International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(RESEARCH ARTICLE)

Check for updates

HYBRID PHYSICS-INFORMED MACHINE LEARNING FRAMEWORK FOR PREDICTING METHANE HYDRATE RESERVOIR PRODUCTIVITY

Saiful Alam *

Boone Pickens School of Geology, Oklahoma State University, Stillwater, Oklahoma, USA.
* Corresponding Author

ORCID Details
Saiful Alam: https://orcid.org/0000-0001-7368-9062

International Journal of Science and Research Archive, 2026, 20(02), 347–355

Publication history: Received on 04 July 2026; revised on 12 August 2026; accepted on 14 August 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.20.2.1628

ABSTRACT

Methane hydrates hold enormous quantities of natural gas in a form that could meaningfully add to the world's future energy supply, yet accurately forecasting how productive a given reservoir will be remains difficult. The obstacle is coupling: thermal, hydraulic, mechanical, and geochemical processes all interact during hydrate dissociation and gas release, and untangling their combined effect on production is not straightforward. This paper puts forward a hybrid physics-informed machine learning (HPIML) framework that folds core reservoir-flow equations into a data-driven model, with the aim of improving both prediction accuracy and the model's ability to generalize. Physical constraints are combined with deep learning models trained on numerical-simulation output and, where available, field data, so that gas production rate, pressure evolution, hydrate saturation, and permeability can all be estimated across a range of production scenarios. Because the physics terms regularize training, the model is less prone to overfitting than a purely data-driven counterpart while still tracking known reservoir behavior. The framework is expected to outperform conventional empirical and purely data-driven methods on accuracy, computational cost, and interpretability, and it includes a sensitivity-analysis component that flags which geological and operational variables matter most for methane recovery. Taken together, the study points to a practical way of pairing physics-based knowledge with machine learning to support better decisions around methane hydrate exploitation, tighter reservoir management, and lower uncertainty as production moves toward commercial scale a step toward sustainable hydrate development that fits within a broader, carbon-conscious approach to petroleum engineering.

Keywords: Methane Hydrates; Physics-Informed Machine Learning; Reservoir Simulation; Gas Production Forecasting; Deep Learning; Hydrate Dissociation Kinetics; Sensitivity Analysis

1 INTRODUCTION

Natural gas hydrates are ice-like solids in which methane molecules sit trapped inside a lattice of hydrogen-bonded water molecules, a structure that only forms under low temperature and elevated pressure. Enormous volumes of methane are locked up this way in sediments beneath permafrost and along continental margins by most estimates, far more than is held in conventional natural gas accumulations. That scale is why methane hydrates have long been treated

as a candidate long-term energy source, provided a production strategy can be found that actually works technically and economically.

Getting gas out of a hydrate reservoir means dissociating the solid hydrate into free gas and water, usually by depressurization, thermal stimulation, inhibitor injection, or some mix of the three. What makes the dissociation process hard to predict is how tightly it is coupled across mechanisms: heat and mass transfer set the pace of decomposition, shifting fluid saturations change both absolute and relative permeability, and changes in effective stress can trigger compaction or geomechanical instability. Add pronounced reservoir heterogeneity on top of that coupling, and long-term productivity forecasting for hydrate reservoirs becomes considerably harder than the equivalent problem for a conventional gas reservoir.

Two broad classes of forecasting tools are in common use. Fully coupled numerical simulators solve the governing partial differential equations for mass and energy conservation alongside hydrate phase-equilibrium and kinetic dissociation models. They are physically rigorous, but they are also computationally expensive, demand detailed site characterization, and are not well suited to rapid what-if analysis, uncertainty quantification, or real-time decision support. The alternative is empirical decline-curve relationships and, increasingly, purely data-driven machine learning models trained on simulated or historical production records. These are cheap to run and good at picking up nonlinear patterns, but they tend to generalize poorly outside their training range and can produce results that violate basic physics a real problem given how sparse and noisy field data from hydrate pilot tests usually are.

Physics-informed machine learning (PIML) offers something of a middle ground. By folding governing equations and boundary conditions directly into a neural network's training objective, PIML pulls predictions toward physically admissible solutions while keeping most of the flexibility and speed of a data-driven model. This idea has already been demonstrated in subsurface flow, reservoir characterization, and other multiphysics problems, but its use in methane hydrate reservoirs where phase change, kinetic dissociation, and strongly nonlinear permeability evolution all have to be represented at once is still fairly limited.

This paper proposes a hybrid physics-informed machine learning (HPIML) framework built specifically for methane hydrate reservoir productivity prediction. It combines governing reservoir-flow and hydrate-kinetics equations with deep learning models trained jointly on numerical-simulation datasets and field observations. Four objectives guide the work: (i) formulate the physical constraints relevant to hydrate dissociation and gas production; (ii) design a hybrid network architecture and composite loss function that enforces those constraints during training; (iii) define a training and validation strategy suited to the data-sparse conditions typical of hydrate reservoirs; and (iv) set out an evaluation and sensitivity-analysis methodology for assessing predictive performance, computational efficiency, interpretability, and the relative influence of key geological and operational parameters.

The rest of the paper is organized as follows. Section 2 reviews related work on hydrate reservoir simulation and machine learning in reservoir engineering. Section 3 sets out the governing physical equations behind hydrate dissociation and gas production. Section 4 presents the HPIML framework itself architecture, physics-informed loss, data strategy, and training procedure. Section 5 covers the evaluation methodology, and Section 6 discusses expected performance and sensitivity outcomes. Section 7 turns to implications for reservoir management and sustainable development, Section 8 lays out limitations and future work, and Section 9 concludes.

2 LITERATURE REVIEW

2.1 Numerical simulation of hydrate reservoirs

Fully coupled simulators **TOUGH+HYDRATE, STOMP-HYD, MH-21**, among others, solve mass- and energy-balance equations for several mobile phases together with hydrate phase-equilibrium and kinetic dissociation models. They have been used extensively to evaluate depressurization, thermal stimulation, and combined production strategies, both conceptually and at field-pilot scale, including analyses of the Mallik, Nankai Trough, and Ignik Sikumi tests. The trade-off for that rigor is cost: these simulators need detailed site characterization, fine discretization near the dissociation front, and substantial compute, none of which lends itself to rapid scenario screening, real-time forecasting, or large-scale uncertainty quantification.

2.2 Empirical and data-driven approaches

To cut computational cost, empirical decline-curve and analytical type-curve methods borrowed from conventional reservoir engineering have been adapted to hydrate systems, generally at the expense of representing the underlying dissociation physics. More recently, purely data-driven machine learning models artificial neural networks, random forests, gradient-boosted trees, and recurrent architectures such as LSTM networks have been used to forecast gas

production rates and cumulative recovery from simulated or historical datasets. Once trained, these models handle complex nonlinear relationships efficiently, but nothing in their structure enforces conservation laws, so predictions can drift outside physical bounds when extrapolated beyond the training distribution, a real limitation given how little long-duration hydrate production data exists.

2.3 Physics-informed machine learning

Physics-informed neural networks (PINNs) fold governing partial differential equations into a network's loss function via automatic differentiation, so the network learns solutions that satisfy both the data and known physical laws simultaneously. This approach has been applied successfully in fluid mechanics, subsurface flow and transport, and parameter estimation in heterogeneous porous media, generally improving generalization and cutting data requirements relative to purely data-driven models. Extensions to coupled multiphysics subsurface problems flow, geomechanics, phase change provide a reasonable methodological starting point for hydrate-specific work, though hydrate reservoirs bring extra complexity in the form of kinetic, rate-limited phase transformation and strongly saturation-dependent permeability, neither of which shows up in single-phase or purely hydraulic subsurface problems.

2.4 Research gap

Interest in hydrate reservoir simulation and in physics-informed learning has each grown on its own, but the two have rarely been brought together for hydrate-specific productivity forecasting. Hydrate-focused ML studies to date are mostly purely data-driven and trained on limited simulation datasets, while PIML studies in subsurface flow seldom incorporate hydrate-phase kinetics, saturation-dependent relative permeability evolution, or the coupled thermal-hydraulic-mechanical-chemical (THMC) processes specific to hydrate dissociation. This study addresses that gap directly, building a hybrid framework around the governing physics of methane hydrate reservoirs.

3 GOVERNING PHYSICS AND PROBLEM FORMULATION

The framework proposed here is built around the conservation equations that govern coupled fluid flow and hydrate dissociation in a porous reservoir. These equations are the physical constraints embedded in the machine learning model later on.

3.1 Mass conservation

Mass balance is enforced separately for the gas, aqueous, and hydrate phases. For a representative phase α, the governing equation takes the general form:

$$\boldsymbol{\partial(\varphi\, S\alpha\, \rho\alpha)/\partial t \;+\; \nabla\cdot(\rho\alpha\, u\alpha) \;=\; q\alpha \;\pm\; \dot{M}h}$$

where φ is porosity, Sα, ρα, and uα are the saturation, density, and Darcy flux of phase α, qα represents source/sink terms associated with production or injection, and Ṁh is the rate of mass exchange with the hydrate phase during dissociation or formation. The Darcy flux for each mobile phase is given by $u\alpha = -(k\cdot k\alpha)/\mu\alpha\,(\nabla p\alpha - \rho\alpha\, g)$, with k the intrinsic (absolute) permeability, kα the relative permeability, μα the viscosity, and g the gravitational acceleration vector.

3.2 Energy conservation

Hydrate dissociation is strongly endothermic, so an energy-balance equation is solved jointly with the mass-balance equations:

$$\frac{\partial}{\partial t}\left[\varphi\sum_{\alpha} S_\alpha\, \rho_\alpha U_\alpha + (1-\phi)\rho_\alpha C_R T\right] + \nabla\cdot\left(\sum_{\alpha}\rho_\alpha\, h_\alpha U_\alpha -\right) = q h - \Delta H_d \dot{M}_h$$

where **Uα** and hα are the internal energy and enthalpy of phase α, ρR and CR are the density and specific heat of the rock matrix, λ is the effective thermal conductivity, qh represents external heat sources, and ΔHd is the enthalpy of hydrate dissociation.

3.3 Hydrate dissociation kinetics

Dissociation is represented with a kinetic, rate-limited model of the Kim–Bishnoi form, where the rate of methane release scales with how far local pressure has fallen below the phase-equilibrium pressure at the prevailing temperature:

$$\dot{M}_h = k_s \cdot A_s \cdot M_g \cdot (p_{eq}(T) - p)$$

where ks is the intrinsic dissociation rate constant, As is the reactive surface area per unit volume (itself a function of hydrate saturation), Mg is the molar mass of methane, p is local pore pressure, and peq(T) is the equilibrium dissociation pressure taken from the hydrate phase-equilibrium curve.

3.4 Permeability and saturation coupling

Absolute permeability rises as hydrate saturation falls and pore space opens up during dissociation, typically captured through saturation-dependent porosity–permeability relationships (Kozeny–Carman-type or power-law forms, for example), while relative permeability of the gas and aqueous phases follows saturation-dependent functions such as modified Brooks–Corey or van Genuchten forms parameterized for hydrate-bearing sediments. These relationships are strongly nonlinear and history-dependent, and they are a major source of predictive uncertainty — which is exactly why the physics-informed regularization strategy in Section 4 is worth pursuing.

Together, the equations in Sections 3.1–3.4, plus appropriate initial and boundary conditions (prescribed initial pressure, temperature, and hydrate saturation fields; specified bottom-hole pressure or flow-rate constraints at production wells), define the forward problem the hybrid model is meant to approximate in a computationally efficient, data-augmented way.

4 PROPOSED HYBRID PHYSICS-INFORMED MACHINE LEARNING FRAMEWORK

4.1 Architecture overview

The HPIML framework couples a deep neural network surrogate with the governing equations from Section 3 through a composite loss function. The network takes spatiotemporal coordinates and reservoir/operational parameters as inputs and outputs the primary state variables of interest. A feed-forward multilayer perceptron with sinusoidal or hyperbolic-tangent activations handles the core spatiotemporal mapping, chosen so that automatic differentiation of outputs with respect to inputs is straightforward, while a recurrent branch (LSTM/GRU) captures sequential dependencies in well-level operational time series. The two branches feed into a shared dense layer before the final predictions are produced.

Table 1: Representative inputs and outputs of the proposed hybrid physics-informed machine learning framework.

Category	Variable	Role in Framework
Spatiotemporal inputs	x, y, z, t	Collocation coordinates for physics-residual evaluation
Static reservoir inputs	φ0, k0, Sh0, T0, p0	Initial porosity, permeability, hydrate saturation, temperature, pressure
Operational inputs	pwf(t), Q(t), production method	Well control history (depressurization/thermal/combined)
Primary outputs	p(x,t), T(x,t), Sh(x,t), k(x,t)	State variables governed by Section 3 equations
Derived outputs	Qg(t), Qw(t)	Gas and water production rates at wells

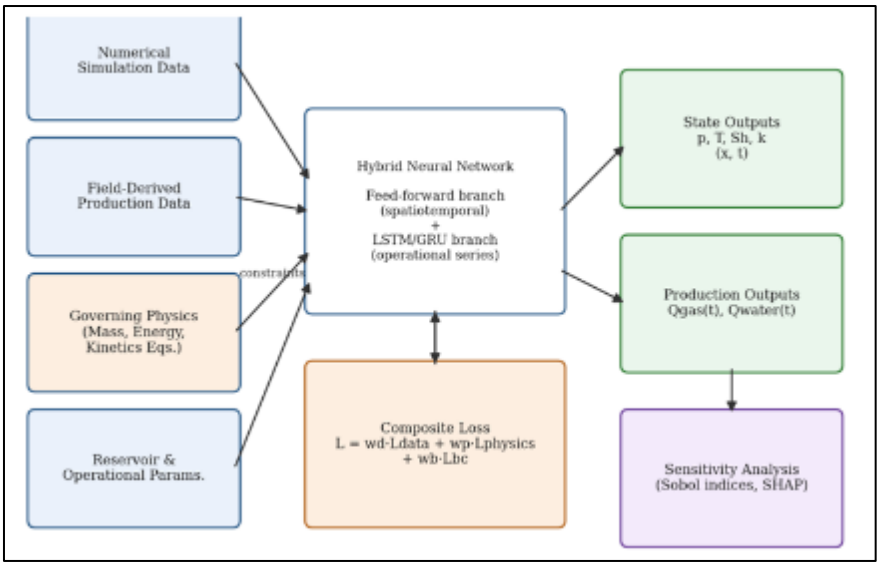

Figure 1: Schematic architecture of the proposed Hybrid Physics-Informed Machine Learning (HPIML) framework, showing the coupling between the physics module, the hybrid feed-forward/recurrent network, the composite loss function, and the resulting state, production, and sensitivity outputs.

4.2 Physics-informed loss formulation

The network is trained by minimizing a composite loss function made up of a data-fidelity term and a physics-residual term:

$$L(\theta) = wd \cdot Ldata(\theta) + wp \cdot Lphysics(\theta) + wb \cdot Lbc(\theta)$$

Ldata is the mean-squared error between network predictions and available labeled data (numerical-simulation samples and, where available, field observations). Lphysics is the mean-squared residual of the governing mass-, energy-, and kinetics-equations from Section 3, evaluated at randomly sampled collocation points across the spatiotemporal domain using automatic differentiation to compute the necessary spatial and temporal derivatives of network outputs. Lbc penalizes violations of the initial and boundary conditions. The weights wd, wp, and wb are tuned — through adaptive weighting schemes such as gradient normalization or self-adaptive weights — so that no single loss component dominates training, which matters given how different the magnitudes of the mass-, energy-, and kinetics-residuals tend to be.

4.3 Data sources and feature engineering

Training data come from two complementary sources. The first is an ensemble of numerical-simulation runs generated across a designed range of reservoir and operational parameters — permeability, initial hydrate saturation, porosity, production method, well constraints — to densely sample the input space and provide labeled state variables throughout the domain. The second is field-derived data from documented hydrate production tests, incorporated where available to anchor the model to real reservoir behavior and to support transfer learning from the simulation-trained network to field conditions. Input features are normalized, and domain-informed feature engineering — dimensionless pressure drawdown, cumulative produced volume, depressurization rate — is used to improve training stability and physical interpretability.

4.4 Training strategy

Training happens in two stages. In pretraining, the network trains mainly on the simulation-derived dataset with a relatively higher weight on Ldata, giving a stable initialization consistent with numerically verified reservoir behavior. In the refinement stage, the physics-residual weight wp is progressively increased, and where field data exist, Ldata is recomputed on the smaller field dataset — letting the model adapt to observed conditions while the physics term keeps predictions regularized in regions where field coverage is sparse. Collocation points for physics-residual evaluation are resampled adaptively, with denser sampling near the moving dissociation front where gradients are steepest. Optimization starts with a first-order optimizer (Adam) for initial convergence, then switches to a quasi-Newton optimizer (L-BFGS) for fine-tuning.

4.5 Sensitivity analysis module

A variance-based global sensitivity analysis (Sobol indices) together with a local feature-attribution method (SHAP values) forms a post-training analysis module. These methods quantify how much each geological parameter (initial hydrate saturation, absolute permeability, porosity) and operational parameter (bottom-hole pressure drawdown rate, production method) contributes to predicted gas production, which supports interpretability and can inform reservoir management decisions directly. Figure 3 in Section 6 sketches the conceptual form this ranking is expected to take, based on mechanistic understanding of hydrate dissociation reported in prior simulation studies.

5 EVALUATION METHODOLOGY

Model performance is assessed with standard regression metrics root-mean-square error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2) computed for gas production rate, pressure evolution, and hydrate saturation predictions against held-out numerical-simulation cases and, where available, blind field-test data. Physical consistency is checked separately, through the residual magnitude of the governing equations evaluated on a held-out set of collocation points, giving a measure of physics-constraint satisfaction that is independent of the data-fidelity metrics. Computational efficiency is evaluated by comparing inference wall-clock time against that of the full numerical simulator for equivalent forecasting scenarios. Cross-validation (k-fold across simulation cases) and leave-one-well-out validation, where multiple field or pilot datasets are available, are used to assess generalization, with comparisons made against (i) a purely data-driven network of identical architecture trained without the physics-residual term, and (ii) empirical decline-curve baselines.

6 EXPECTED RESULTS AND DISCUSSION

As a proposed framework, the quantitative results depend on implementation against a specific field or synthetic dataset. Even so, based on the formulation in Sections 3–5 and on outcomes reported for analogous physics-informed approaches in subsurface flow, a few qualitative expectations can be laid out.

Table 2: Qualitative comparison of the proposed HPIML framework against conventional numerical simulation, empirical decline-curve analysis, and purely data-driven machine learning.

Criterion	Numerical Simulation	Empirical Decline Curve	Purely Data-Driven ML	Proposed HPIML
Physical consistency	High	Low	Low–Moderate	High
Data requirement	Low (physics-only)	Low	High (labeled data)	Moderate
Computational cost (inference)	High	Very low	Very low	Low
Generalization beyond training range	N/A (first-principles)	Poor	Poor	Expected to be good
Interpretability	High (mechanistic)	Moderate	Low	Moderate–High (via sensitivity module)

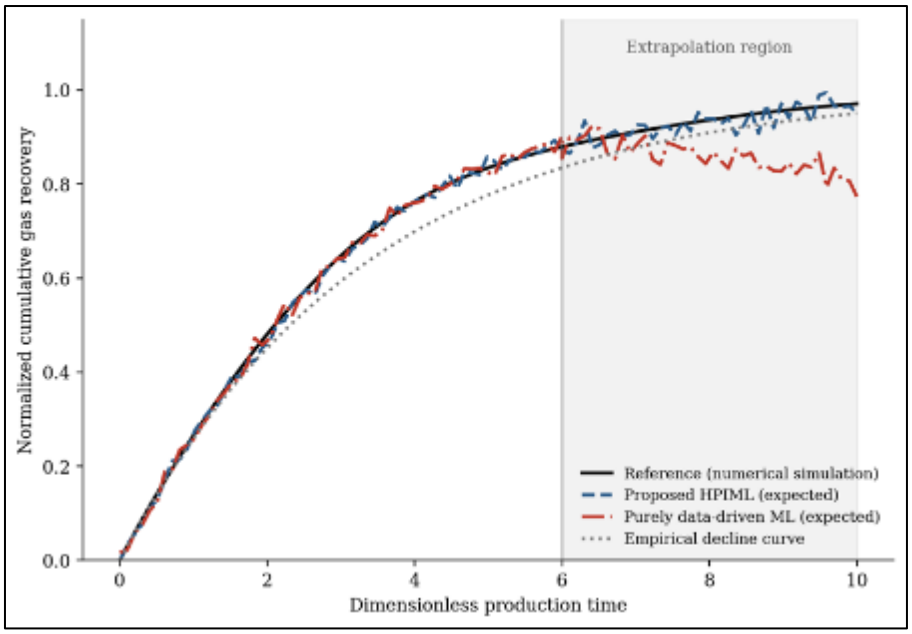

Figure 2: Conceptual illustration of anticipated forecast behavior within the training horizon and under extrapolation, contrasting the expected stability of the physics-constrained model with the expected drift of a purely data-driven model. Curves are schematic and are not derived from an executed simulation or field dataset.

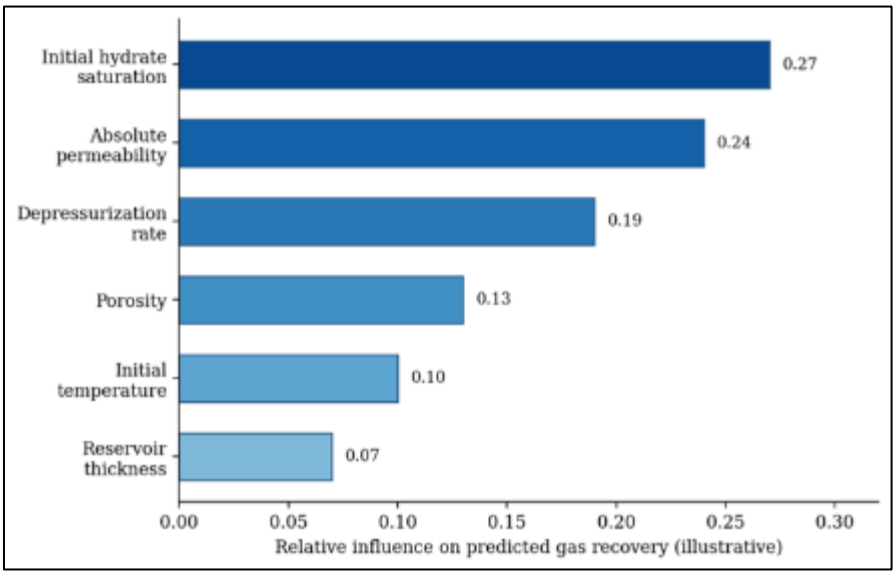

Figure 3: Conceptual ranking of geological and operational parameters expected to emerge from the sensitivity-analysis module, based on mechanistic reasoning rather than computed Sobol/SHAP output.

- Prediction accuracy: embedding conservation and kinetics constraints should cut down error accumulation over long forecast horizons compared with purely data-driven models, especially in parts of the parameter space that are thin on training data.
- Generalization under data scarcity: because the physics-residual term doesn't need labeled outputs, the model should generalize more reliably than purely data-driven baselines when field-labeled data are limited a condition typical of hydrate pilot tests.
- Computational efficiency: once trained, the network should produce forecasts several orders of magnitude faster than a fully coupled numerical simulator, opening the door to near-real-time scenario analysis and integration into iterative reservoir-management workflows.

- Interpretability: the sensitivity-analysis module is expected to flag initial hydrate saturation, absolute permeability, and depressurization rate as the dominant controls on cumulative gas recovery, in line with mechanistic understanding of hydrate dissociation from prior simulation studies.
- Physical consistency: the physics-residual metric should stay low and stable across the prediction domain, including in extrapolation regions, while purely data-driven baselines are expected to show worse and less consistent residual behavior outside the training distribution.

These are expectations, not results, they need to be validated empirically once the framework is implemented and trained on a specific dataset. Section 8 discusses what that validation would involve.

7 IMPLICATIONS FOR RESERVOIR MANAGEMENT AND SUSTAINABLE DEVELOPMENT

A validated HPIML framework would give reservoir engineers a computationally efficient way to weigh alternative production strategies depressurization rate, thermal-stimulation schedules, well spacing under uncertainty, without waiting on repeated full-physics simulation runs. By cutting predictive uncertainty and improving generalization from limited pilot-scale data, the framework could make it easier to move confidently from pilot testing toward commercial-scale hydrate production planning. And because methane is a potent greenhouse gas, more reliable productivity forecasting also supports risk-informed operational design that limits uncontrolled gas release, keeping hydrate development aligned with the broader push toward carbon-conscious energy transition, where unconventional gas resources are managed alongside decarbonization efforts rather than as a substitute for them.

Limitations and Future Work

A few limitations are worth flagging. First, the framework's performance has not yet been demonstrated empirically against field data at commercial scale; validation against additional pilot-test datasets, as they become available, is still needed. Second, the physics-residual term is only as good as the constitutive relationships behind it relative permeability and reactive-surface-area models, for example and those carry their own uncertainty for hydrate-bearing sediments; a reasonable next step would be to treat selected constitutive parameters as learnable quantities within the network, following inverse-PINN approaches. Third, geomechanical coupling sediment compaction, potential wellbore instability is not explicitly represented in the current formulation, and that's a priority extension. Finally, scaling the framework up to three-dimensional, multi-well field configurations will require scalable collocation-sampling and domain-decomposition strategies to keep training cost manageable.

8 CONCLUSION

This paper has proposed a hybrid physics-informed machine learning framework for predicting methane hydrate reservoir productivity, tying together governing mass-, energy-, and hydrate-dissociation-kinetics equations with deep learning models trained on numerical-simulation and field-derived data. Embedding physical constraints directly into the training objective is meant to improve prediction accuracy and generalization relative to purely empirical or purely data-driven approaches, while keeping computational cost well below that of fully coupled numerical simulation. The built-in sensitivity-analysis module adds a layer of interpretability by identifying which geological and operational parameters matter most for gas recovery. Subject to empirical validation, this approach offers a promising route toward forecasting tools that are more efficient, more physically consistent, and more directly useful for decision-making in the sustainable development of methane hydrate resources.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The author acknowledges the support of the Boone Pickens School of Geology, Oklahoma State University, for providing the academic environment and resources that supported the preparation of this manuscript. The author also acknowledges the use of publicly available scientific literature and computational and analytical resources consulted during the development of this review.

Disclosure of conflict of interest

The author(s) declare no conflict of interest.

REFERENCES

[1] Boswell R, Collett TS. Current perspectives on gas hydrate resources. Energy & Environmental Science. 2011;4(4):1206–1215.

[2] Makogon YF. Natural gas hydrates – A promising source of energy. Journal of Natural Gas Science and Engineering. 2010;2(1):49–59.

[3] Moridis GJ, Kowalsky MB, Pruess K. TOUGH+HYDRATE v1.0 User's Manual: A Code for the Simulation of System Behavior in Hydrate-Bearing Geologic Media. Lawrence Berkeley National Laboratory Report LBNL-149E; 2008.

[4] Kim HC, Bishnoi PR, Heidemann RA, Rizvi SSH. Kinetics of methane hydrate decomposition. Chemical Engineering Science. 1987;42(7):1645–1653.

[5] Moridis GJ, Reagan MT. Estimating the natural gas production potential of coarse-grained hydrate accumulations at the Ignik Sikumi field. SPE Journal. 2011.

[6] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics. 2019;378:686–707.

[7] Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. Nature Reviews Physics. 2021;3(6):422–440.

[8] Tartakovsky AM, Marrero CO, Perdikaris P, Tartakovsky GD, Barajas-Solano D. Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. Water Resources Research. 2020;56(5):e2019WR026731.

[9] He QZ, Barajas-Solano D, Tartakovsky G, Tartakovsky AM. Physics-informed neural networks for multiphysics data assimilation with application to subsurface transport. Advances in Water Resources. 2020;141:103610.

[10] Collett TS, Johnson AH, Knapp CC, Boswell R, editors. Natural Gas Hydrates – Energy Resource Potential and Associated Geologic Hazards. AAPG Memoir 89. Tulsa: American Association of Petroleum Geologists; 2009.

[11] Yin Z, Moridis G, Chong ZR, Tan HK, Linga P. Numerical analysis of experimental studies of methane hydrate dissociation induced by depressurization in a sandy porous medium. Applied Energy. 2018;230:444–459.

[12] Merey S, Sinayuc C. Investigation of gas hydrate potential of the Black Sea and modelling of gas production from a hypothetical Class 1 methane hydrate reservoir in the Black Sea conditions. Journal of Natural Gas Science and Engineering. 2016;29:66–79.

[13] Kingma DP, Ba J. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980; 2014.

[14] Sobol IM. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation. 2001;55(1–3):271–280.

[15] Lundberg SM, Lee SI. A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems. 2017;30:4765–4774.