow variation in effective wall friction due to deposition/erosion	Very slow; cumulative operating stress and maintenance windows
Interfacial organization mode	Mixtures of annular, slug, or stratified-like behavior reflected in closures	Slow to medium; responds to upstream mixing and campaigns
Sensor bias and gain	Offset and sensitivity changes in pressure and temperature channels	Slow; driven by calibration events and aging
Heat-transfer effectiveness	Deviation in energy exchange due to fouling or insulation changes	Very slow; correlated with thermal duty cycles

Table 4. Illustrative components of the latent drift vector θ_t capturing closure and sensor evolution.

Loss term	Representative form	Purpose	Emphasis
Telemetry likelihood $\mathcal{L}_{\text{like}}$	$-\log p_\psi(y_{0:T} \mid x_{0:T}, \theta_{0:T})$	Fit mechanistic-causal model to observed signals	Data-consistent predictions
Physics residual penalty $\mathcal{L}_{\text{phys}}$	$\sum_t \|\mathcal{R}(x_{t+1}, x_t, u_t, \theta_t)\|_W^2$	Encourage conservation-consistent latent trajectories	Backbone consistency
Invariance penalty $\mathcal{L}_{\text{inv}}$	Dependence measure $\mathcal{I}(r_t; e \mid \varpi_t)$	Enforce environment-invariant structural noise	Cross-environment stability
Adversarial loss $\mathcal{L}_{\text{adv}}$	Environment classification loss for d_α	Approximate conditional independence via an adversary	Residual decorrelation
Regularization terms	Parameter norms and drift smoothness penalties	Control overfitting and enforce slow drift	Model parsimony

Table 5. Key components of the joint objective used for invariant mechanism identification.

(ϕ, ψ, ω), and the model is

$$\begin{aligned} x_{t+1} &= F_\phi(x_t, u_t, \theta_t) \; + \; w_t, \\ \theta_{t+1} &= G_\omega(\theta_t, u_t) \; + \; \xi_t, \\ y_t &= M_\psi(x_t, \theta_t) \; + \; \eta_t. \end{aligned} \tag{3.1}$$

The mechanistic backbone is embedded in F_ϕ by representing it as the sum of a known discretized conservation operator and a learned correction that is constrained to preserve admissibility. Specifically, one writes

$$F_\phi(x, u, \theta) = \; F_{\text{phys}}(x, u) \; + \; \Delta_\phi(x, u, \theta), \tag{3.2}$$

where F_{phys} is a fixed entropy-consistent two-fluid update on the discretized network, and Δ_ϕ is a learnable closure-and-mismatch correction modulated by θ. The correction

Residual type	Parents (conditioning set)	Invariance condition	Primary use
State transition residual ε^x_{t+1}	(x_t, u_t, θ_t)	$\varepsilon^x_{t+1} \perp e \mid x_t, u_t, \theta_t$	Stabilize physical dynamics across policies
Drift residual ε^θ_{t+1}	(θ_t, u_t)	$\varepsilon^\theta_{t+1} \perp e \mid \theta_t, u_t$	Route environment effects through drift modes
Measurement residual ε^y_t	(x_t, θ_t)	$\varepsilon^y_t \perp e \mid x_t, \theta_t$	Separate sensor drift from physics and policy
Multi-step residual summaries	Windowed $(x_{t:t+k}, u_{t:t+k})$	Approximate invariance over horizons	Guard against short-horizon overfitting

Table 6. Residuals and associated invariance conditions used to constrain the causal mechanisms.

is constrained so that it does not violate positivity and dissipation requirements. This decomposition reduces the burden on learning and improves interpretability of invariance constraints because structural noise can be interpreted relative to a physically grounded predictor.

To enforce conservation-consistent behavior without requiring full-field supervision, the framework uses a residual penalty based on the mechanistic discretization. Define a residual operator

$$\mathcal{R}(x_{t+1}, x_t, u_t, \theta_t) = \; x_{t+1} - F_{\text{phys}}(x_t, u_t) - \Delta_\phi(x_t, u_t, \theta_t). \tag{3.3}$$

Under the model, $\mathcal{R}$ corresponds to process noise w_t. A residual penalty encourages $\mathcal{R}$ to be small in a weighted norm consistent with process noise assumptions. This penalty plays a dual role: it regularizes latent inference by discouraging arbitrary latent trajectories that fit measurements but violate physics, and it provides a consistent definition of structural noise for invariance constraints [16].

Invariance constraints are imposed by requiring that, for each environment, the conditional distribution of residuals given parents be the same. A tractable surrogate is to penalize dependence between residuals and environment labels after conditioning. Let r_t denote a residual summary statistic derived from $\mathcal{R}$, such as a set of weighted residual components that are most informative given sparse measurements. The invariance objective penalizes the mutual dependence between r_t and e given a representation of parents. Because conditional independence is difficult to enforce directly, the framework uses a practical penalty based on the ability of an environment classifier to predict e from residuals and parent representations. The identification objective is

$$\begin{aligned} \min_{\phi,\psi,\omega} \quad & \sum_{e\in\mathcal{E}} \mathbb{E}_e\Big[\ell\big(y_t, \hat{y}_t\big)\Big] \\ & + \; \lambda_{\text{phys}}\, \mathbb{E}\Big[\|\mathcal{R}\|^2_W\Big] \\ & + \; \lambda_{\text{inv}}\, \mathcal{I}(r_t;\, e \mid \varpi_t), \end{aligned} \tag{3.4}$$

where ℓ is a telemetry likelihood loss, $\hat{y}_t = M_\psi(x_t, \theta_t)$ is the predicted measurement, W is

Scenario	Network topology	Source of shift	Evaluation focus
Annular segment	Single conduit with inlet and outlet boundaries	Policy change in outlet pressure modulation and friction/sensor drift	Out-of-environment uncertainty calibration
Branched junction	Y-shaped network with mixing junction and outlet leg	Upstream mixing and policy differences across branches	Intervention ranking under confounding
Sensor drift regime	Fixed physics with two calibration periods	Pressure and temperature bias/gain changes	Separation of physics and telemetry transformations
Environment partition study	Same physical setup, varied environment labeling	Coarse vs. fine environment granularity	Sensitivity of invariance to partitioning

Table 7. Computational study scenarios used to assess robustness and intervention ranking stability.

a residual weighting matrix, $\mathcal{I}$ denotes a conditional dependence measure, and ϖ_t denotes a representation of parent variables, such as (x_t, u_t, θ_t) or a compressed sufficient statistic. The invariance penalty discourages residuals that leak environment information beyond what is explained by parent variables, encouraging stable mechanisms across environments.

Latent variables x_t and θ_t are inferred jointly with parameters. The framework uses a variational approximation to the posterior over latent trajectories, parameterized by an inference model $q_\gamma(x_{0:T}, \theta_{0:T} \mid y_{0:T}, u_{0:T}, e)$. The overall learning objective becomes a variational bound on the likelihood augmented by residual and invariance penalties. The variational structure matters because invariance should be enforced on structural noise under the inferred latent parents, not under arbitrary point estimates [17]. Enforcing invariance at the distribution level reduces the risk that the method satisfies invariance by artificially inflating posterior uncertainty in one environment and collapsing in another.

The mechanistic backbone imposes admissibility constraints that must hold during inference and learning. In particular, volume fractions and phase masses must remain within bounds, and thermodynamic variables must remain in physically meaningful ranges. The framework enforces these constraints by projecting latent state samples onto an admissible set during inference updates. Projection is implemented as a differentiable map so that gradients can flow through it. This is essential because invariance constraints can otherwise be satisfied by pushing latent states into regions where the residual definition changes qualitatively, a form of degeneracy that would undermine counterfactual validity. By constraining latent states to remain admissible, the residual and invariance penalties remain comparable across environments.

A crucial modeling decision is the definition of environments. Environments can represent different operators, different control policies, different seasons, different facility configurations, or different sensor calibration periods [18]. If environments are defined too coarsely, invariance can be violated for benign reasons and the model may become overly conservative. If environments are defined too finely, invariance constraints become weak because each environment has too little data. The framework defines environments in a way that captures policy and facility-state changes that are expected to induce shifts in input distributions and in latent drift distributions. Environment labels can also be treated

Planning quantity	Definition (schematic)	Role in decision-making	Risk treatment
Expected cost $J(\bar{u}_{0:T})$	Throughput, energy, and wear aggregated over horizon	Objective for economic and operational performance	Mean over posterior and disturbance samples
Constraint violation probability	$\mathbb{P}(c_k(\bar{u}_{0:T}) > 0)$	Safety and admissibility for pressure and holdup	Chance constraints with empirical quantiles
Intervention ranking score	Ordering induced by J and violation probabilities	Selection among candidate actuation trajectories	Robust ranking across ensemble
Extrapolation distance	Distance of $\bar{u}_{0:T}$ from observed policy support	Guard against mechanism extrapolation	Penalization in planning objective

Table 8. Decision metrics in the counterfactual intervention planning formulation.

as latent, inferred from metadata and telemetry statistics, but the present development assumes that a reasonable environment partition is available.

The invariance penalty can also be localized. Rather than enforcing invariance on all residual components, which may be unrealistic under model mismatch, the framework enforces invariance on residual components associated with stable conservation mechanisms and on measurement residual components that should be stable across environments given correct sensor modeling. This allows the model to represent genuine environment-dependent changes through θ_t while still enforcing stability of core mechanisms. For example, if friction scaling truly changes due to deposition, that change should be captured by θ_t, and residuals conditioned on θ_t should remain invariant. The invariance constraint thus encourages the model to route environment dependence through explicit drift variables rather than through hidden, environment-specific modifications of the transition operator [19].

The causal interpretation of invariance is that residuals represent exogenous disturbances, and exogenous disturbances should not be predictable from environment labels once causal parents are accounted for. This interpretation is central for counterfactual validity: when computing counterfactual outcomes under new policies, one must assume that exogenous disturbance distributions remain comparable, or at least that changes are represented explicitly. Invariance-based identification thus provides a disciplined basis for extrapolating under policy changes.

Nevertheless, invariance is not sufficient for full causal identification. If the data lack excitation that varies u_t independently of latent confounders, then causal effects of u_t on outcomes may remain ambiguous even under invariance. The next section addresses this identifiability issue by analyzing when the causal effect of interventions can be estimated from observational multi-environment data and how interventional design can improve identifiability while respecting operational constraints.

Aspect	Limitation	Consequence	Mitigation strategy
Excitation coverage	Limited variation of u_t along some directions	Partial identifiability of action effects	Safe experiment design and use of multi-policy data
Environment partitioning	Ambiguous definitions of environment labels	Weak or noisy invariance constraints	Use operational metadata and sensitivity diagnostics
Drift parameterization	Misspecified latent drift modes θ_t	Environment effects leak into core mechanisms	Physics-guided drift bases and model selection
Model mismatch	Simplified backbone or closure representation	Residual invariance cannot be fully achieved	Localized invariance and residual diagnostics
Computational load	High-dimensional state and ensembles	Expensive training and planning rollouts	Structure-exploiting solvers and reduced backbones

Table 9. Limitations and mitigation strategies associated with the invariant structural causal modeling approach.

4 Identifiability and Interventional Design Under Policy-Induced Confounding

Counterfactual prediction and intervention planning rely on identifying how outcomes change when control actions are changed. In a networked multiphase system, data are typically collected under feedback policies that adapt actions based on telemetry and operational context [20]. This induces confounding because u_t is correlated with latent state and drift variables, and these latent variables also influence future outcomes. A causal model makes this confounding explicit, but it does not automatically resolve it. This section analyzes identifiability of intervention effects in the proposed framework and develops an interventional design strategy that improves identifiability by exploiting safe excitation and multi-environment variation.

Consider a target quantity of interest $J(u_{0:T})$ defined as a functional of the state trajectory under a planned action sequence, such as expected constraint margin, expected energy consumption, or expected throughput. The causal effect of an action change is the difference between counterfactual expectations under different do-interventions. If the causal model were fully known, this could be computed directly. In practice, the model is learned from data, and the key question is whether the data identify the relevant causal pathways.

A useful abstraction is to consider a reduced causal model for a subset of variables that capture the dominant effects relevant to the outcome. Let h_t denote a latent holdup-related summary, let p_t denote a pressure-drop-related summary, and let θ_t denote drift

variables capturing friction and sensor bias. A simplified structural model is [21]

$$\begin{aligned} h_{t+1} &= f_h(h_t, u_t, \theta_t) \; + \; \epsilon^h_{t+1}, \\ p_{t+1} &= f_p(p_t, h_t, u_t, \theta_t) \; + \; \epsilon^p_{t+1}, \\ y_t &= g(p_t, h_t, \theta_t) \; + \; \epsilon^y_t. \end{aligned} \tag{4.1}$$

Even in this reduced model, u_t may be confounded with h_t because the policy chooses u_t based on y_t, which depends on h_t and p_t. If h_t is imperfectly observed, then conditioning on y_t may not block confounding fully. The causal effect of u_t on future p_{t+1} is therefore not identifiable from observational data unless additional assumptions hold.

Multi-environment invariance provides one such assumption: if environments correspond to different policies, then the distribution of u_t given latent state can vary across environments. This variation can help identify causal effects because it provides data where u_t changes in different ways relative to latent state. Intuitively, if one policy opens a valve aggressively in response to pressure changes while another policy opens it conservatively, then the conditional distribution of u_t given y_t differs, and this can help disentangle the direct effect of u_t from the effect of latent state. In the causal framework, this corresponds to having multiple environments with different policy mechanisms π_e. If the structural mechanisms mapping (h_t, u_t, θ_t) to h_{t+1} and p_{t+1} are invariant, then combining data across environments can improve identification of those mechanisms.

A formal identifiability condition can be described in terms of overlap and rank of sensitivity operators. For a local linearization, suppose that in environment e one can approximate

$$\begin{aligned} \delta x_{t+1} &= A_t \, \delta x_t \; + \; B_t \, \delta u_t \; + \; C_t \, \delta \theta_t \; + \; \varepsilon_{t+1}, \\ \delta y_t &= H_t \, \delta x_t \; + \; D_t \, \delta \theta_t \; + \; \eta_t, \end{aligned} \tag{4.2}$$

with matrices $(A_t, B_t, C_t, H_t, D_t)$ depending on the operating point [22]. The ability to identify the effect of δu_t depends on whether δu_t varies independently of δx_t and $\delta \theta_t$ in the data, and whether B_t has directions that influence measured outcomes. Under a policy, δu_t may lie in a subspace correlated with $H_t \delta x_t$ due to feedback. Across environments, the union of these subspaces may be richer, increasing identifiability.

However, policy variation may still be insufficient if all policies are similar or if operational constraints prevent excitation in certain directions. This motivates interventional design: deliberately injecting small, safe perturbations into control actions to excite informative modes. In multiphase networks, such excitation must be constrained by safety limits on pressure and by the risk of inducing undesirable regimes. The design problem is therefore to choose an excitation signal δu_t that improves identifiability while respecting constraints. The causal framework provides a principled objective: maximize expected reduction in posterior uncertainty about causal effect parameters relevant to the outcome.

Let β denote parameters that govern the causal response of interest, such as components of B_t or a parameterization of how u_t influences holdup dynamics [23]. The posterior uncertainty in β under the current model can be represented by a covariance matrix Σ_β. Under a planned excitation experiment, the expected posterior covariance can be approximated by a Fisher-information update. A tractable objective is to maximize a scalar measure of information gain, such as the log determinant reduction. The interventional

design objective can be written as

$$\begin{aligned} \max_{\delta u_{0:T}} \quad & \mathbb{E}\Big[\log\det(\Sigma_\beta) \;-\; \log\det(\Sigma_{\beta|\delta u})\Big] \\ \text{s.t.} \quad & \mathbb{P}\big(c(x_t, u_t + \delta u_t) \le 0\big) \ge 1 - \delta, \end{aligned} \tag{4.3}$$

where $c(\cdot)$ encodes safety constraints and δ is a risk tolerance. Because exact evaluation is difficult, the framework uses a conservative approximation based on posterior predictive quantiles and linearized constraints. The key is that excitation is designed not to optimize performance directly but to improve causal identifiability so that future intervention planning is more reliable.

The design must also account for sensor transformations and drift variables [24]. If sensor bias is uncertain, then excitation signals may be misinterpreted as state changes rather than as bias changes. The causal model represents sensor drift explicitly through θ_t, and the design objective includes information gain about both causal response parameters and sensor drift parameters when those parameters are confounded. This avoids experiments that would be informative only if sensors were perfectly calibrated.

A distinctive advantage of the causal-invariance framework is that it can use naturally occurring policy variation as quasi-experimental variation. In many facilities, operators switch policies between operating modes, such as startup, steady production, and shutdown. These mode switches can provide intervention-like variation without deliberate excitation, provided that environments are labeled and invariance is enforced. The framework can then estimate causal response parameters from observational multi-environment data and reserve deliberate excitation for remaining ambiguous directions. This is operationally appealing because deliberate excitation can be costly or risky.

Identifiability also depends on the choice of latent representation [25]. If θ_t is too expressive, it can absorb the effect of u_t, making causal effects underidentified. If it is too restrictive, it can force environment dependence into the transition operator, violating invariance. The framework resolves this by constraining θ_t to represent plausible drift modes with slow dynamics, and by enforcing that residuals remain invariant across environments only after conditioning on θ_t. This encourages θ_t to explain environment-dependent drift while leaving u_t effects stable.

The sections analysis implies that causal effects relevant to intervention planning are partially identifiable from multi-environment observational data when policy variation provides sufficient independent variation and when latent drift is modeled appropriately. When this is insufficient, safe excitation designed using information gain objectives can improve identifiability. The next section turns these principles into computational algorithms that implement invariant identification, latent inference, and counterfactual intervention planning under admissibility constraints and risk criteria.

5 Algorithms for Invariant Learning and Counterfactual Intervention Planning

The computational challenge is to implement invariant mechanism learning and counterfactual planning in high-dimensional networked multiphase systems with sparse observations and admissibility constraints. The framework must support inference of latent state and drift variables, enforcement of invariance across environments, and evaluation of counterfactual outcomes under planned interventions [26]. This section presents a computational strategy built around a constrained variational inference loop combined with an invariance adversary and a risk-aware planning stage that propagates causal uncertainty.

The mechanistic backbone provides a differentiable transition operator F_{phys} that can be evaluated on the discretized network. The learned correction Δ_ϕ is parameterized as a constrained function of (x_t, u_t, θ_t). Rather than learning an unconstrained neural mapping, the correction is represented in a basis that respects locality and dissipation. A typical representation is a low-rank field over edges whose coefficients are functions of θ_t and local nondimensional groups derived from x_t and u_t. This representation prevents the correction from encoding arbitrary environment-specific features. The correction is also passed through monotone transforms to enforce sign constraints on dissipative components. The resulting operator remains differentiable for gradient-based learning.

Latent inference uses a recognition model q_γ that outputs distributions over (x_t, θ_t) given telemetry, inputs, and environment labels. Because x_t is high dimensional, the recognition model uses a structured parameterization: it outputs a low-dimensional latent code that is decoded into a physically admissible state through a constrained decoder that enforces volume fraction bounds and thermodynamic consistency [27]. This is implemented by decoding into entropy-variable space and mapping to conservative variables, followed by a differentiable projection that enforces positivity. The projection is used both in inference and in forward prediction, ensuring that latent samples remain admissible.

The learning objective combines telemetry likelihood, physics residual penalty, and invariance penalty. Telemetry likelihood is computed under a sensor model that includes drift variables θ_t, capturing calibration offsets and filtering parameters. The physics residual penalty is computed from $\mathcal{R}$, weighted by a covariance reflecting expected process variability. The invariance penalty is implemented through an adversary that attempts to predict environment labels from residual summaries. Specifically, define a residual summary $r_t = \Pi_r(\mathcal{R}_t)$ and a parent summary $\varpi_t = \Pi_\varpi(x_t, u_t, \theta_t)$. An adversary network d_α tries to predict e from (r_t, ϖ_t), while the mechanism learner tries to make this prediction difficult by minimizing adversary accuracy, thereby reducing dependence. The min–max objective is

$$\min_{\phi,\psi,\omega,\gamma} \max_{\alpha} \quad \mathcal{L}_{\text{like}} + \lambda_{\text{phys}} \mathcal{L}_{\text{phys}} + \lambda_{\text{inv}} \, \mathbb{E}\Big[\ell_{\text{adv}}\big(d_\alpha(r_t, \varpi_t), e\big)\Big], \tag{5.1}$$

where ℓ_{adv} is a classification loss. The adversary approximates a dependence measure: if environment is predictable from residuals given parents, invariance is violated [28]. Minimizing adversary predictability encourages invariance. Because the adversary can exploit trivial information such as time index correlated with environment, the framework randomizes time windows and includes only physically meaningful residual summaries, reducing leakage.

The physics residual penalty is computed from a one-step or multi-step rollout of the mechanistic backbone plus correction. A multi-step penalty improves stability by discouraging corrections that are accurate for one step but unstable over longer horizons. However, multi-step rollout can be expensive. The framework uses a mixed strategy: one-step residual penalties are computed at every time step, and occasional multi-step penalties are computed on randomly selected windows. This yields a training signal that targets both local consistency and long-horizon stability.

A key algorithmic issue is balancing invariance against fit. If invariance is enforced too strongly, the model may be forced to ignore genuine environment-dependent drift that is not captured by θ_t, leading to poor fit and potentially invalid counterfactual predictions [29]. If invariance is too weak, the model may encode environment-specific correlations

and fail to generalize. The framework addresses this by gradually increasing λ_{inv} during training, allowing the model to first fit core dynamics and then to allocate environment dependence into explicit drift variables under invariance pressure. This schedule encourages a representation in which stable mechanisms are learned early and drift variables are used to explain environment differences as invariance becomes more strict.

Counterfactual prediction is performed by sampling from the posterior over current latent state and drift variables, applying a do-intervention on control actions, and propagating forward through the learned transition operator while sampling exogenous disturbances. The do-intervention replaces the policy-generated u_t by a planned $\bar{u}_t$. The counterfactual distribution of outcomes is approximated by Monte Carlo samples:

$$\hat{p}(g) = \frac{1}{N} \sum_{j=1}^{N} \delta(g - g^{(j)}), \tag{5.2}$$

where $g^{(j)}$ is the outcome computed from the jth rollout. Because Monte Carlo rollouts can be expensive, the framework uses low-variance sampling by sharing random seeds across candidate interventions, enabling consistent comparison of interventions.

Intervention planning is formulated as a risk-aware optimization problem. Let $J(\bar{u}_{0:T})$ denote an expected cost functional combining throughput and energy, and let $c_k(\bar{u}_{0:T})$ denote constraint functionals such as maximum pressure and minimum holdup margin. The planning objective is

$$\begin{aligned} \min_{\bar{u}_{0:T}} \quad & \mathbb{E}\Big[J(\bar{u}_{0:T})\Big] \\ \text{s.t.} \quad & \mathbb{P}\Big(c_k(\bar{u}_{0:T}) \le 0\Big) \ge 1 - \delta_k, \end{aligned} \tag{5.3}$$

where probabilities are under the counterfactual distribution conditioned on current telemetry. This chance-constrained problem is approximated using quantiles computed from rollouts [30]. The resulting optimization is solved by gradient-based methods by differentiating through rollouts and projections. Differentiability is ensured by using smooth projections and by maintaining the mechanistic backbone as a differentiable operator. Because gradients through long rollouts can be noisy, the framework uses a receding-horizon approach and truncates gradients when uncertainty dominates.

A crucial feature of causal uncertainty is that different mechanism parameterizations can satisfy invariance and fit but imply different intervention effects. The framework represents this ambiguity by maintaining an ensemble of mechanism parameter sets that are consistent with the data. Planning then optimizes a robust objective that accounts for ensemble variation. A conservative robust objective is to minimize expected cost under the worst-case ensemble member within a plausible set, while a less conservative alternative is to penalize variance in predicted outcomes across ensemble members. The planning objective becomes

$$\begin{aligned} \min_{\bar{u}_{0:T}} \quad & \frac{1}{M} \sum_{m=1}^{M} \mathbb{E}\Big[J_m(\bar{u}_{0:T})\Big] \\ & + \lambda_{\text{rob}} \operatorname{Var}\Big(J_m(\bar{u}_{0:T})\Big), \end{aligned} \tag{5.4}$$

where m indexes ensemble members [31]. This formulation discourages interventions whose estimated effect depends strongly on uncertain mechanism details, improving robustness.

The algorithm must also ensure that invariance holds during planning. If the planner se-

lects interventions far outside the observed input distribution, invariance assumptions may fail because mechanisms are extrapolated. The framework therefore includes an extrapolation penalty based on the distance between planned inputs and the observed support under the environments. This distance can be measured in a latent input embedding space. The planner then trades off performance against extrapolation risk, preventing unrealistic interventions.

The overall computational pipeline thus consists of an invariant learning stage and a planning stage. Invariant learning estimates mechanisms and latent variables from multi-environment telemetry while enforcing admissibility and invariance [32]. Planning uses the learned mechanisms to evaluate counterfactual outcomes of candidate interventions and selects robust interventions under uncertainty and constraints. The next section evaluates this pipeline in computational studies designed to stress generalization under policy shifts and to test whether invariance improves intervention ranking stability when environments change.

6 Computational Studies: Out-of-Environment Calibration and Intervention Ranking Stability

Computational studies are designed to evaluate whether invariant structural mechanism learning improves robustness of counterfactual predictions and intervention planning under environment changes common in multiphase network operations. The evaluation emphasizes two metrics that are particularly relevant for decision-making: calibration of predicted uncertainty under out-of-environment policy shifts and stability of intervention ranking when policies and facility conditions change. The studies use simulated telemetry from a high-resolution two-fluid network model with closure drift and sensor drift processes that create realistic distribution shifts, and then train and test the proposed framework across multiple environments defined by differing policies and drift regimes.

The first study considers a single annular segment operated under two different policies. Policy A modulates outlet pressure aggressively in response to inlet flow fluctuations, producing a telemetry dataset with frequent pressure transients and a relatively broad distribution of control actions. Policy B maintains outlet pressure within a tighter band, producing a narrower distribution of control actions and fewer large transients. Both policies operate on the same underlying physical segment, but the resulting telemetry distributions differ [33]. Closure drift is introduced as a slow change in effective wall friction scaling correlated with cumulative operating stress, and sensor drift is introduced as a slow pressure offset drift correlated with facility condition. These drift processes create environment-dependent distributions of latent variables even when the structural transition mechanism remains stable.

Models are trained on one policy environment and evaluated on the other. A non-invariant baseline model that directly predicts measurements from past telemetry and actions can fit in-environment well but exhibits systematic bias when evaluated on the other policy, particularly in predicting pressure response to action changes because the baseline has learned policy-specific correlations between actions and latent holdup states. The invariant causal framework reduces this bias by enforcing that structural residuals remain independent of environment labels given inferred parents, which discourages encoding policy-specific correlations into the transition mechanism. In evaluation, predictive intervals for pressure responses maintain closer to nominal coverage under policy shift, while baseline intervals become miscalibrated and under-cover. Calibration is evaluated by comparing the fraction of observed pressures that fall within predicted 90% credible

intervals across held-out trajectories.

Intervention ranking stability is evaluated by defining a set of candidate valve-opening trajectories and computing the predicted distribution of a risk metric, such as probability of exceeding a pressure limit over a horizon. Under the truth simulator, interventions can be ranked by their true risk [34]. The models are evaluated by how often they preserve the correct ordering of interventions, particularly among the top-risk and bottom-risk sets. The invariant framework exhibits higher ranking stability under policy shift because it has learned a mechanism representation that is less dependent on the policy-induced input distribution. The baseline model, which is accurate in predicting typical behavior under its training policy, can rank interventions incorrectly when the intervention deviates from that policy distribution, a failure that is particularly problematic when planning new operating strategies.

The second study considers a branched network with a junction and an outlet segment. Two environments correspond to different upstream mixing conditions that affect effective interfacial drag and slip behavior, modeled as an environment-dependent distribution of a latent drift variable θ_t. Policies also differ: one policy prioritizes throughput and thus operates closer to constraint limits, while the other maintains larger margins. The telemetry includes boundary pressures and a small number of differential pressure measurements. The key challenge is confounding: the policy that operates near constraints produces data where control actions are correlated with near-constraint states, and the environment-dependent drift produces additional correlations between telemetry patterns and facility state. A model that fails to represent drift explicitly can attribute environment differences to action effects, leading to incorrect counterfactual predictions [35].

The invariant framework explicitly represents drift via θ_t and enforces invariance of residuals across environments conditioned on θ_t. In evaluation, the inferred drift variables correlate with the true drift process, and the residual dependence on environment labels is reduced. Counterfactual predictions of pressure response under a standardized intervention remain more consistent across environments compared to baselines. Importantly, when the data are insufficient to uniquely identify drift versus action effects, the framework expresses this ambiguity as increased posterior uncertainty, which in planning leads to more conservative intervention choices. This behavior is evaluated by measuring how predicted risk distributions widen in ambiguous conditions and by assessing whether the planner avoids interventions that have high uncertainty in constraint satisfaction.

The third study stresses sensor drift confounding. A single environment is used for physics, but the sensor calibration state changes between two periods, producing two environments in telemetry space. A baseline model trained on one calibration period can fail on the other, and a naive invariant model can also fail if it tries to enforce invariance by ignoring calibration changes [36]. The proposed framework includes sensor drift variables in θ_t and a sensor model M_ψ that maps state to measurements with calibration parameters. Invariance is enforced on structural noise conditioned on θ_t, enabling the model to represent calibration shifts as drift rather than as physics changes. In evaluation, counterfactual predictions of physical outcomes, such as holdup proxies computed from the state, remain consistent across calibration environments, and intervention ranking based on physical risk measures remains stable. This illustrates that invariance must be paired with explicit representation of plausible environment-dependent drift mechanisms; otherwise invariance can enforce an incorrect stability assumption.

A qualitative diagnostic used in the studies is residual environment predictability. After training, an environment classifier is trained on residual summaries conditioned on parent summaries. Under successful invariance, classifier accuracy should approach chance. The

invariant framework reduces this accuracy relative to baselines, indicating that residuals have become less environment-informative. When residual predictability remains high, it signals either model mismatch or insufficient drift representation, and planning is expected to be less reliable [37]. This diagnostic provides a practical tool for assessing whether causal invariance assumptions are plausibly satisfied in a given deployment.

A key concern is whether invariance can be satisfied trivially by inflating process noise or by choosing latent representations that wash out environment differences. The framework mitigates this by including a physics residual penalty and an admissibility constraint, which restrict trivial solutions that ignore dynamics. Moreover, intervention ranking evaluation reveals whether invariance is achieved in a way that improves decision-relevant outcomes rather than merely reducing a classifier signal. In the studies, models that reduce environment predictability but degrade intervention ranking are interpreted as over-regularized; the training schedule that gradually increases invariance weight reduces the frequency of such failures.

The computational studies also explore sensitivity to environment partitioning. If environments are defined by very short time windows, invariance constraints become noisy and can encourage the model to overfit. If environments are defined too coarsely, invariance becomes weak and generalization can degrade. The studies show that moderate environment granularity that captures policy changes and calibration periods yields the strongest improvements [38]. While this dependence is expected, it underscores that the causal framework relies on meaningful environment distinctions and that operational metadata can be valuable.

Overall, the computational evidence supports the thesis that invariant structural mechanism learning improves out-of-environment calibration and intervention ranking stability in multiphase network settings where policy changes and drift processes create structured distribution shifts. The framework does not eliminate all ambiguity; when data do not identify causal effects, the framework expresses uncertainty and penalizes interventions that rely on poorly identified mechanisms. This behavior is essential for safe planning, where incorrect ranking can lead to constraint violations. The conclusion summarizes the contribution and outlines implications for deploying causal invariance as a technical tool in multiphase operations.

7 Conclusion

This paper developed an invariant structural causal modeling framework for counterfactual prediction and intervention planning in networked gas–liquid two-fluid systems under policy-induced distribution shifts. The approach represents multiphase network telemetry as arising from stable structural mechanisms coupled with latent drift variables capturing closure and sensor evolution, and it enforces invariance across environments by penalizing dependence between structural residuals and environment labels once causal parents are accounted for. A constrained mechanistic backbone provides conservation-consistent structure and admissibility constraints, while a causal latent layer allocates environment dependence into explicit drift variables rather than into environment-specific transition behavior. Counterfactual operators are defined by do-interventions on control actions and policies, and a risk-aware planning formulation selects interventions that are robust to remaining causal ambiguity and extrapolation risk [39].

The framework addresses a central operational challenge: data are generated under feedback policies that induce confounding between actions and latent state, and facility conditions drift over time. Multi-environment invariance provides a disciplined constraint

that reduces reliance on spurious correlations specific to a policy distribution, thereby improving generalization under policy change. When identifiability remains limited, the framework represents this limitation as posterior uncertainty and penalizes interventions whose estimated benefit depends strongly on uncertain mechanisms. Computational studies on annular segments and branched networks indicated improved out-of-environment calibration and more stable intervention ranking compared to non-invariant baselines, especially under sensor drift and closure drift scenarios.

The approach has limitations that are intrinsic to the problem class. Invariance cannot compensate for fundamental lack of excitation in directions needed to identify causal effects, and thus safe interventional design remains important in some deployments. Moreover, the quality of invariance enforcement depends on environment partitioning and on the adequacy of the drift variable representation. These limitations are not failures of the framework but reflect the inherent challenges of causal learning under sparse sensing in multiphase systems. Within those constraints, the papers contribution is to provide a technically explicit and physically grounded pathway from multi-environment telemetry to counterfactual reasoning and robust intervention planning, using invariance as a measurable proxy for stable causal mechanism identification in multiphase network operations [40].

References

[1] M. Ripepe et al., "Infrasonic early warning system for explosive eruptions," *Journal of Geophysical Research: Solid Earth*, vol. 123, no. 11, pp. 9570–9585, Nov. 23, 2018.

[2] S. B. Pawar, "Process engineering aspects of vertical column photobioreactors for mass production of microalgae," *ChemBioEng Reviews*, vol. 3, no. 3, pp. 101–115, Jun. 7, 2016.

[3] K. Manikonda et al., "Application of machine learning classification algorithms for two-phase gas-liquid flow regime identification," in *Abu Dhabi International Petroleum Exhibition and Conference*, SPE, 2021, D041S121R004.

[4] N. Qin, C.-j. Lu, and J. Li, "Numerical investigation of characteristics of water-exit ventilated cavity collapse," *Journal of Shanghai Jiaotong University (Science)*, vol. 21, no. 5, pp. 530–540, Sep. 29, 2016.

[5] H. Joos, P. Spichtinger, and U. Lohmann, "Orographic cirrus in a future climate," *Atmospheric Chemistry and Physics*, vol. 9, no. 20, pp. 7825–7845, Oct. 20, 2009.

[6] K. Manikonda et al., "Horizontal two-phase flow regime identification with machine learning classification models," in *International Petroleum Technology Conference*, IPTC, 2022, D011S021R002.

[7] R. Erfani, H. Zare-Behtash, and K. Kontis, "Influence of shock wave propagation on dielectric barrier discharge plasma actuator performance," *Journal of Physics D: Applied Physics*, vol. 45, no. 22, pp. 225 201–, May 16, 2012.

[8] U. Borges and M. Jamiolahmady, "Well test analysis in tight gas reservoirs," *EUROPEC/EAGE Conference and Exhibition*, pp. 1178–1198, Jun. 8, 2009.

[9] A. Efhaima and M. H. A. Dahhan, "Hydrodynamics study in a fluidized bed reactor using gamma ray computed tomography (ct) technique," *International Journal of Petrochemical Science & Engineering*, vol. 1, no. 4, Nov. 29, 2016.

[10] S. Jasper, D. P. Gradzki, R. Bracke, J. Hussong, M. Petermann, and R. Lindken, "Nozzle cavitation and rock erosion experiments reveal insight into the jet drilling process," *Chemie Ingenieur Technik*, vol. 93, no. 10, pp. 1610–1618, Jul. 22, 2021.

[11] F. Zhu, Z. Lu, and S. Wang, "Flame spread and extinction over a thick solid fuel in low-velocity opposed and concurrent flows," *Microgravity Science and Technology*, vol. 28, no. 2, pp. 87–94, Nov. 25, 2015.

[12] Q. Wang, H. Wang, and S. Xin, "Suitability of an mrmce (multi-resolution minimum cross entropy) algorithm for online monitoring of a two-phase flow," *Measurement Science and Technology*, vol. 22, no. 10, pp. 104 010–, Sep. 9, 2011.

[13] S. Byrdina et al., "Dipolar selfpotential anomaly associated with carbon dioxide and radon flux at syabrubensi hot springs in central nepal," *Journal of Geophysical Research: Solid Earth*, vol. 114, no. B10, Oct. 3, 2009.

[14] F. Bao, H. Hao, Z. Yin, and C. Tu, "Numerical study of nanoparticle deposition in a gaseous microchannel under the influence of various forces," *Micromachines*, vol. 12, no. 1, pp. 47–, Jan. 1, 2021.

[15] S. Calvari, A. Bonaccorso, and G. Ganci, "Anatomy of a paroxysmal lava fountain at etna volcano: The case of the 12 march 2021, episode," *Remote Sensing*, vol. 13, no. 15, pp. 3052–, Aug. 3, 2021.

[16] E. V. Fischer et al., "Nitric acid phase partitioning and cycling in the new england coastal atmosphere," *Journal of Geophysical Research: Atmospheres*, vol. 111, no. D23, Oct. 25, 2006.

[17] A. Myers-Pigg, P. Louchouarn, and R. Teisserenc, "Flux of dissolved and particulate low-temperature pyrogenic carbon from two high-latitude rivers across the spring freshet hydrograph," *Frontiers in Marine Science*, vol. 4, Feb. 15, 2017.

[18] D. Düzenli, E. Seker, S. Senkan, and I. Onal, "Epoxidation of propene by high-throughput screening method over combinatorially prepared cu catalysts supported on high and low surface area silica," *Catalysis Letters*, vol. 142, no. 10, pp. 1234–1243, Jul. 26, 2012.

[19] S. Srinivasan et al., "A machine learning framework for rapid forecasting and history matching in unconventional reservoirs.," *Scientific reports*, vol. 11, no. 1, pp. 21 730–21 730, Nov. 5, 2021.

[20] M. Rodal and M. Schlutow, "Waves in the gas centrifuge: Asymptotic theory and similarities with the atmosphere," *Journal of Fluid Mechanics*, vol. 928, Oct. 5, 2021.

[21] S.-Z. Qi, X.-H. Tan, X.-P. Li, Z. Meng, Y.-J. Xu, and D. Tang, "Transient pressure analysis of inclined well in continuous triple-porosity reservoirs with dual-permeability behavior," *Frontiers in Energy Research*, vol. 8, Dec. 1, 2020.

[22] M. Haider, K. Schwaiger, F. Holzleithner, and R. Eisl, "A comparison between passive regenerative and active fluidized bed thermal energy storage systems," *Journal of Physics: Conference Series*, vol. 395, no. 1, pp. 012 053–, Nov. 26, 2012.

[23] A. Costa, Y. Suzuki, and T. Koyaguchi, "Understanding the plume dynamics of explosive super-eruptions.," *Nature communications*, vol. 9, no. 1, pp. 654–654, Feb. 13, 2018.

[24] C. Liu, Z. Li, and H. Wang, "Drag force and transport property of a small cylinder in free molecule flow: A gas-kinetic theory analysis.," *Physical review. E*, vol. 94, no. 2, pp. 023 102–023 102, Aug. 10, 2016.

[25] Q. Tian, C. Yudong, W. Luo, P. Liu, and B. Ning, "Transient flow of a horizontal well with multiple fracture wings in coalbed methane reservoirs," *Energies*, vol. 13, no. 6, pp. 1498–, Mar. 22, 2020.

[26] R. Rahman, H. Zhu, and A. Yu, "Numerical analysis of effects of specularity coefficient and restitution coefficient on the hydrodynamics of particles in a rotating drum," *Processes*, vol. 10, no. 1, pp. 167–167, Jan. 15, 2022.

[27] J. M. Mogollón, A. W. Dale, I. L'Heureux, and P. Regnier, "Journal of geophysical research - impact of seasonal temperature and pressure changes on methane gas production, dissolution, and transport in unfractured sediments," *Journal of Geophysical Research*, vol. 116, no. 3, pp. 03 031–, Sep. 9, 2011.

[28] X. Liu et al., "Measuring technologies for cfb solid circulation rate: A review and future perspectives," *Energies*, vol. 15, no. 2, pp. 417–417, Jan. 6, 2022.

[29] S. Govindarajan, A. Mishra, and A. Kumar, "Darcy-and pore-scale issues associated with multi-phase fluid flow through a petroleum reservoir," *Earth Science Malaysia*, vol. 4, no. 2, pp. 108–117, Jun. 10, 2020.

[30] D. A. Knopf, U. Pöschl, and M. Shiraiwa, "Radial diffusion and penetration of gas molecules and aerosol particles through laminar flow reactors, denuders, and sampling tubes," *Analytical chemistry*, vol. 87, no. 7, pp. 3746–3754, Mar. 16, 2015.

[31] M. A. Abd-Rabbo, R. Y. Sakr, M. A. Mohammad, and M. M. Mandour, "Numerical study of heat transfer from elliptic tube in a fluidized bed," *Journal of Engineering Research and Reports*, pp. 1–9, Nov. 20, 2019.

[32] S. Pouget, M. Bursik, C. Johnson, A. J. Hogg, J. C. Phillips, and R. S. J. Sparks, "Interpretation of umbrella cloud growth and morphology: Implications for flow regimes of short-lived and long-lived eruptions," *Bulletin of Volcanology*, vol. 78, no. 1, pp. 1–19, Dec. 11, 2015.

[33] C. S. J. Shaw, J. Eyzaguirre, B. J. Fryer, and J. E. Gagnon, "Regional variations in the mineralogy of metasomatic assemblages in mantle xenoliths from the west eifel volcanic field, germany," *Journal of Petrology*, vol. 46, no. 5, pp. 945–972, Jan. 28, 2005.

[34] M. Wang, X. Dou, R. Ming, L. Weiqiang, W. Zhao, and C. Tan, "Semi-analytical rate decline solutions for a refractured horizontal well intercepted by multiple reorientation fractures with fracture face damage in an anisotropic tight reservoir," *Energies*, vol. 14, no. 22, pp. 7482–, Nov. 9, 2021.

[35] D. Gilberg and K. Steiner, "Size segregation in compressible granular shear flows of binary particle systems," *Granular Matter*, vol. 22, no. 2, pp. 1–15, Apr. 2, 2020.

[36] C. López-Muñoz, J. R. Garcia-Cascales, F. Velasco, and R. Otón-Martínez, "An energetic model for detonation of granulated solid propellants," *Energies*, vol. 12, no. 23, pp. 4459–, Nov. 22, 2019.

[37] J.-M. Lim et al., "Ultra-high throughput synthesis of nanoparticles with homogeneous size distribution using a coaxial turbulent jet mixer.," *ACS nano*, vol. 8, no. 6, pp. 6056–6065, May 28, 2014.

[38] J. Goral et al., "Confinement effect on porosity and permeability of shales," *Scientific reports*, vol. 10, no. 1, pp. 49–49, Jan. 8, 2020.

[39] C. A. Brown, "Jet-surface interaction test: Far-field noise results," *Journal of Engineering for Gas Turbines and Power*, vol. 135, no. 7, pp. 071 201–, Jun. 10, 2013.

[40] P. Zhang, L. Hu, J. N. Meegoda, and S. Gao, "Micro/nano-pore network analysis of gas flow in shale matrix," *Scientific reports*, vol. 5, no. 1, pp. 13 501–13 501, Aug. 27, 2015.