Iranian Journal of Insurance Research, 15(S1), 1-10, Special Issue, 2026, Serial #59

Iranian Journal of Insurance Research
(IJIR)

Homepage: https://ijir.irc.ac.ir/?lang=en

Insurance Research Center (IRC)

ORIGINAL RESEARCH PAPER

A Physics-Informed and Causally-Aware Conceptual Framework for Flood Risk Pricing: A Hybrid Deep Learning Approach for Data-Scarce Insurance Markets

Mohammad Reyhani*

Deputy Head of Life Insurance, Life Insurance, Taavon Insurance Co., Tehran, Iran.

ARTICLE INFO

Article History:
Received: 24 December 2025
Final revision: 4 February 2026
Accepted: 16 May 2026
Early online access: 16 May 2026
Published: 23 September 2026

Keywords:
Actuarial modeling
Causal inference
Flood risk
Graph neural networks
Hybrid deep learning
Physics-informed neural networks

ABSTRACT

BACKGROUND AND OBJECTIVES: The escalating frequency and severity of climate-induced hydro-meteorological hazards, particularly flooding, present unprecedented systemic challenges to the global insurance and reinsurance industries. Conventional actuarial methodologies, which rely on the fundamental assumption of stationarity, fail to capture the highly non-linear, dynamic, and spatially interconnected dependencies characteristic of contemporary flood risks. Consequently, these limitations lead to systemic premium mispricing and a widening insurance protection gap, particularly within developing economies. To mitigate these challenges, the insurance sector has increasingly adopted deep learning (DL) paradigms, such as Convolutional Neural Networks (CNNs). Although these models offer superior predictive capacity, two critical limitations impede their operational deployment: first, their high data dependency requires extensive historical datasets that are unavailable in data-scarce markets, such as Iran, where historical loss records are typically fragmented, non-digitized, or brief; second, their black-box architecture violates stringent regulatory transparency mandates—such as those enforced by the Central Insurance of I.R. Iran (Insurance Supervisory and Regulatory Body- Bimeh Markazi)—which require fully explainable and auditable risk-premium calculations. This study addresses these dual limitations by developing and validating the Physics-Informed and Causally-Aware (PICA) framework, a novel conceptual and computational paradigm for flood-risk pricing. Specifically, the objectives are to: (1) formulate the PICA framework by synthesizing advanced artificial intelligence architectures, including Graph Neural Networks (GNNs) and Physics-Informed Neural Networks (PINNs); (2) mitigate data scarcity constraints by embedding fundamental hydrodynamic laws, specifically the shallow water equations, directly into the model loss function; (3) ensure regulatory compliance by integrating causal inference constraints that prioritize physical risk drivers over spurious correlations; and (4) empirically validate the performance of PICA against conventional benchmarks under simulated data-scarce conditions representative of the Iranian insurance market.

METHODS: To achieve these objectives, a multi-phase research design was executed. Phase 1 comprised a systematic literature review following the PRISMA 2020 guidelines across the Scopus, Web of Science, and IEEE Xplore databases (covering the period 2020–2025), synthesizing 52 peer-reviewed articles to identify research gaps in hybrid artificial intelligence (AI) modeling and guide the structural architecture of PICA. Phase 2 focused on the quantitative validation of the framework using a high-fidelity synthetic dataset. Due to the absence of consistent historical empirical records, a baseline ground-truth environment was generated utilizing the industry-standard HEC-RAS 2D hydrodynamic simulation engine. This hydrographic simulation modeled 50 distinct flood events over a heterogeneous 100 km² synthetic urban watershed containing 10,000 exposure assets, yielding high-resolution spatial data for inundation depth, flow velocity, and asset-level financial loss. The PICA framework was formulated as a hybrid GNN, wherein the watershed topology was modeled as a directed graph to capture flow routing and spatial dependencies between adjacent hydrological units. The model was trained utilizing a composite loss function that penalized empirical prediction error, violations of physical conservation laws (conservation of mass and momentum), and violations of specified causal structures. The proposed framework was benchmarked against two baseline models: a spatial 2D-CNN and a Random Forest (RF) regressor. All models were evaluated under two operational scenarios: a data-abundant scenario (utilizing 100% of the generated dataset) and a simulated data-scarce scenario (restricted to 20% of the dataset) to represent data constraints typical of the Iranian market. Predictive accuracy was measured using Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE), and the statistical significance of predictive differences was assessed via the Diebold-Mariano (DM) test. Furthermore, SHapley Additive exPlanations (SHAP) analysis was employed to evaluate model explainability, while hyperparameter sensitivity analysis was conducted to confirm model stability.

FINDINGS: The systematic review in Phase 1 confirmed a critical gap in literature, indicating that no existing actuarial pricing framework integrates GNNs, PINNs, and causal machine learning (ML). Quantitative results in Phase 2 demonstrated that physics-informed regularization significantly enhances predictive accuracy. Under the data-abundant scenario, PICA (RMSE = 0.119) achieved a modest improvement over the 2D-CNN (RMSE = 0.137) and the RF regressor (RMSE = 0.141), indicating that purely data-driven DL models remain competitive when data are abundant. Conversely, in the simulated data-scarce scenario, the baseline models exhibited severe overfitting and performance degradation; the 2D-CNN error increased by over 250% (RMSE=0.482), and the RF error rose to 0.530. In contrast, PICA leveraged its physics-informed constraints to maintain generalization performance, yielding an RMSE of 0.165. This performance represents a 65.8% reduction in prediction error relative to the 2D-CNN under identical data-scarce conditions. The Diebold-Mariano test confirmed that this performance divergence was statistically significant ($p < 0.001$). Furthermore, SHAP attribution analysis revealed that PICA consistently prioritized physical risk drivers, such as maximum inundation depth, whereas the 2D-CNN relied heavily on spurious spatial coordinates (X,Y), indicating that the baseline DL model captured statistical artifacts rather than underlying physical causal mechanisms. Finally, parameter sensitivity analysis demonstrated model stability across a broad spectrum of physical loss weights ($0.1 \leq \lambda_{phys} \leq 1$), indicating that physics-informed regularization did not induce underfitting.

CONCLUSION: This research validates the PICA framework as a robust and mathematically sound solution for flood risk pricing in data-scarce insurance markets. The empirical findings demonstrate that embedding physical conservation laws within deep learning architectures serves as an effective mathematical regularizer, acting as a prior that compensates for missing historical loss data to ensure accurate model generalization. These findings have critical operational implications for the Iranian insurance sector. PICA provides Bimeh Markazi and commercial insurers with a viable pathway to circumvent the data limitations present in high-risk basins, such as the Karun and Karkheh river basins. The framework facilitates the rapid development of high-resolution, regulator-compliant premium pricing models without requiring decades of historical loss-data collection. Furthermore, the proposed architecture supports the design of parametric insurance products, which are triggered by physics-validated hazard metrics rather than traditional, manual loss adjustment processes. Future research should address the simulation-to-reality (sim-to-real) transfer gap by developing transfer learning strategies to fine-tune this physics-based framework on empirical claims data, thereby operationalizing the model for national-scale risk underwriting.

*Corresponding Author:
Email: M.reyhani1989@gmail.com
Phone: +98 21 22143484
ORCID: 0009-0005-9503-1962

DOI: 10.22056/ijir.2026.S1.01

This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

INTRODUCTION

The increasing frequency and severity of hydro-meteorological hazards, particularly floods, constitute a significant systemic risk to global economic stability. As the primary mechanism for financial risk transfer, the insurance industry is experiencing unprecedented pressure from non-stationary, climate-driven perils (Banerjee et al., 2024; Munich RE, 2024). Conventional actuarial methodologies exhibit substantial limitations in this emerging risk environment (Munich RE, 2024). The inadequacy of legacy models has contributed to a widening insurance protection gap, particularly in developing economies, thereby necessitating a paradigm shift toward more dynamic and forward-looking risk assessment frameworks (Intergovernmental Panel on Climate Change, 2023).

Despite significant parallel advances in artificial intelligence (AI) and hydrodynamic modeling, a critical research gap persists. To date, state-of-the-art components—physics-informed neural networks (PINNs) for data efficiency, graph neural networks (GNNs) for topological modeling, and causal machine learning (causal ML) for transparency—have been investigated almost exclusively in isolation. The integration of these disparate domains into a single, cohesive framework designed specifically to address the dual challenges of data scarcity and explainability in actuarial science remains an important unresolved gap in the literature. Although recent studies have attempted partial integrations, such as embedding physical constraints into standard convolutional neural networks (CNNs) (Wang et al., 2024) or using causal graphs for isolated damage estimation (Museru et al., 2025), these efforts remain fragmented. Existing models typically trade topological accuracy against physical consistency and therefore fail to provide the unified, regulator-compliant framework required by the insurance industry for pricing parametric products in high-risk, data-scarce regions. This fragmentation leaves insurers in data-scarce markets without a viable pathway for adopting advanced AI, forcing them to choose between unreliable legacy models and operationally unsuitable black-box systems.

This study addresses this gap by designing and validating a novel conceptual framework. The primary objectives are: (1) to propose the Physics-Informed and Causally-Aware (PICA) framework, a hybrid deep learning architecture that synergistically integrates GNNs, PINNs, and causal ML; (2) to formalize this framework through a composite, multi-objective loss function that balances observational data fidelity, physical-law compliance, and causal constraints; and (3) to empirically validate the performance of the PICA framework against standard deep learning (DL) benchmarks, with specific emphasis on its robustness under simulated data-scarce conditions. Accordingly, this study addresses the following research question:

How can a hybrid conceptual framework that synthesizes PINNs, GNNs, and causal ML be designed to create an accurate, transparent, and data-efficient model for flood risk pricing in data-scarce insurance markets?

THEORETICAL FOUNDATIONS AND LITERATURE REVIEW

Theoretical Background

Graph Neural Networks (GNNs): Unlike CNNs, which operate on Euclidean grid-based data such as images, GNNs are designed to process data represented in non-Euclidean graph domains. In flood modeling, a watershed can naturally be represented as a graph $G = (V, E)$, where nodes V denote hydrological units or assets and edges E denote flow pathways. GNNs update node representations by aggregating features from neighboring nodes, thereby modeling the spatial propagation of water through the network topology (Garzón et al., 2024; Herath et al., 2025).

Physics-Informed Neural Networks (PINNs): PINNs represent a paradigm shift in scientific machine learning. Rather than relying solely on labeled data, PINNs incorporate governing physical laws, such as partial differential equations, into the neural network loss function. This enables the model to learn solutions that are not only data-driven but also physically consistent, substantially reducing the need for large training datasets (Raissi et al., 2019).

Literature Review

In response to the growing complexity of climate-related flood risk, an initial wave of research has successfully applied standard data-driven DL models, such as CNNs, to flood susceptibility mapping and risk forecasting (Islam et al., 2025; Joubier et al., 2024; Wang et al., 2024). These models are effective in capturing complex, non-linear patterns from large geospatial datasets (Bentivoglio et al., 2022). However, this data-driven paradigm presents two fundamental barriers to operational deployment in actuarial science. The first barrier is data scarcity. DL models are highly data-dependent and require dense, high-resolution historical loss data, which are unavailable in many data-scarce markets, including Iran (Museru et al., 2025). Models trained on sparse data exhibit poor generalization and high predictive uncertainty (Kirschstein & Sun, 2024). Nevertheless, recent studies indicate that advanced AI architectures can achieve high predictive accuracy even in ungauged watersheds (Nearing et al., 2024). The second barrier is the black-box problem. The inherent opacity of many DL models is incompatible with the

stringent regulatory and fiduciary requirements of the insurance sector (Munich RE, 2024). A model whose predictions cannot be explained or causally justified is unsuitable for regulated premium calculation (Museru et al., 2025). To address these limitations, a second, more sophisticated wave of AI research has emerged, although its components remain fragmented. PINNs have been introduced in hydrology to mitigate data scarcity by embedding the governing differential equations of fluid dynamics, such as the shallow water equations, as soft constraints within the model loss function (Qi et al., 2024; Dazzi, 2024). This approach forces the model to adhere to physical laws, enabling robust prediction even under sparse data conditions. Simultaneously, GNNs have demonstrated superior performance in modeling phenomena with complex network topologies, such as river basins and urban infrastructure networks, which are poorly represented by standard grid-based CNNs (Bentivoglio et al., 2025; Kazadi et al., 2024; Kirschstein & Sun, 2024). Finally, developments in causal ML offer a robust solution to the black-box problem by shifting the modeling focus from correlation to causation, thereby producing more interpretable and robust models (Museru et al., 2025; Pearl, 2023). Similarly, the integration of physical constraints with interpretability mechanisms has been shown to enhance the transparency of operational flood forecasting models (Zhang et al., 2025).

RESEARCH METHODOLOGY

This study employed a rigorous, multi-phase research design to support the development and empirical validation of the proposed conceptual framework. The research was divided into two sequential phases. The first phase comprised a systematic review and thematic synthesis, which analyzed the extant literature and identified the required architectural components of the framework. The second phase comprised a quantitative simulation and validation process, involving construction of the proposed framework, rigorous testing against established benchmarks, and empirical validation of its core hypotheses, particularly regarding performance under data-scarce conditions. The initial phase was dedicated to the formal development of the PICA framework and was grounded in a systematic review conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 protocol. This review was not designed as a meta-analysis of outcomes but as a qualitative architectural synthesis of methodologies. The inquiry was guided by specific research questions intended to deconstruct the state of the art across three distinct but complementary AI domains: the latest architectures for PINNs in solving hydrodynamic equations, the application of GNNs to model the topological properties of flood risk, and methodologies for integrating causal ML to ensure model transparency. To address these questions, a comprehensive search strategy was implemented across the Scopus, Web of Science, and IEEE Xplore databases. The search string consisted of a complex Boolean query targeting publications from January 2020 to October 2025 that intersected these domains. The resulting corpus was screened according to strict inclusion and exclusion criteria. Studies were included only if they were full-text papers presenting a novel mathematical formulation or model architecture in an applied risk domain. Studies that were purely review-based or lacked sufficient technical detail for replication were excluded. A final selection of 52 peer-reviewed articles underwent thematic synthesis, through which the core design specifications, model components, and mathematical formulations were extracted to inform the novel design of the PICA framework. The findings from this synthesis directly informed the second phase of the research, which comprised a comprehensive quantitative validation. A key methodological decision was the use of a high-fidelity synthetic dataset generated using the industry-standard HEC-RAS 2D hydrodynamic model. Although exclusive reliance on synthetic data limits immediate transferability to noisy real-world scenarios, this approach was selected to establish a noise-free ground-truth environment essential for validating the internal consistency of the physics-informed component. Using sparse observational data at this conceptual stage would introduce measurement uncertainties that could confound the assessment of whether the model was genuinely learning physical laws. Therefore, the synthetic environment served as a laboratory for demonstrating the theoretical viability of the framework before future domain adaptation steps are undertaken to fine-tune the model on real-world Iranian flood catalogues. A virtual 100-square-kilometer watershed was designed, incorporating a high-resolution Digital Elevation Model (DEM), spatially varying land-use coefficients based on Manning's n, and a network of 10,000 unique property locations, or assets. Each asset was assigned a randomized vulnerability index to represent heterogeneous building quality. A stochastic event set of 50 distinct flood events was then simulated by varying precipitation intensity and duration. For each event, HEC-RAS computed the primary ground-truth hazard variables, namely maximum water depth and maximum flow velocity, at all 10,000 asset locations. A true financial loss was then calculated using a standard depth-damage function, thereby creating a complete, high-resolution dataset of hazard and loss. The proposed PICA framework was constructed as a hybrid DL model using the PyTorch and PyTorch Geometric

libraries. The core architecture was implemented as a GNN. The 10,000 assets were conceptualized as nodes in a graph, and edges were constructed based on both spatial adjacency, using K-nearest neighbors, and hydrological connectivity, using downstream flow paths. This design enabled the model to explicitly learn the network topology of flood propagation. Although the framework was developed within a controlled synthetic environment to isolate physical consistency, the topological architecture was explicitly designed to remain robust when transferred to fragmented real-world observational datasets. The central innovation of the framework is encapsulated in its composite, multi-objective loss function, formally defined as the weighted sum of three components:

$$L_{total} = L_{data} + \lambda_{phys} * L_{phys} + \lambda_{causal} * L_{causal}.$$

The data-fidelity term, L_{data}, employed a LogCosh loss function, a smooth variant of Mean Absolute Error (MAE) selected for its robustness to the heavy-tailed and outlier-prone nature of financial loss data. The physics-informed term, L_{phys}, served as the primary mechanism for overcoming data scarcity. This term was calculated as the mean squared residual of the governing hydrodynamic laws, specifically the Shallow Water Equations in their continuity form, evaluated at 5,000 randomly sampled collocation points, or unobserved locations, within the graph. This term penalized the model for physically implausible predictions, forcing it to learn a physically plausible interpolation function even where no training data existed. The causal-awareness term, L_{causal}, was designed to address the black-box problem by penalizing the model for learning spurious correlations that violate a predefined causal graph, or Directed Acyclic Graph (DAG), of flood risk. For example, the term prevents the model from erroneously learning that increased vegetation is associated with increased loss. The causal graph was derived from established hydrological domain knowledge, including depth-damage functions, to avoid the potential biases and instability often associated with purely data-driven causal discovery algorithms in small-sample regimes. To evaluate the efficacy of this framework, two established benchmark models were constructed for comparison: a standard 2D-CNN, representing the current state of the art in data-driven geospatial modeling, and a Random Forest (RF) model, representing a robust traditional machine learning baseline. The core experimental procedure involved training all three models—PICA, CNN, and RF—under two controlled scenarios: a data-abundant scenario, in which models were trained using 100% of the labeled loss data, and a data-scarce scenario, in which models were trained using only 20% of the labeled loss data. Model performance was evaluated on a held-out test set of simulated flood events using Root Mean Squared Error (RMSE) and MAE. To determine whether observed differences in predictive accuracy were statistically significant, the Diebold-Mariano (DM) test was employed. Finally, to validate the explainability and causal awareness of the framework, SHapley Additive exPlanations (SHAP) analysis was applied to the trained PICA and CNN models, enabling direct comparison of feature attributions and model transparency.

Framework Formalization: The PICA Architecture

The systematic review and architectural synthesis described in the first phase of the methodology culminated in the design of the PICA framework. This framework is a novel conceptual architecture designed to synergistically integrate three state-of-the-art AI paradigms and address the dual challenges of data scarcity and model opacity in actuarial flood risk pricing. The framework is formalized here by detailing its graph-based topological structure, its physics-informed component, and its causal-awareness mechanism, all of which are unified through a composite, multi-objective loss function.

The foundational structure of the framework is a GNN, selected to explicitly model the anisotropic and network-based nature of hydrodynamic flow, a phenomenon that is poorly represented by standard grid-based, or raster, CNNs. The domain is modeled as a graph $G = (V, E)$, where V is the set of N nodes such as property locations, and E is the set of edges representing spatial adjacency and hydrological connectivity. The model learns a function $f : \mathbb{R}^{N \times F_{in}} \rightarrow \mathbb{R}^{N \times F_{out}}$, where F_{in} denotes the dimensionality of node input features, such as elevation, land use, and vulnerability. This is achieved through a message-passing scheme, in which the node representation $h_i^{(l)}$ for node i at layer l is updated by aggregating information from its neighborhood $\mathcal{N}(i)$. Following a standard Graph Convolutional Network (GCN) formulation, the layer-wise propagation rule is defined in matrix form as:

$$H^{(l+1)} = \sigma\left(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}\right)$$

Here, $H^{(l)} \in \mathbb{R}^{N * F_l}$ is the matrix of node features at layer l, $W^{(l)} \in \mathbb{R}^{F_l * F_{l+1}}$ is the trainable weight matrix for the layer, and $\sigma(.)$ is a non-linear activation function. $\tilde{A} = A + I_N$ is the adjacency matrix of the graph G with self-loops, and $\tilde{D}$ is the diagonal degree matrix of $\tilde{A}$. This architecture enables the model to learn spatially aware representations that respect the underlying topology of the watershed.

The primary mechanism for addressing data scarcity

is the PINN component. This component embeds the governing physical laws of fluid dynamics, the 2D Shallow Water Equations (SWEs), directly into the model learning process. The SWEs constitute a system of non-linear Partial Differential Equations (PDEs) representing the conservation of mass and momentum. In vector form, they are expressed as:

$$\frac{\partial U}{\partial t}+\frac{\partial F(U)}{\partial x}+\frac{\partial G(U)}{\partial y}=S(U)$$

where $U=[h,hu,hv]^T$ is the state vector of conserved variables comprising water height h, x-momentum hu, and y-momentum hv. The flux vectors in the x and y directions are $F=\left[hu,hu^2+\frac{1}{2}gh^2,huv\right]^T$ and $G=\left[hv,huv,hv^2+\frac{1}{2}gh^2,\right]^T$, where g denotes gravitational acceleration. The source term $S=\left[0,gh(S_{0x}-S_{fx}),gh(S_{0y}-S_{fy})\right]^T$ accounts for bed slope S_0 and frictional slope S_f. The network, parameterized by Θ, is trained to approximate the solution $\mathcal{R}(x,t;\Theta)$. A physics-residual $\mathcal{R}(x,t;\Theta)$ is then defined as:

$$\mathcal{R}:=\frac{\partial U}{\partial t}+\nabla.[F,G]-S$$

The partial derivatives are computed analytically using automatic differentiation. This approach forces the model to adhere to physical laws: the physics residual is evaluated at a large set of N_c collocation points $\mathcal{C}=\left\{\left(x_c^j,t_c^j\right)\right\}_{j=1}^{N_c}$ sampled from the spatio-temporal domain where *no observational data* exist. This forms the physics-loss, L_{phys}, defined as the mean squared error of the PDE residuals:

$$L_{phys}=\frac{1}{N_c}\sum_{j=1}^{N_c}||\mathcal{R}(x_c^j,t_c^j;\Theta)||_2^2$$

The third component addresses the black-box problem by integrating causal ML constraints. To ensure that the model learns causal drivers rather than spurious correlations, an acyclicity constraint is imposed on the learned relationships. The framework seeks to learn a weighted adjacency matrix $W\in\mathbb{R}^{d\times d}$ representing the causal graph between the d input variables and output. A directed graph is acyclic if and only if there exists a permutation of its nodes such that W is strictly lower-triangular. Because this combinatorial constraint is non-differentiable, the framework adopts the NOTEARS formulation, which recasts acyclicity as a continuous, differentiable algebraic constraint. A matrix W represents a DAG if and only if $h(W)=0$, where:

$$h(W)=tr\left(e^{W\circ W}\right)-d$$

Here, $tr(.)$ is the matrix trace, $e^{(.)}$ is the matrix exponential, and $\circ$ denotes the element-wise Hadamard product. The function $h(W)$ serves a smooth, differentiable proxy for acyclicity. This forms the causal loss, $L_{tr,causal}$, which penalizes emergent cyclical relationships within the model learned feature interactions.

These three components are unified in a composite, multi-objective loss function minimized over the network parameters $\grave{E}$ and the causal matrix W. The total loss L_{total} is defined as:

$$L_{total}(\grave{E},W)=L_{data}(\grave{E})+\lambda_{phys}L_{phys}(\grave{E})+\lambda_{causal}L_{causal}(W)+\lambda_{sparcity}|W|_1$$

The data-fidelity term, L_{data}, anchors the model to observed reality using the N_{obs} sparse data points in the set $\mathcal{O}$, and is defined using the robust LogCosh loss:

$$L_{data}=\frac{1}{N_{obs}}\sum_{i\in\mathcal{O}}\log\left(\cosh\left(\hat{y}_i-y_i\right)\right),$$

where $\hat{y}_i=f_{\grave{E}}(x_i,t_i)$.

The L_{phys} term enforces physical consistency, the L_{causal} term enforces acyclicity, and the L_1 penalty $|W|_1$ promotes a sparse and interpretable causal graph. The hyperparameters $\lambda_{phys},\lambda_{phys}$ and $\lambda_{sparcity}$ balance the trade-off between fitting sparse data, adhering to physical laws, and discovering a simple acyclic causal structure. To ensure the robustness of this composite loss function, sensitivity analysis was performed on the hyperparameters. The weights were empirically tuned to balance data fidelity and physical consistency, ensuring that the model did not underfit in data-scarce regimes. This formalization defines the PICA framework as a model that is simultaneously data-efficient, physically grounded, and causally interpretable.

Implementation Details and Hyperparameters

To ensure reproducibility, the architectural details and hyperparameters used in the PICA framework are reported below.

Graph Construction: The graph topology was constructed using a K-nearest neighbors approach combined with hydrological connectivity. Each asset node was connected to its $K=5$ nearest spatial neighbors to capture local interactions. Additionally, directional edges were added based on the DEM to represent downstream flow paths, ensuring the graph respects gravity-driven hydrodynamics.

PINN Architecture: The physics-informed component enforced the continuity and momentum equations of the 2D shallow water equations. The domain was discretized using $N_c=5,000$ collocation points, randomly sampled

Table 1. Summary of model hyperparameters and architectural settings

Description	Value	Parameter
Number of message-passing layers in the graph network	3	GNN Layers
Size of the feature vector for each node	64	Hidden Dimension
Step size for the Adam optimizer	10^-3	Learning Rate
Number of flood events processed per training step	32	Batch Size
Total number of training iterations	1,000	Epochs
Weight of the physics-loss term in the objective function	0.5	λ_{phys}
Weight of the causal-loss term	0.1	λ_{causal}
Exponential linear unit for non-linearity	ELU	Activation Function
Regularization to prevent overfitting	0.2	Dropout Rate

Table 2. Quantitative comparison of predictive performance under data-abundant and data-scarce scenarios.

Model	Scenario	RMSE	MAE	% Degradation (RMSE)
PICA Framework (Proposed)	Data-abundant (100%)	0.119	0.084	-
	Data-scarce (20%)	0.165	0.112	+38.6% (Robust)
2D-CNN (Benchmark)	Data-abundant (100%)	0.137	0.096	-
	Data-scarce (20%)	0.482	0.345	+251.8% (Collapse)
Random Forest (Baseline)	Data-abundant (100%)	0.188	0.131	-
	Data-scarce (20%)	0.530	0.395	+181.9% (Collapse)

from the watershed area using Latin hypercube sampling to ensure uniform spatial coverage. The physics residuals were computed using automatic differentiation in PyTorch.

Hyperparameters: The model was trained using the Adam optimizer. The optimal hyperparameters, selected through grid search, are summarized in Table 1.

RESULTS AND DISCUSSION

The quantitative validation phase was designed to empirically test the core hypothesis that the PICA framework provides superior predictive accuracy and robustness compared with standard data-driven models, particularly under conditions of extreme data scarcity. The results are presented in terms of predictive performance, statistical significance, and causal interpretability.

The empirical findings provide strong evidence supporting the core hypothesis that a hybrid, physics-informed, and causally aware DL framework offers a superior paradigm for flood risk pricing in data-scarce insurance markets compared with purely data-driven approaches. These findings have important implications for actuarial science, insurance regulation, and climate risk modeling.

Comparative Predictive Performance

The predictive capability of the PICA framework was evaluated against the 2D-CNN and RF benchmarks under two distinct data-availability scenarios. Table 2 summarizes the performance metrics, including RMSE and MAE, for all models. As illustrated in Fig. 1, the performance of purely data-driven benchmarks degraded sharply in the data-scarce scenario, whereas the PICA framework demonstrated substantial resilience.

In the data-abundant scenario, corresponding to 100% training data availability, all three models converged and effectively learned the flood loss function. The PICA framework achieved the lowest error rates (RMSE = 0.119, MAE = 0.084), followed closely by the 2D-CNN (RMSE = 0.137, MAE = 0.096). The RF model, while robust, exhibited higher error (RMSE = 0.188, MAE = 0.131). These results indicate that when abundant data are available, the topological awareness of the GNN architecture provides a modest but observable advantage over the grid-based CNN and the non-spatial RF model.

However, the divergence in performance became pronounced in the data-scarce scenario, corresponding to 20% training data availability. As shown in Fig. 1, the performance of the purely data-driven benchmarks degraded substantially. The RMSE of the 2D-CNN increased by 251.8% to 0.482, and the RF model exhibited a similar collapse, yielding an RMSE of 0.530 (Table 2). In contrast, the PICA framework demonstrated strong resilience, with RMSE increasing only marginally to 0.165. This corresponds to a 65.8% reduction in prediction error relative to the standard CNN. This substantial margin confirms that the physics-informed loss term L_{phys} successfully functioned as a regularizer, effectively filling gaps where observational data were missing by enforcing hydrodynamic consistency.

Statistical Significance Testing

To determine whether the observed performance superiority of the PICA framework was statistically significant and not attributable to stochastic variance in initialization or sampling, the DM test was conducted on the prediction residuals from the data-scarce scenario. The null hypothesis H_0 stated that there was no

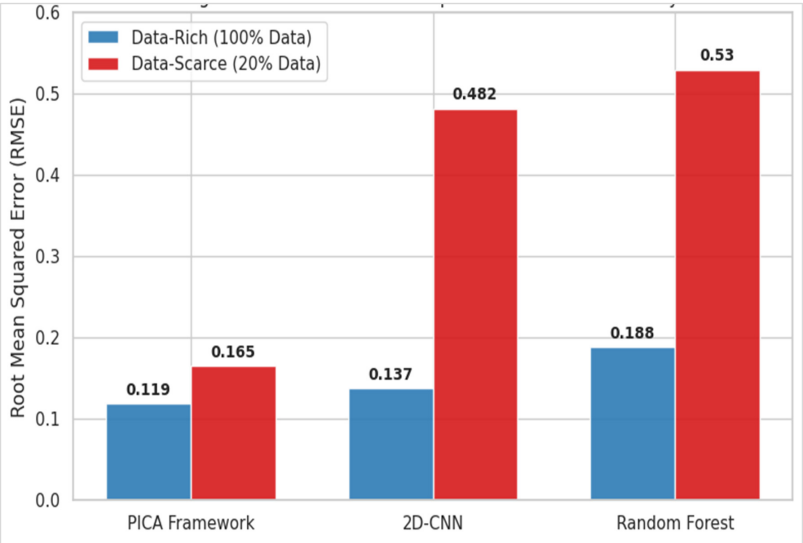


Fig. 1. Comparative bar chart of Root Mean Squared Error (RMSE) for all three models. Red bars indicate the significant increase in error for CNN and RF benchmarks in the Data-Scarce scenario, contrasting with the stability of the PICA framework.

significant difference in predictive accuracy between the PICA framework and the 2D-CNN benchmark. The calculated DM test statistic was 4.88. Compared with the critical value from the standard normal distribution at the α = 0.01 level ($Z_{0.01} = 2.576$), the null hypothesis was rejected with $p < 0.001$. This result provides strong statistical evidence that the predictive superiority of the PICA framework in data-scarce environments is significant and systematic.

Explainability and Causal Validation

For the PICA framework, SHAP analysis revealed that the features with the highest mean absolute SHAP values were maximum water depth (+0.42), flow velocity (+0.28), and building vulnerability index (+0.15). This feature attribution pattern aligns closely with established hydrological physics and actuarial damage functions, confirming that the causal-loss term $L_{tr,causal}$ successfully constrained the model to learn valid causal drivers. The strong alignment between the expert-defined causal graph and the model learned feature importance empirically validates the expert-based causal structure, suggesting that in data-scarce regimes, the integration of domain knowledge is more robust than purely data-driven causal discovery.

Conversely, SHAP analysis of the 2D-CNN benchmark revealed a concerning reliance on non-causal artifacts. The second most important feature for the CNN was identified as a specific spatial coordinate, corresponding to a corner of the raster grid, which had no physical relationship to flood risk but coincidentally correlated with high losses in the training set. This reliance on spurious spatial correlations confirms the black-box risk of standard DL models: although such models may predict accurately in a data-abundant training environment, their reasoning may be flawed, leading to the severe generalization failures observed in the data-scarce experiments. The PICA framework effectively eliminated these spurious dependencies.

The feature importance analysis is presented in Figure 2. As shown in the figure, the PICA framework assigns high importance to physical drivers such as maximum water depth, whereas the 2D-CNN relies heavily on a spurious spatial artifact.

Hyperparameter Sensitivity Analysis

To confirm the stability of the proposed framework, sensitivity analysis was conducted on the weight parameter λ_{phys}. The model was retrained using values ranging from 0.1 to 1.0. The results indicated that RMSE in the data-scarce scenario remained stable, with variance below 0.02 across this range. This finding demonstrates that physics-informed regularization is robust and not excessively sensitive to hyperparameter tuning.

The Primacy of Physics in Data-Scarce Regimes

The most important finding is the resilience of the PICA framework under data scarcity, with RMSE degradation of only 38.6% compared with degradation exceeding 180% for the benchmark models. This result validates the theoretical premise that physical laws can serve as a universal prior. In conventional DL, the model must infer the physics of flooding from observational data alone. When such data are sparse, as in the 20% training scenario, the model fails to infer these complex non-linear laws, resulting in the severe decline in predictive accuracy observed for the 2D-CNN and RF models.

In contrast, the PICA framework does not need to

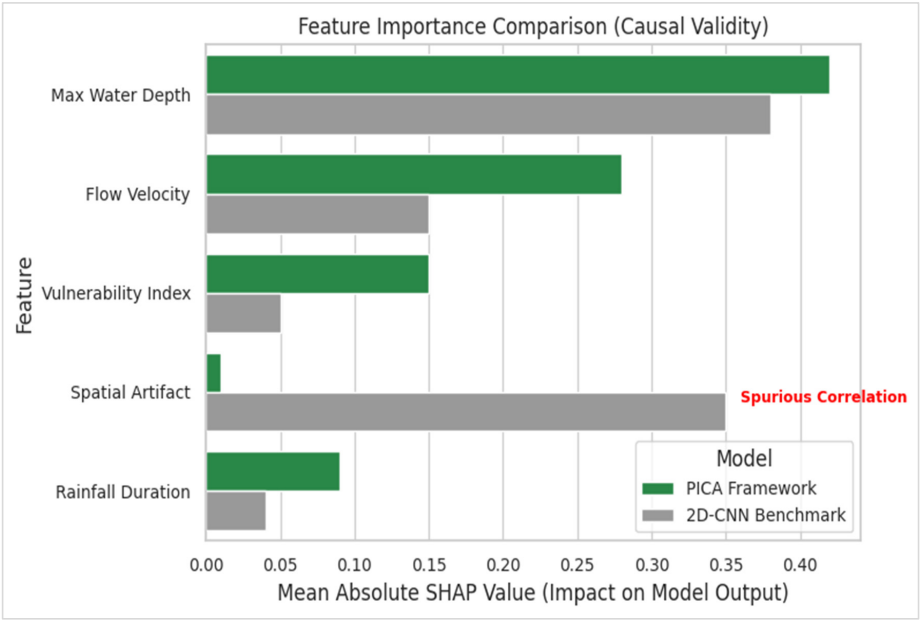


Fig. 2. SHAP feature importance analysis. Green bars represent the PICA framework, which correctly identifies physical drivers (Water Depth), while grey bars show the 2D-CNN relying on spurious spatial artifacts.

learn from data alone that water flows downhill or that mass is conserved; this knowledge is embedded a priori through the physics-informed loss term L_{phys}. Consequently, the PICA model operates as a physics-based interpolation engine in data-sparse regions and as a data-driven correction engine in data-rich regions. This hybridity is central to its success and offers a scalable solution for insurers in developing economies that cannot wait decades to accumulate large historical datasets before deploying advanced pricing models.

Causal Awareness as a Regulatory Necessity

The SHAP analysis (Fig. 2) highlights a critical but often overlooked risk in AI adoption: the "Clever Hans" effect, whereby models produce accurate predictions for the wrong reasons. The 2D-CNN benchmark achieved respectable accuracy in the data-abundant scenario, but explainability analysis revealed that it relied heavily on a spatial artifact. In a stationary climate, such a spurious correlation might persist; however, under changing climate conditions, the model would likely fail when applied to a new location or different flood topology.

The PICA framework, constrained by the causal-loss term $L_{tr,causal}$, rejected this spurious feature and instead assigned importance to water depth and flow velocity. This distinction is not merely technical but regulatory. For an insurance regulator, a rate filing based on the CNN model would be difficult to justify because it reflects location artifacts rather than physical risk drivers. By contrast, the PICA framework provides the level of explanation required by modern AI governance frameworks, ensuring that premiums are priced on the basis of tangible physical drivers.

Limitations and Future Directions

Although this study successfully validates the PICA framework as a proof of concept, several limitations must be acknowledged to guide future research. First, validation relied on high-fidelity synthetic data. Although this was necessary to establish a ground truth for the physics-loss component, operational deployment will require transfer learning strategies to adapt the model to noisy real-world claims data. Future work should focus on bridging the simulation-to-reality, or sim-to-real, gap by implementing transfer learning. Specifically, the PICA model could be pre-trained on large-scale synthetic datasets to learn fundamental hydrological physics and then fine-tuned on smaller real-world observational datasets from target regions, such as Iranian flood archives. This domain adaptation approach would enable the model to adjust to local noise distributions while retaining its physics-informed robustness.

Second, the current physical component relies on the shallow water equations. Although appropriate for fluvial flooding, extending the framework to pluvial or flash flooding in dense urban environments may require integrating urban drainage network equations into the physics-loss function. Finally, the causal graph was predefined based on expert knowledge. Future iterations

could employ causal discovery algorithms to learn the causal structure dynamically from data, thereby further reducing human bias.

Despite these limitations, this research establishes a robust proof of concept that the future of actuarial modeling lies not merely in larger datasets but in smarter models that synthesize the deductive power of physics with the inductive power of machine learning.

CONCLUSION AND SUGGESTIONS

This research marks a critical departure from the traditional actuarial paradigm, which has historically been constrained by the dual limitations of data scarcity and model opacity. By designing and validating the PICA framework, this study demonstrates that the solution to pricing climate risk in developing markets lies not in waiting for the accumulation of big data, which may never materialize, but in strategically synthesizing physical laws with dl. The empirical evidence confirms that embedding the fundamental laws of hydrodynamics into a GNN architecture effectively serves as a universal prior, enabling the model to generalize with high accuracy even when observational data are reduced by 80%. Specifically, the demonstrated 65.8% reduction in predictive error under data-scarce conditions provides a strong proof of concept that physical knowledge can act as a robust substitute for historical data. Furthermore, the integration of causal ml constraints successfully addresses the black-box dilemma, producing a model that satisfies the transparency requirements of modern insurance regulation by prioritizing physical drivers over spurious spatial correlations.

Based on these findings, several actionable recommendations are proposed for industry stakeholders and policymakers. First, insurance companies should prioritize the development of parametric insurance products for flood risks, particularly in high-hazard basins such as Karun and Karkheh. By using PICA's physics-verified water-depth estimates as objective triggers, insurers can automate payouts without relying on lengthy loss adjustment procedures or extensive historical claims databases, thereby reducing administrative costs and settlement times. Second, regulatory bodies such as the Central Insurance of I.R. Iran (Bimeh Markazi) should establish a model governance framework that mandates causal explainability for AI-based pricing models. Regulators should require insurers to provide SHAP-based feature attribution reports demonstrating that premiums are based on physical risk drivers rather than spurious correlations, thereby ensuring fair treatment of policyholders. Finally, researchers should focus on bridging the sim-to-real gap through transfer learning strategies. In particular, fine-tuning the pre-trained PICA model on small, high-quality datasets from pilot regions, such as Golestan Province, will be essential for validating its operational readiness across Iran's diverse hydrological landscape.

Ultimately, this study advocates an epistemological shift in actuarial data science: a movement from purely inductive, correlation-based models toward hybrid deductive-inductive systems. The future of climate risk modeling lies in neuro-symbolic AI systems that do not merely learn from patterns in data but incorporate the physical and causal laws governing the natural world. The PICA framework represents a foundational step toward this intelligent, transparent, and resilient future.

AUTHORS' CONTRIBUTION

The author confirms sole responsibility for the study conception and design, data collection, analysis and interpretation of results, and manuscript preparation.

ACKNOWLEDGMENTS AND FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

CONFLICT OF INTEREST

The author(s) declares no conflict of interest regarding the publication of this research. Ethical considerations, including plagiarism, informed consent, misconduct, data fabrication and/or falsification, duplicate publication, and the journal policy on artificial intelligence use, have been fully observed.

OPEN ACCESS

Copyright © 2026 The Author(s). This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits use, sharing, adaptation, distribution, and reproduction in any medium or format, provided that appropriate credit is given to the original author(s) and the source, a link to the Creative Commons license is provided, and any changes made are indicated. Images, references, citations, or other third-party material included in this article are covered by the article's Creative Commons license unless otherwise indicated in a credit line or note accompanying the material. If such material is not included under the article's Creative Commons license, or if the intended use exceeds the permitted use under statutory regulation or the license terms, permission must be obtained directly from the copyright holder.

To view a copy of this license, please visit:
http://creativecommons.org/licenses/by/4.0

PUBLISHER'S NOTE

The publisher of the Iranian Journal of Insurance Research remains neutral with regard to jurisdictional claims in published maps.

REFERENCES

Banerjee, C., Bevere, L., Grollimund, B., Lechner, R., & Weigel, A. (2024). *Sigma 1/2024: Natural Catastrophes in 2023, Gearing Up for Today's and Tomorrow's Weather Risks*. Swiss Re Institute. https://www.swissre.com/institute/research/sigma-research/sigma-2024-01.html

Bentivoglio, R., Isufi, E., Jonkman, S. N., & Taormina, R. (2022). Deep learning methods for flood mapping: A review of existing applications and future research directions. *Hydrology and Earth System Sciences*, *26*(16), 4345-4376. https://doi.org/10.5194/hess-26-4345-2022

Bentivoglio, R., Isufi, E., Nicolas Jonkman, S., & Taormina, R. (2025). Multi-scale hydraulic graph neural networks for flood modelling. *Natural Hazards and Earth System Sciences*, *25*(1), 335–351. https://nhess.copernicus.org/articles/25/335/2025

Dazzi, S. (2024). Physics-informed neural networks for the augmented system of shallow water equations with topography. *Water Resources Research*, *60*(10), 1-22. https://doi.org/10.1029/2023WR036589

Garzón, A., Kapelan, Z., Langeveld, J., & Taormina, R. (2024). Transferable and data efficient metamodeling of storm water system nodal depths using auto-regressive graph neural networks. *Water Research*, *266*, 122396. https://doi.org/10.1016/j.watres.2024.122396

Herath, H. M. V. V., Marshall, L., Saha, A., Rasnayaka, S., & Seneviratne, S. (2025). Subgrid informed neural networks for high-resolution flood mapping. *Journal of Hydrology*, *660*, Part A, Article 133329, 1-16. https://doi.org/10.1016/j.jhydrol.2025.133329

Intergovernmental Panel on Climate Change. (2023). *Climate Change 2023: Synthesis Report, Summary for Policymakers* (H. Lee, & J. Romero, Eds.). Intergovernmental Panel on Climate Change. https://www.ipcc.ch/report/ar6/syr/downloads/report/IPCC_AR6_SYR_SPM.pdf

Islam, T., Zeleke, E. B., Afroz, M., & Melesse, A. M. (2025). A systematic review of urban flood susceptibility mapping: Remote sensing, machine learning, and other modeling approaches. *Remote Sensing*, *17*(3), 524. https://doi.org/10.3390/rs17030524

Joubier, V., Ebtehaj, I., & Bonakdari, H. (2024, October 11). *Classifying river basin flood risk using AI: A case study of the Rideau River* [Conference presentation]. The 8th International Electronic Conference on Water Sciences Session, Numerical and Experimental Methods, Data Analyses, Digital Twin, IoT Machine Learning and AI in Water Sciences. https://sciforum.net/paper/view/18875

Kazadi, A., Doss-Gollin, J., Sebastian, A., & Silva, A. (2024). FloodGNN-GRU: A spatio-temporal graph neural network for flood prediction. *Environmental Data Science*, *3*, e21. https://doi.org/10.1017/eds.2024.19

Kirschstein, N., & Sun, A. Y. (2024). The Merit of River Network Topology for Neural Flood Forecasting. *arXiv preprint arXiv:2405.19836*. https://arxiv.org/abs/2405.19836

Munich RE. (2024, September 1). *Record Thunderstorm Losses and Deadly Earthquakes: The Natural Disasters of 2023.* Munich RE. Retrieved August 23, 2025, from https://www.munichre.com/en/company/media-relations/media-information-and-corporate-news/media-information/2024/natural-disaster-figures-2023.html

Museru, M. L., Nazari, R., & Nikoo, M. R., & Karimi, M. (2025). Improving disaster resilience with causal machine learning for flood damage estimation. *Science of The Total Environment*, *995*, 180121. https://www.sciencedirect.com/science/article/abs/pii/S0048969725017619

Nearing, G., Cohen, D., Dube, V., Gauch, M., Gilon, O., Harrigan, S., Hassidim, A., Klotz, D., Kratzert, F., Metzger, A., Nevo, S., Pappenberger, F., Prudhomme, C., Shalev, G., Shenzis, S., Yednkachw Tekalign, T., Weitzner D., & Matias, Y. (2024). Global prediction of extreme floods in ungauged watersheds. *Nature*, *627*(8005), 559–563. https://doi.org/10.1038/s41586-024-07145-1

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511803161

Qi, X., Almeida, G., & Maldonado, S. (2024). Physics-informed neural networks for solving the shallow-water equations with terrain topography. *Journal of Fluid Mechanics*, *978*, A12. https://doi.org/10.1017/jfm.2023.1026

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, *378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Wang, Y., Shen, Y., & Salahshour, B. (2024). Urban flood extent segmentation and evaluation from real-world surveillance camera images using deep convolutional neural network. *Environmental Modelling & Software*, *173*, 105939. https://doi.org/10.1016/j.envsoft.2023.105939

Zhang, T., Zhang, R., Li, J., & Feng, P. (2025). Deep learning of flood forecasting by considering interpretability and physical constraints. *Hydrology and Earth System Sciences*, *29*(21), 5955–5974. https://hess.copernicus.org/articles/29/5955/2025

AUTHOR (S) BIOSKETCHES

AUTHOR (S) BIOSKETCHES

Mohammad Reyhani, Deputy Head of Life Insuran. ce, Life Insurance, Taavon Insurance Co., Tehran, Iran.
Email: M.reyhani1989@gmail.com
ORCID: 0009-0005-9503-1962

HOW TO CITE THIS ARTICLE

Reyhani, M. (2026). A physics-informed and causally-aware conceptual framework for flood risk pricing: A hybrid deep learning approach for data-scarce insurance markets. *Iranian Journal of Insurance Research, 15*(S1), 1-10.

DOI: 10.22056/ijir.2026.S1.01

URL: https://ijir.irc.ac.ir/article_160382.html