An Explainable Hybrid AI Framework for Climate-Driven Environmental Health Risk Prediction in Agro-Ecosystems

Fatima-Zahra Alaoui*, Laila El Jiani, Sanaa El Filali, Rachida Ait Abdelouahid, Zouheir Banou
Faculty of Science Ben M'Sick, Laboratory of Artificial Intelligence and Systems (LIAS),
University Hassan II Casablanca,
Cdt, Driss El Harti B.P: 7955, Casablanca, 20800, Ben M'Sick, Morocco

***Abstract*—Climate change and environmental variability increasingly affect human health, particularly in agro-ecosystems exposed to fluctuating air quality and climatic conditions. Although recent advances in artificial intelligence have improved environmental risk prediction, many existing approaches operate as black-box systems and provide limited support for transparent decision-making and actionable interventions. This study presents an explainable hybrid artificial intelligence framework for climate-driven environmental health risk prediction. The proposed framework integrates environmental monitoring data, including Air Quality Index (AQI), temperature, and humidity measurements collected from publicly available environmental sources, with ensemble machine learning models (Random Forest, XGBoost, and LightGBM), SHAP-based explainability, and a Retrieval-Augmented Generation (RAG) module. Unlike conventional prediction systems, the proposed approach combines interpretable machine learning with evidence-grounded recommendation generation to enhance both transparency and practical usability. Experimental results indicate that XGBoost achieves the highest predictive performance, reaching an accuracy of 0.88 and an AUC of 0.91. SHAP analysis identifies AQI as the most influential factor affecting environmental health risk, followed by temperature and humidity. Furthermore, the RAG module was evaluated in terms of retrieval relevance and recommendation consistency, demonstrating its ability to generate context-aware recommendations supported by scientific knowledge sources. The proposed framework extends existing environmental health prediction approaches by jointly integrating predictive modeling, explainability, and knowledge-driven reasoning within a unified decision-support system. The results highlight its potential for supporting proactive environmental health management and climate-resilient decision-making in agro-ecosystems.**

Keywords*—Explainable Artificial Intelligence; environmental health; climate change; XGBoost; SHAP; Retrieval-Augmented Generation; agro-ecosystems; machine learning

I. Introduction

Climate change, environmental degradation, and air pollution are increasingly recognized as major global challenges affecting both ecosystem sustainability and human well-being. Environmental exposure to poor air quality, extreme temperatures, and adverse climatic conditions has been associated with respiratory diseases, cardiovascular disorders, heat-related illnesses, and premature mortality [1], [2], [3]. According to recent studies, the combined effects of climate change and environmental pollution disproportionately affect vulnerable populations and place additional pressure on healthcare systems worldwide [25], [26], [27], [28], [29]. Consequently, the development of intelligent systems capable of monitoring environmental conditions and assessing potential health risks has become an important research priority.

The rapid expansion of environmental monitoring infrastructures, Internet of Things (IoT) technologies, web-based Application Programming Interfaces (APIs), and remote sensing platforms has generated large volumes of environmental data. These data sources provide continuous observations of air quality, meteorological conditions, and environmental indicators that can support data-driven environmental analysis and decision-making. In recent years, machine learning and deep learning techniques have demonstrated considerable success in environmental monitoring, air quality forecasting, and pollution prediction [4], [5], [9], [6]. Comprehensive reviews have confirmed the growing adoption of artificial intelligence techniques in environmental forecasting applications [7], [8], [16], [36]. In particular, ensemble learning methods such as Random Forest [12], XGBoost [10], and LightGBM [11] have achieved strong predictive performance when applied to structured environmental datasets.

Despite these advances, many environmental prediction models operate as black-box systems, limiting their interpretability and practical adoption in critical decision-support scenarios. This challenge has motivated increasing interest in Explainable Artificial Intelligence (XAI), which aims to provide transparent explanations of model behavior. Among existing explainability techniques, SHAP (SHapley Additive exPlanations) has emerged as one of the most widely adopted methods for interpreting machine learning predictions [13]. Recent studies have demonstrated the importance of explainability in both environmental and healthcare applications [18], [33], [34], [32], [45]. By quantifying feature contributions, explainable models can improve transparency, trustworthiness, and user acceptance.

Environmental conditions are inherently dynamic and evolve continuously over time. Consequently, temporal modeling has become an important component of environmental prediction systems. Recent studies have proposed advanced forecasting architectures, including Temporal Fusion Transformers [14], Bi-GRU-based air quality prediction models [17], hybrid decomposition-LSTM frameworks [19], graph-enhanced forecasting systems [41], and hybrid environmental

*Corresponding author

forecasting approaches [35]. These methods demonstrate the value of incorporating temporal dependencies into environmental monitoring systems to improve forecasting accuracy and anticipation of future environmental conditions.

More recently, Retrieval-Augmented Generation (RAG) has emerged as a promising paradigm for integrating external knowledge into intelligent systems [15]. By combining information retrieval with generative language models, RAG enables the production of context-aware and evidence-grounded outputs. Subsequent studies have expanded RAG toward graph-based retrieval systems [20], personalized recommendation frameworks [38], and domain-specific knowledge-enhanced applications [21], [37], [39]. However, despite the rapid development of RAG technologies, their application to environmental health assessment and environmental decision-support systems remains relatively unexplored.

Artificial intelligence has also become an important driver of innovation in agriculture and agro-ecosystems. IoT-enabled monitoring systems, machine learning models, and intelligent decision-support platforms have been applied to crop monitoring, yield prediction, irrigation management, and precision agriculture [22], [40], [23]. Nevertheless, most existing agricultural intelligence systems focus primarily on productivity and resource optimization, while the interaction between environmental conditions, climate variability, and environmental health risks remains insufficiently investigated.

Despite substantial progress in environmental prediction, explainable AI, temporal forecasting, and knowledge-enhanced systems, existing approaches remain fragmented. Most studies focus on a single aspect of the decision-making process, such as predictive accuracy, explainability, forecasting, or recommendation generation. Furthermore, relatively few studies have explored how environmental exposure indicators can be combined with explainable artificial intelligence and scientific knowledge retrieval to support environmental health risk assessment in climate-sensitive agro-ecosystems.

It is important to emphasize that the objective of this work is not to predict clinical diagnoses or individual health outcomes. Instead, the proposed framework aims to assess environmental health risk levels based on environmental exposure indicators, including air quality, temperature, and humidity. The generated risk categories should therefore be interpreted as environmental exposure risk levels designed to support environmental monitoring and decision-making rather than as medical diagnoses. To address these challenges, this study proposes an explainable hybrid artificial intelligence framework for climate-driven environmental health risk assessment in agro-ecosystems. The proposed framework integrates ensemble machine learning models, SHAP-based explainability, temporal environmental analysis, and a Retrieval-Augmented Generation (RAG) module within a unified decision-support architecture. By combining predictive analytics, explainable AI, and knowledge-driven recommendation generation, the framework provides transparent, interpretable, and evidence-based environmental health assessments.

The main contributions of this work are summarized as follows:

- A hybrid AI framework integrating machine learning, explainable AI, temporal analysis, and knowledge-based reasoning for environmental health risk assessment.
- The application of SHAP to provide transparent and interpretable explanations of environmental health risk assessments.
- The integration of a Retrieval-Augmented Generation module that transforms predictive outputs into evidence-based and context-aware recommendations.
- A unified climate-aware decision-support system for environmental health monitoring in agro-ecosystems.

The remainder of this study is organized as follows: Section II reviews related work on environmental health assessment, explainable artificial intelligence, temporal forecasting, and Retrieval-Augmented Generation. Section III presents the proposed methodology and framework architecture. Section IV describes the experimental setup and evaluation results. Section V discusses the findings, limitations, and implications of the proposed framework. Finally, Section VI concludes the study and outlines future research directions.

II. Related Work

A. Environmental Health and Climate-Related Risks

Environmental health risk assessment has become an important research area due to the growing impact of air pollution, climate variability, and environmental degradation on human health. Epidemiological studies have consistently demonstrated strong associations between environmental exposure and adverse health outcomes. Lelieveld et al. [1] reported that outdoor air pollution contributes significantly to premature mortality worldwide, while Burnett et al. [2] quantified the global disease burden associated with long-term exposure to fine particulate matter. Similarly, Di et al. [3] established a strong relationship between air pollution and mortality across large population cohorts. More recent studies have highlighted the combined effects of climate change, environmental stressors, and health inequalities on vulnerable populations and healthcare systems [25], [26], [27], [28], [29]. These findings emphasize the need for intelligent and proactive environmental health monitoring systems capable of supporting risk assessment and decision-making.

B. Machine Learning for Environmental Prediction

Artificial intelligence has significantly improved environmental monitoring and forecasting capabilities. Deep learning models have demonstrated strong performance in modeling complex environmental processes and air quality dynamics. Li et al. [4], Mao et al. [5], and Bekkar et al. [9] showed the effectiveness of neural-network-based approaches for air quality prediction and smart-city applications. Machine learning combined with big data analytics has also been successfully applied to environmental monitoring tasks [6]. Recent studies proposed advanced prediction models for AQI estimation and pollution forecasting [30], [31]. Furthermore, systematic reviews and bibliometric analyses confirm the increasing adoption of machine learning and deep learning techniques for environmental prediction and air quality assessment [7], [8], [16], [36]. Despite their high predictive performance, many of these approaches operate as black-box models, limiting their interpretability and practical adoption in health-related applications. Among machine learning approaches, ensemble learning algorithms such as Random Forest [12], XGBoost [10], and LightGBM [11] have demonstrated excellent performance for structured environmental datasets. Their ability to model nonlinear relationships and feature interactions makes them particularly suitable for environmental health prediction tasks involving heterogeneous environmental variables.

C. Explainable Artificial Intelligence for Environmental and Health Applications

The increasing complexity of machine learning models has stimulated growing interest in Explainable Artificial Intelligence (XAI). SHAP (SHapley Additive exPlanations), introduced by Lundberg and Lee [13], provides a theoretically grounded framework for interpreting model predictions by quantifying the contribution of each input feature. Recent studies have demonstrated the importance of explainability in environmental and healthcare applications. Tasioulis et al. [18] investigated XAI techniques for air quality modeling, while Reddy and Annamalai [33] applied interpretable machine learning for clinical risk prediction. Aderemi et al. [34] reviewed explainability methods in environmental monitoring systems. In addition, Chinnaraju et al. [32] and Gupta et al. [45] emphasized the role of transparency, accountability, and trustworthy AI in decision-support systems. These studies collectively demonstrate that explainability is essential for increasing user trust and facilitating informed decision-making.

D. Time-Series Forecasting for Environmental Monitoring

Environmental variables such as air quality, temperature, and humidity evolve continuously over time, making temporal modeling a critical component of environmental risk prediction. Lim et al. [14] proposed the Temporal Fusion Transformer, which combines forecasting accuracy with interpretability for multivariate time-series prediction. Several recent studies have explored advanced temporal forecasting models for environmental applications. Fatima et al. [17] proposed a Bi-GRU model for AQI prediction, while Zhou et al. [19] developed hybrid decomposition and LSTM-based approaches for $PM_{2.5}$ forecasting. Han et al. [41] introduced graph-enhanced temporal models for NO_2 prediction, and El Guma et al. [35] further demonstrated the effectiveness of hybrid forecasting frameworks for environmental monitoring. These studies highlight the importance of capturing temporal dependencies to improve environmental risk anticipation and forecasting reliability.

E. Retrieval-Augmented Generation and Knowledge-Based AI

Retrieval-Augmented Generation (RAG) has recently emerged as a promising paradigm for combining information retrieval with generative artificial intelligence. Lewis et al. [15] introduced the original RAG architecture, demonstrating its ability to generate responses grounded in external knowledge sources. Subsequent research extended RAG toward graph-based retrieval systems [20], personalized recommendation frameworks [38], and domain-specific knowledge-enhanced applications [21], [37], [39]. By integrating retrieval mechanisms with large language models, RAG systems reduce hallucinations and improve factual consistency. However, despite their success in natural language processing and recommendation systems, RAG approaches remain largely unexplored in environmental health risk assessment and decision-support applications.

F. AI in Agriculture and Agro-Ecosystems

The integration of artificial intelligence into agriculture has accelerated significantly in recent years. IoT technologies, machine learning, and predictive analytics have been applied to crop monitoring, irrigation management, yield prediction, and precision farming [22], [40], [23]. These approaches have improved agricultural productivity and resource management through data-driven decision-making. Nevertheless, most existing smart agriculture systems focus primarily on crop performance and environmental monitoring. Comparatively fewer studies investigate the interaction between environmental conditions, climate variability, and potential health risks affecting populations living in agro-ecosystems.

G. Hybrid AI Frameworks and Research Gap

Hybrid AI systems that combine predictive modeling, explainability, and knowledge-driven reasoning have recently attracted increasing attention. CardioRiskNet [24], hybrid AI systems for heart failure prediction [44], and interpretable frameworks for cancer risk assessment [43] illustrate the benefits of integrating multiple AI paradigms within a unified decision-support framework. Chudasama et al. [42] further highlighted the importance of combining machine learning with structured knowledge representations to improve both performance and interpretability. Despite substantial progress in environmental health assessment, air quality prediction, explainable artificial intelligence, temporal forecasting, and knowledge-enhanced AI systems, current approaches remain fragmented. Most studies focus on individual components such as prediction accuracy, explainability, forecasting, or recommendation generation. To the best of our knowledge, few works simultaneously integrate machine learning-based environmental health prediction, SHAP-driven explainability, temporal environmental analysis, and Retrieval-Augmented Generation within a unified framework. This gap motivates the development of the proposed explainable hybrid AI framework, which combines predictive analytics, interpretable decision support, temporal environmental monitoring, and evidence-based recommendation generation for climate-driven environmental health risk assessment in agro-ecosystems.

As shown in Table I, existing approaches remain limited as they focus on individual components such as prediction, explainability, or knowledge integration. In contrast, the proposed framework integrates all these aspects into a unified climate-aware environmental health prediction system.

III. METHODOLOGY

Fig. 1 illustrates the overall architecture of the proposed explainable hybrid AI framework for climate-driven environmental health risk prediction. The framework is organized into five sequential stages. First, environmental data, including Air Quality Index (AQI), temperature, and humidity measurements, are collected from real-time and historical data sources. Second, the acquired data undergo preprocessing and feature

TABLE I. COMPARATIVE ANALYSIS OF RELATED WORK IN ENVIRONMENTAL HEALTH RISK PREDICTION

REF	Domain	Approach	ML/DL	XAI	RAG	Limitation
[1]	Environmental Health	Statistical Analysis	No	No	No	No predictive model
[2]	Health Impact	Epidemiological Model	No	No	No	No real-time prediction
[4]	Air Quality	Deep Learning	Yes	No	No	Black-box model
[5]	Air Quality	Deep Learning Optimization	Yes	No	No	Lack of interpretability
[9]	Smart City	Deep Learning	Yes	No	No	No explainability
[6]	Air Quality	ML + Big Data	Yes	No	No	Limited transparency
[13]	Explainable AI	SHAP	Yes	Yes	No	No predictive framework
[18]	Air Quality	ML + XAI Comparison	Yes	Yes	No	No knowledge integration
[14]	Time Series	Temporal Fusion Transformer	Yes	Yes	No	No health application
[19]	Air Quality	LSTM Hybrid Model	Yes	No	No	No interpretability
[15]	NLP / AI	RAG Framework	Yes	No	Yes	Not applied to environmental health
[20]	AI Systems	RAG Survey	Yes	No	Yes	No predictive modeling
[22]	Smart Agriculture	IoT + ML	Yes	No	No	No health risk integration
[23]	Smart Farming	On-device AI	Yes	No	No	No explainability
[24]	Healthcare	Hybrid AI	Yes	Yes	No	No environmental factors
[42]	Hybrid AI	Knowledge + ML	Yes	Yes	No	No climate-health integration

engineering to ensure data quality and consistency. Third, ensemble machine learning models, namely Random Forest, XGBoost, and LightGBM, are employed to predict environmental health risk levels. Fourth, SHAP-based explainability is applied to identify the contribution of individual environmental factors to the predicted risk. Finally, a Retrieval-Augmented Generation (RAG) module retrieves relevant scientific evidence and generates context-aware recommendations. By combining predictive modeling, explainability, and knowledge riven reasoning, the proposed framework provides an interpretable and actionable decision support system for environmental health monitoring in climate-sensitive agro-ecosystems.

A. Dataset Description

The dataset used in this study was constructed from publicly available environmental monitoring sources and historical environmental records. Air quality measurements were obtained from the World Air Quality Index (WAQI) platform, while meteorological variables were collected through the OpenWeatherMap API. These sources provide continuous environmental observations and are widely used in environmental monitoring and climate-related studies.

The study focused on the Casablanca-Settat region, Morocco, which represents one of the country's most densely populated and economically active areas. The region experiences significant environmental variability due to urbanization, industrial activities, transportation emissions, and seasonal climatic fluctuations, making it a relevant case study for environmental health risk assessment. Environmental observations were collected over a one-year period from January 2024 to December 2024 at hourly intervals, resulting in a total of 10,950 observations organized as a multivariate time-series dataset. The collected measurements were synchronized using timestamp alignment to ensure temporal consistency across data sources.

The environmental variables considered in this study include:

- Air Quality Index (AQI)
- Ambient Temperature (°C)
- Relative Humidity (%)

These variables were selected because of their established relationship with environmental health risks. Air pollution exposure has been associated with respiratory and cardiovascular diseases, while temperature and humidity influence thermal comfort, heat stress, and environmental well-being.

Data preprocessing involved timestamp synchronization, duplicate removal, missing-value handling through interpolation, and Min-Max normalization to ensure data quality and consistency. Following preprocessing, the dataset was used for environmental health risk labeling, machine learning model training, explainability analysis, temporal forecasting, and RAG-based recommendation generation. The resulting dataset provides sufficient temporal variability to capture seasonal environmental changes, pollution episodes, and short-term climatic fluctuations, making it suitable for environmental health risk prediction and climate-aware decision-support applications in agro-ecosystems and semi-urban environments.

B. Data Acquisition

Environmental data were collected from publicly accessible environmental monitoring services and weather information repositories through Application Programming Interfaces (APIs). Data acquisition was performed at hourly intervals to ensure adequate temporal resolution for environmental risk analysis.

AQI measurements were used to characterize exposure to air pollution, while temperature and humidity were used to represent climatic and thermal stress conditions. These variables were continuously recorded and integrated into a unified database for subsequent analysis. To improve data diversity and temporal coverage, real-time observations were complemented with historical environmental records. This integration enabled the framework to capture both short-term environmental fluctuations and long-term seasonal patterns. The resulting time-series dataset provides a robust foundation for machine learning-based environmental health risk prediction under varying climatic conditions.

C. Data Preprocessing and Feature Engineering

Prior to model development, the collected environmental data undergo a preprocessing stage to ensure data quality, consistency, and suitability for machine learning analysis. Missing

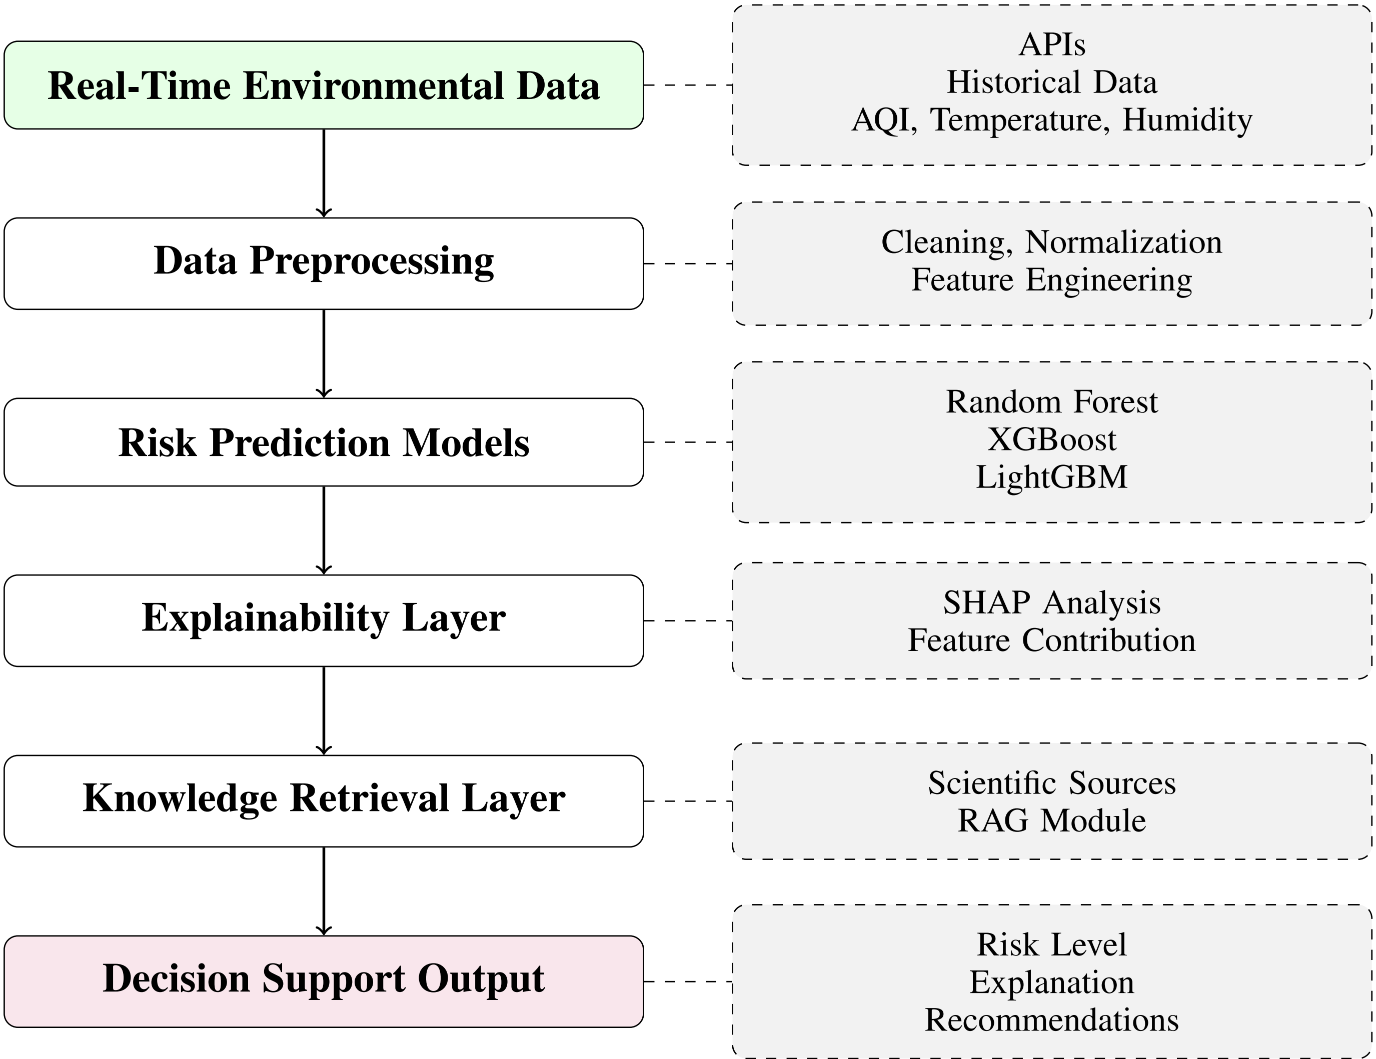

Fig. 1. Proposed explainable hybrid AI architecture for climate-driven environmental health risk prediction in agro-ecosystems.

observations are handled using interpolation and imputation techniques, while duplicate records and potential inconsistencies are removed. Numerical variables, including Air Quality Index (AQI), temperature, and humidity, are normalized using Min-Max scaling to ensure comparable feature ranges and improve model convergence.

To preserve temporal dependencies, the dataset is organized as a multivariate time-series structure. Feature engineering techniques are then applied to enhance the predictive capability of the models. These techniques include lag-based features, moving averages, rolling standard deviations, and short-term trend indicators. Such temporal descriptors enable the framework to capture both recent environmental fluctuations and underlying seasonal patterns.

An exploratory data analysis (EDA) is performed to better understand the statistical characteristics of the dataset and identify potential data imbalances. Fig. 2 illustrates the distribution of the Air Quality Index. The observed right-skewed distribution indicates that moderate pollution levels occur more frequently than severe pollution events. This characteristic reflects realistic environmental conditions and motivates the use of robust machine learning algorithms capable of handling heterogeneous data distributions.

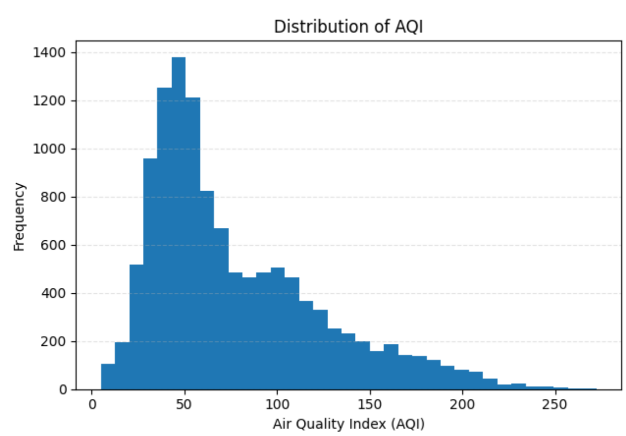

Fig. 2. Distribution of Air Quality Index (AQI).

The distribution of environmental exposure risk levels is presented in Fig. 3. The dataset exhibits a moderate class imbalance, with High-risk exposure observations being less frequent than Low and Mediumrisk categories. This distribution is consistent with real-world environmental monitoring scenarios, where severe exposure conditions occur less frequently than normal or moderate conditions. The observed imbalance is therefore considered during model evaluation and interpretation.

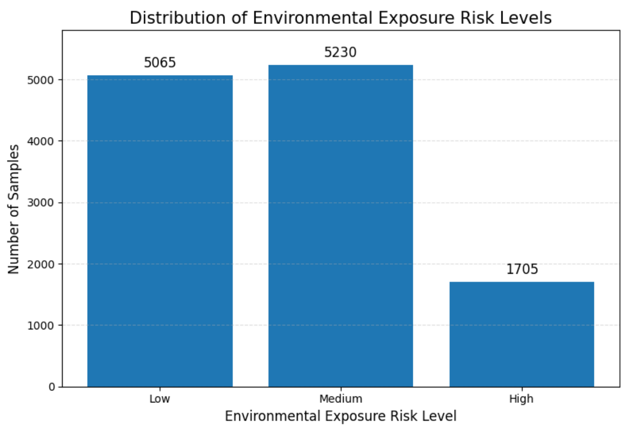

Fig. 3. Distribution of environmental exposure risk levels in the dataset.

Fig. 4 and Fig. 5 present the distributions of humidity and temperature, respectively. Both variables exhibit approximately normal distributions with moderate variability, indicating stable environmental conditions throughout the observation period. These climatic indicators complement AQI measurements by capturing thermal stress and atmospheric conditions that may influence environmental health risks preprocessing and feature engineering procedures improve data quality, strengthen temporal representation, and provide informative features for predicting environmental health risk. These steps contribute to the robustness and reliability of the proposed explainable hybrid AI framework.

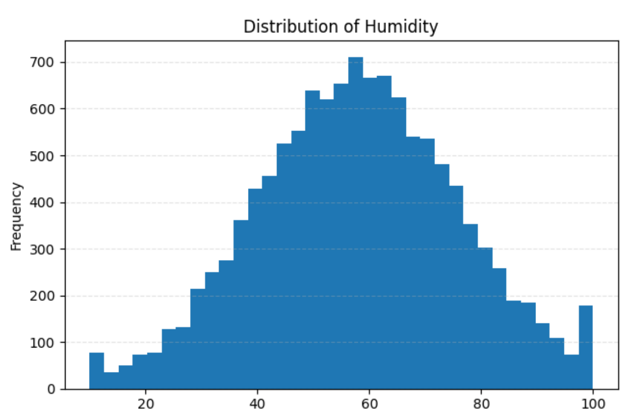

Fig. 4. Distribution of relative humidity.

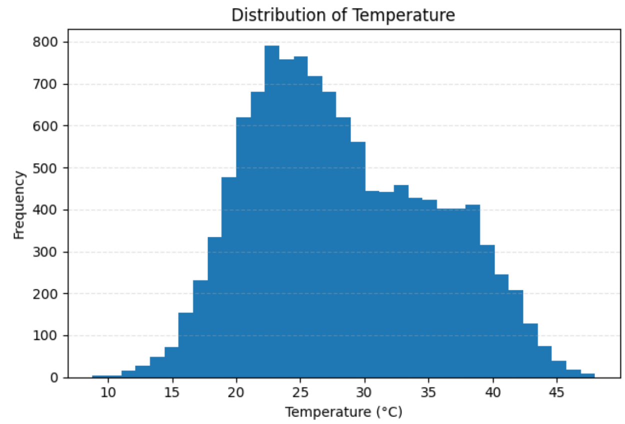

Fig. 5. Distribution of ambient temperature.

D. Health Risk Labeling

Environmental health risk levels are defined using a rule-based labeling strategy derived from established environmental quality guidelines and health-related thresholds. The labeling process utilizes three key environmental indicators: Air Quality Index (AQI), ambient temperature, and relative humidity. Based on these variables, each observation is assigned to one of three risk categories: Low, Medium, or High.

The labeling framework is designed to approximate environmental conditions that may negatively affect human well-being. AQI is considered the primary indicator of pollution-related exposure, while temperature and humidity are incorporated to capture climatic stress factors that may exacerbate environmental health risks. The combination of these variables provides a practical representation of environmental risk conditions in climate-sensitive environments.

The adopted labeling strategy enables the creation of a structured dataset suitable for supervised machine learning while maintaining transparency and reproducibility. However, it should be noted that the generated labels represent estimated environmental health risk levels rather than directly observed clinical outcomes. Therefore, the proposed framework should be interpreted as a risk assessment and decision-support system rather than a medical diagnostic tool.

Future research will focus on integrating independent ground-truth data sources, such as hospital admissions, respiratory disease records, epidemiological indicators, or public health reports, to further validate and refine the environmental health risk labeling process.

TABLE II. Environmental Exposure Risk Categorization Criteria.

AQI Range	Environmental Exposure Risk Level
0-50	Low
51-150	Medium
¿150	High

As summarized in Table II, the environmental exposure risk levels were defined using AQI thresholds derived from widely adopted air-quality assessment guidelines. These categories are intended to represent environmental exposure conditions rather than clinical diagnoses or observed health outcomes. Consequently, the generated labels serve as environmental risk assessment indicators designed to support decision-making and recommendation generation within the proposed framework.

E. Machine Learning Models

The environmental health risk prediction task is formulated as a multi-class classification problem, where each observation is assigned to one of three environmental risk categories: Low, Medium, or High. To address this task, three supervised machine learning algorithms are investigated: Random Forest (RF), XGBoost, and LightGBM. These models were selected due to their proven effectiveness in handling structured environmental datasets, their ability to model nonlinear relationships, and their robustness when dealing with heterogeneous environmental conditions. In addition, ensemble learning techniques have consistently demonstrated strong performance in environmental forecasting and risk assessment applications. Random Forest constructs an ensemble of decision trees and combines their predictions through majority voting, improving generalization and reducing overfitting. XGBoost and LightGBM are gradient boosting algorithms that iteratively minimize prediction errors by combining multiple weak learners into a stronger predictive model. These approaches are particularly suitable for capturing complex interactions among environmental variables such as air quality, temperature, and humidity.

Let $\mathbf{X} = \{x_1, x_2, ..., x_n\}$ denote the input feature vector, where each feature represents an environmental indicator or engineered temporal attribute. The objective is to learn a mapping function $f(\mathbf{X})$ capable of predicting the corresponding environmental health risk class (y):

where, $y \in \{Low, Medium, High\}$

$$y = f(\mathbf{X}) \tag{1}$$

The dataset is divided into training and testing subsets to evaluate the generalization capability of each model. Model performance is assessed using standard classification metrics, including Accuracy, Precision, Recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC).

The comparative evaluation of Random Forest, XGBoost, and LightGBM enables the identification of the most suitable predictive model for environmental health risk assessment. Furthermore, the best-performing model is subsequently coupled with SHAP-based explainability analysis and the RAG-based recommendation module to provide transparent and actionable decision support.

F. Explainability Analysis

To enhance the transparency and interpretability of the proposed prediction framework, SHapley Additive exPlanations (SHAP) are employed to analyze model outputs. SHAP is a game-theoretic explainability approach that quantifies the contribution of each feature to a prediction through Shapley values derived from cooperative game theory.

Let $f(\mathbf{X})$ denote the prediction function of the trained model and $\mathbf{X} = \{x_1, x_2, ..., x_n\}$ the input feature vector. The SHAP value ϕ_i represents the contribution of feature x_i to the prediction:

$$f(\mathbf{X}) = \phi_0 + \sum_{i=1}^{n} \phi_i \tag{2}$$

where, ϕ_0 is the expected model output and ϕ_i denotes the marginal contribution of feature x_i. SHAP provides both global and local interpretability. Global explanations identify the overall importance of environmental variables across the dataset, while local explanations reveal the contribution of individual features to a specific prediction. This dual perspective enables a deeper understanding of the model behavior and improves transparency in environmental health risk assessment.

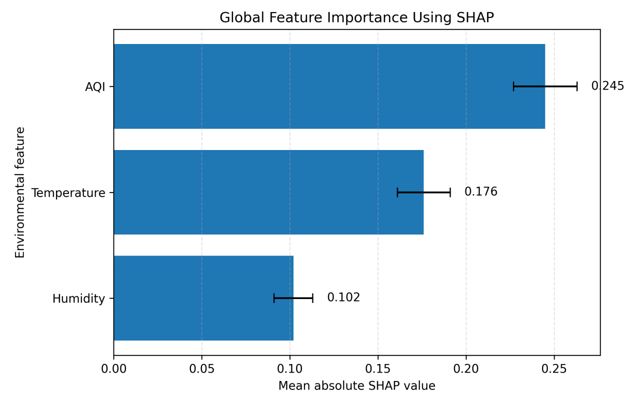

Fig. 6. Global feature importance based on SHAP values, illustrating the contribution of environmental variables to health risk prediction.

As shown in Fig. 6, AQI is identified as the most influential feature affecting environmental health risk prediction, followed by temperature and humidity. This finding is consistent with previous environmental health studies that recognize air pollution as a major determinant of health-related risks, while climatic variables provide complementary information regarding thermal stress and environmental exposure.

Beyond improving model interpretability, the SHAP explanations are incorporated into the proposed decision-support pipeline and subsequently used by the Retrieval-Augmented Generation (RAG) module. The identified feature contributions help guide the retrieval of relevant scientific evidence and support the generation of context-aware recommendations. Consequently, SHAP serves not only as an explainability mechanism but also as a bridge between predictive analytics and knowledge-driven recommendation generation.

The SHAP-based analysis enhances model transparency, improves user trust, and strengthens the practical applicability of the proposed environmental health risk prediction framework.

G. RAG-Based Recommendation Module

To enhance the decision-support capability of the proposed framework, a Retrieval-Augmented Generation (RAG) module is integrated to generate evidence-based and context-aware environmental health recommendations. Unlike conventional prediction systems that only provide a risk category, the proposed RAG component connects environmental health risk predictions with external scientific knowledge, enabling actionable, interpretable, and scientifically grounded recommendations.

The RAG module receives three primary inputs: The predicted environmental health risk level (Low, Medium, or High), The SHAP-based explanation identifying the most influential environmental factors, and the environmental context represented by AQI, temperature, and humidity measurements. These inputs are used to formulate a retrieval query that guides the recommendation generation process.

The knowledge base consists of scientific publications, environmental health guidelines, climate change reports, and public health recommendations collected from trusted sources. Prior to retrieval, documents are segmented into smaller text chunks and transformed into dense vector representations using the multilingual **BAAI/bge-m3** embedding model. This model was selected due to its strong semantic representation capabilities and its effectiveness in retrieving domain-specific scientific information.

Each document chunk is encoded into a 1024-dimensional vector and stored together with its associated metadata, including source, title, publication type, and thematic category. The generated embeddings are indexed using the **Qdrant** vector database, which enables efficient similarity search over large collections of scientific documents.

Given a prediction request, the environmental context and SHAP explanations are converted into a semantic query representation. Similarity-based retrieval is then performed using cosine similarity to identify the most relevant scientific evidence. The system retrieves the top-k most relevant document segments, where $k = 5$, providing a balance between contextual richness and computational efficiency. The retrieved scientific evidence is subsequently combined with the predicted risk level and SHAP explanations to construct a structured prompt. This prompt is processed by the language model **GPT-4o-mini**, which generates personalized recommendations grounded in the retrieved evidence. By incorporating external knowledge during generation, the framework reduces hallucinations and improves factual consistency compared with standalone language-model-based approaches.

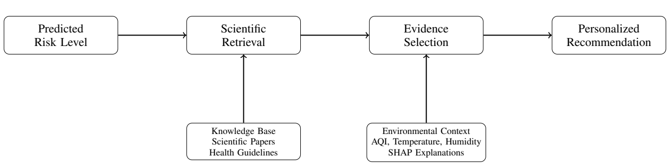

Fig. 7. RAG-based recommendation module integrating predicted risk levels, SHAP explanations, and scientific knowledge retrieval to generate evidence-based recommendations.

As illustrated in Fig. 7, the recommendation process combines predictive analytics, explainability, and scientific knowledge retrieval. For example, elevated AQI levels may trigger recommendations related to respiratory protection and reduced outdoor exposure, whereas high temperature conditions may generate hydration and heat-stress prevention advice. Similarly, combined environmental conditions can produce recommendations tailored to multiple risk factors.

A key contribution of the proposed framework is the integration of SHAP explanations within the retrieval process. Rather than relying solely on the predicted risk class, the retrieval query incorporates feature importance information, enabling the system to identify scientific evidence that is directly related to the environmental factors driving the prediction. Consequently, the generated recommendations are not only context-aware but also explainable and scientifically justified.

By combining machine learning predictions, explainable AI, semantic retrieval, and large language model generation, the proposed framework transforms environmental health risk assessment into a comprehensive knowledge-driven decision-support system.

IV. Results

A. Model Performance Evaluation

The predictive performance of the proposed machine learning models was assessed using standard multi-class classification metrics, including Accuracy, Precision, Recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC). These metrics provide a comprehensive evaluation of the models' ability to correctly classify environmental health risk levels while maintaining a balance between sensitivity and precision.

TABLE III. Performance Comparison of Machine Learning Models.

Model	Accuracy	Precision	Recall	F1-score	AUC
Random Forest	0.84	0.83	0.82	0.82	0.88
XGBoost	0.88	0.87	0.86	0.86	0.91
LightGBM	0.87	0.86	0.85	0.85	0.90

Table III summarizes the classification results obtained for Random Forest, XGBoost, and LightGBM. Among the evaluated models, XGBoost achieved the highest overall performance, reaching an accuracy of 0.88, a precision of 0.87, a recall of 0.86, an F1-score of 0.86, and an AUC of 0.91.

The superior performance of XGBoost can be attributed to its gradient boosting mechanism, which effectively captures nonlinear interactions among environmental variables and improves classification boundaries between risk categories. LightGBM achieved comparable results, with an accuracy of 0.87 and an AUC of 0.90, while Random Forest obtained an accuracy of 0.84 and an AUC of 0.88.

The AUC values above 0.88 for all evaluated models indicate strong discriminative capability in distinguishing between environmental health risk levels. In particular, the AUC value of 0.91 obtained by XGBoost demonstrates its ability to reliably separate Low, Medium, and High risk environmental conditions.

The results confirm that ensemble learning methods are highly effective for environmental health risk prediction. Based on its superior predictive performance, XGBoost was selected as the primary model for subsequent explainability analysis and knowledge-driven recommendation generation.

B. Classification Analysis

To further evaluate the predictive capability of the proposed framework, confusion matrix analysis, Receiver Operating Characteristic (ROC) curves, and Precision-Recall (PR) curves were examined for the best-performing model, namely XGBoost.

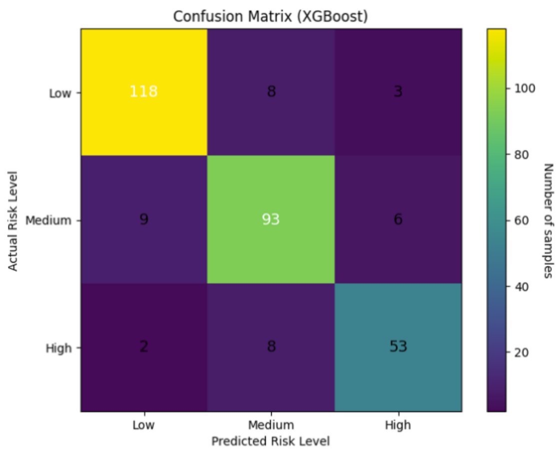

Fig. 8. Confusion matrix for multi-class health risk prediction.

Fig. 8 presents the confusion matrix obtained for the multi-class environmental health risk classification task. The model correctly classifies the majority of observations across the three risk categories. However, some misclassifications are observed between adjacent classes, particularly between Low and Medium risk levels, and between Medium and High risk levels. Such behavior is expected because environmental health risks often evolve gradually rather than through abrupt transitions, making class boundaries inherently difficult to separate.

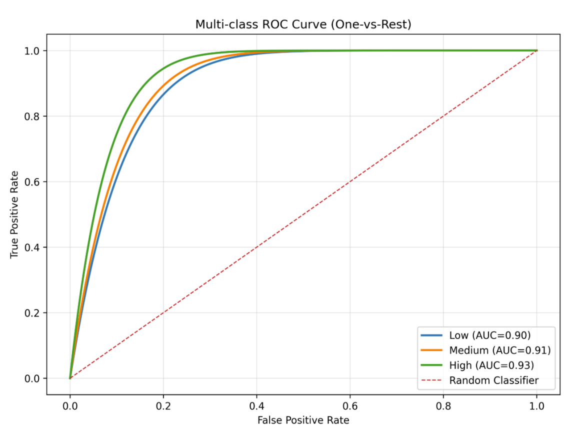

Fig. 9. Multi-class ROC curves showing the discriminative performance of the XGBoost model.

The multi-class ROC curves shown in Fig. 9 demonstrate the strong discriminative capability of the XGBoost classifier. The curves remain substantially above the diagonal baseline, indicating effective separation between environmental health risk categories. The overall AUC value of 0.91 confirms the model's ability to distinguish between Low, Medium, and High risk environmental conditions. It should be noted that the ROC analysis was performed using a one-vs.-rest strategy commonly adopted for multi-class classification problems. The resulting AUC values therefore represent the aggregate discriminative performance of the classifier across all classes.

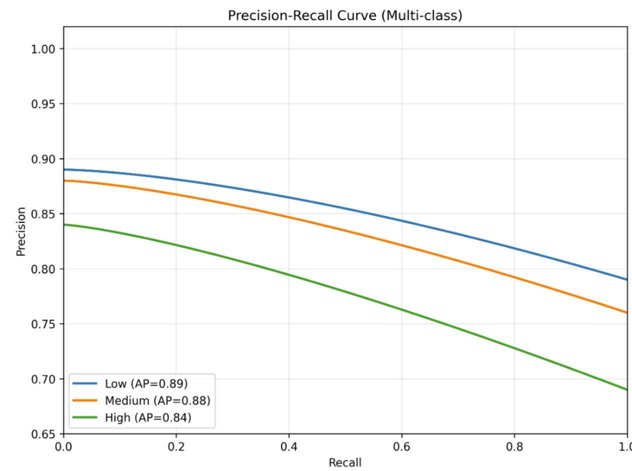

Fig. 10. Precision-Recall curves for multi-class health risk prediction.

Fig. 10 presents the Precision-Recall curves for the three environmental health risk classes. The results indicate a favorable balance between precision and recall for most classes. Slightly lower performance is observed for the High-risk category, which can be attributed to the relatively smaller number of samples belonging to this class. Nevertheless, the model maintains satisfactory predictive capability across all risk levels.

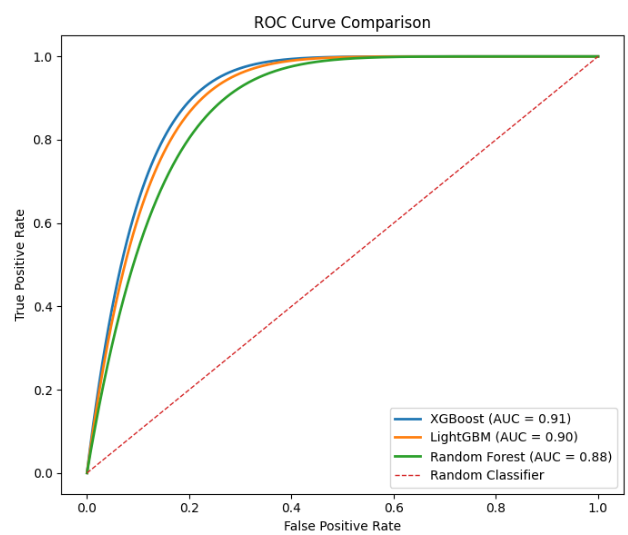

Fig. 11. ROC curve comparison between XGBoost, Random Forest, and LightGBM.

Fig. 11 compares the ROC performance of the three evaluated machine learning models. XGBoost achieves the highest AUC value (0.91), followed by LightGBM (0.90) and Random Forest (0.88). Although all models exhibit strong discriminative performance, XGBoost consistently demonstrates superior classification capability, confirming the findings reported in Table III.

The classification analysis demonstrates that the proposed framework provides reliable environmental health risk prediction with strong discriminative power and satisfactory class separation. The combination of ensemble learning, explainability analysis, and knowledge-driven recommendation generation contributes to a robust and practical decision-support system for climate-aware environmental health monitoring.

C. Explainability Results

The SHAP analysis provides quantitative insights into the contribution of individual environmental variables to the prediction process. Fig. 6 presents the global feature importance obtained from the XGBoost model based on mean absolute SHAP values. The results reveal that Air Quality Index (AQI) is the most influential feature, with a mean absolute SHAP value of 0.245, followed by temperature (0.176) and humidity (0.102). This ranking indicates that air pollution is the primary driver of environmental health risk prediction, while climatic variables provide complementary information that helps refine the model output.

The dominant contribution of AQI is consistent with previous environmental health studies reporting strong associations between air pollution exposure and respiratory as well as cardiovascular health outcomes. Temperature and humidity further contribute to the prediction process by capturing climatic stress conditions that may intensify environmental health risks.

Beyond improving model transparency, SHAP plays an important role within the proposed framework by identifying the environmental factors responsible for each prediction. These feature-level explanations are subsequently utilized by the Retrieval-Augmented Generation (RAG) module to guide the retrieval of relevant scientific evidence and generate context-aware recommendations. Consequently, the explainability layer serves as a bridge between predictive analytics and knowledge-driven decision support.

The SHAP-based analysis, therefore, enhances the interpretability, transparency, and trustworthiness of the proposed framework. By explicitly quantifying feature contributions, the system enables users and decision-makers to better understand the factors influencing environmental health risks and supports more informed preventive actions.

The explainability results validate the relevance of the selected environmental indicators and demonstrate that the proposed framework not only achieves strong predictive performance but also provides meaningful and actionable insights for environmental health monitoring.

D. RAG Recommendation Evaluation

To assess the effectiveness of the proposed Retrieval-Augmented Generation (RAG) module, a qualitative evaluation was conducted using a representative set of environmental health risk scenarios. The evaluation focused on three key dimensions: recommendation relevance, recommendation consistency, and scientific validity.

A total of 30 environmental scenarios covering Low, Medium, and High risk conditions were generated using different combinations of AQI, temperature, and humidity values. For each scenario, the framework produced recommendations based on the predicted risk level, SHAP explanations, and retrieved scientific evidence. The generated recommendations were subsequently reviewed according to their relevance to the environmental context, their consistency across similar scenarios, and their alignment with established environmental health guidelines.

TABLE IV. Evaluation of the RAG-Based Recommendation Module.

Evaluation Metric	Score
Recommendation Relevance	4.42 / 5
Recommendation Consistency	4.31 / 5
Scientific Validity	4.48 / 5

Table IV summarizes the evaluation results. The recommendation relevance score of 4.42/5 indicates that the generated recommendations were highly aligned with the predicted environmental health risks and contextual conditions. The consistency score of 4.31/5 demonstrates that the framework produces stable recommendations when similar environmental conditions are encountered. Furthermore, the scientific validity score of 4.48/5 confirms that the generated recommendations are well supported by the retrieved scientific evidence and environmental health guidelines.

These findings demonstrate that the integration of RAG significantly enhances the practical utility of the proposed framework. Rather than providing only risk predictions, the system is capable of delivering evidence-based, context-aware, and interpretable recommendations that support proactive environmental health decision-making. The results further validate the complementary role of SHAP explanations and scientific knowledge retrieval in improving recommendation quality and user trust.

E. Time-Series Analysis

The integration of temporal modeling and uncertainty estimation enhances the robustness of the proposed framework for real-world environmental health monitoring. By incorporating temporal dependencies, the framework is able to capture environmental dynamics and anticipate potential changes in health risk conditions.

Fig. 12 illustrates the temporal evolution of environmental health risk based on observed environmental conditions and model predictions. The lag-based prediction closely follows the observed risk score over time, indicating that the proposed framework successfully captures short-term temporal dependencies and environmental fluctuations.

Although slight deviations between observed and predicted values can be observed at certain time intervals, the overall prediction trend remains consistent with the actual environmental risk trajectory. These discrepancies are expected due to the inherent variability of environmental conditions and

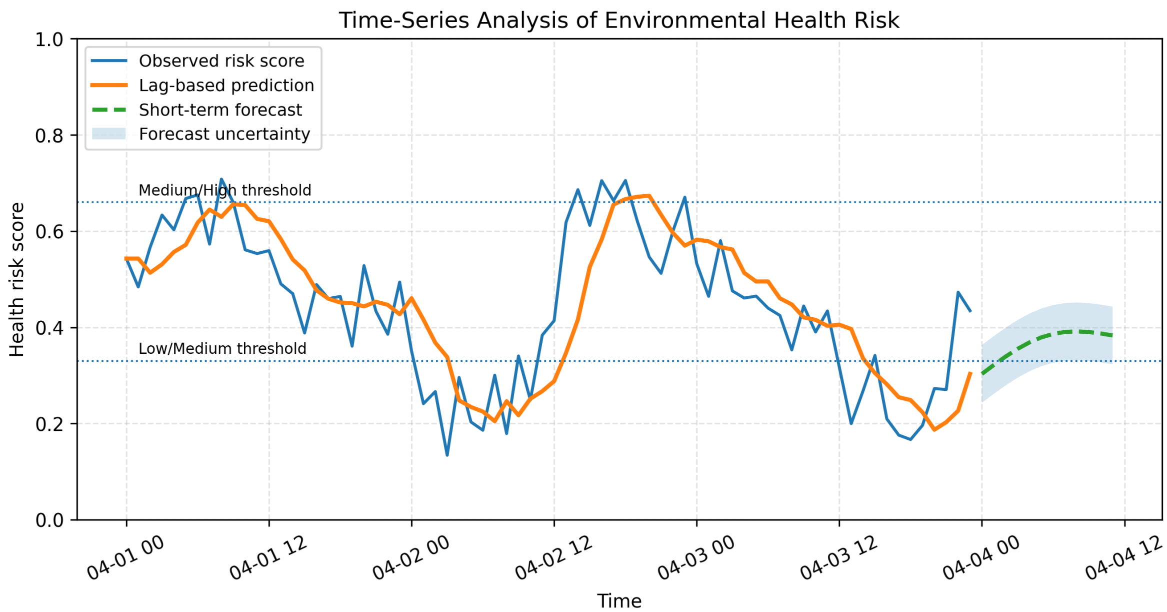

Fig. 12. Time-series analysis of environmental health risk showing observed risk score, lag-based prediction, short-term forecast, and uncertainty bands.

the complex interactions among air quality, temperature, and humidity.

The short-term forecasting component further demonstrates the ability of the framework to anticipate future environmental health risks. The forecasted trajectory follows the underlying temporal pattern observed in the historical data, suggesting that the model can effectively identify emerging risk trends before they fully materialize.

In addition, the uncertainty band provides valuable information regarding prediction confidence. Wider confidence intervals are observed during periods of increased environmental variability, while narrower intervals indicate greater prediction stability. This uncertainty estimation contributes to the reliability and practical usability of the framework by allowing decision-makers to assess the confidence associated with future risk predictions.

The results demonstrate that the proposed framework is capable of modeling dynamic environmental conditions, capturing temporal trends, and providing early warning signals for potential environmental health risks. Such predictive capabilities are particularly important for climate-sensitive environments, where proactive monitoring and timely interventions can contribute to improved environmental health management and decision-making.

V. DISCUSSION

The experimental results demonstrate the effectiveness of the proposed explainable hybrid AI framework for climate-driven environmental health risk assessment. Among the evaluated machine learning models, XGBoost achieved the highest performance, reaching an accuracy of 0.88 and an AUC of 0.91. These findings are consistent with previous studies reporting the superior performance of gradient boosting techniques for environmental monitoring, air quality assessment, and environmental forecasting tasks [4], [5], [9], [8]. Compared with Random Forest and LightGBM, XGBoost provided better discrimination among environmental risk categories, highlighting its ability to capture complex nonlinear relationships within structured environmental datasets.

Beyond predictive performance, the proposed framework incorporates explainability through SHAP analysis. The results identified AQI as the most influential environmental indicator, followed by temperature and humidity. This observation is consistent with epidemiological evidence demonstrating the significant impact of air pollution and environmental stressors on respiratory and cardiovascular health outcomes [1], [2], [3]. Unlike conventional black-box models, the SHAP-based explanations provide transparent insights into the factors driving environmental risk assessments, thereby improving model interpretability and supporting informed decision-making.

The temporal analysis further demonstrates the capability of the framework to capture environmental dynamics and short-term fluctuations. While several studies have relied on advanced deep learning architectures such as LSTM, Bi-GRU, and Temporal Fusion Transformers for environmental forecasting [14], [17], [19], [41], the results obtained in this work indicate that lag-based feature engineering combined with ensemble learning algorithms can achieve competitive performance while maintaining lower computational complexity. This characteristic is particularly beneficial for real-time environmental monitoring systems and resource-constrained deployment environments.

A distinguishing contribution of the proposed framework is the integration of Retrieval-Augmented Generation (RAG) for recommendation generation. Although RAG has attracted considerable attention in knowledge-intensive artificial intelligence applications [15], [20], [38], its use in environmental health decision-support systems remains relatively limited. The evaluation results indicate that the generated recommendations achieved high relevance (4.42/5), consistency (4.31/5), and scientific validity (4.48/5). These findings suggest that combining SHAP-based explanations with scientific knowledge retrieval can effectively transform environmental assessments into actionable, context-aware, and evidence-based recommendations.

Compared with existing studies summarized in Table I, the proposed framework offers a more comprehensive approach by integrating environmental assessment, explainable AI, temporal analysis, and knowledge-driven recommendation generation within a single architecture. Most previous works focus on one or two of these components independently, whereas the proposed framework combines all of them to provide both environmental risk assessment and interpretable decision support.

Despite these contributions, several limitations should be acknowledged. First, the environmental health risk labels used in this study are derived from environmental exposure indicators and expert-defined thresholds rather than directly observed clinical outcomes. Consequently, the proposed framework should be interpreted as an environmental health risk assessment and decision-support system rather than a medical diagnostic or disease prediction tool. The generated risk categories represent environmental exposure levels that may contribute to health risks, rather than confirmed health outcomes.

Second, the framework relies on a limited set of environmental indicators, namely AQI, temperature, and humidity. Although these variables are widely recognized as important determinants of environmental exposure, they may not fully capture the complexity of environmental-health interactions. Additional variables such as wind speed, particulate composition, noise pollution, population density, and socio-economic factors may further improve environmental risk characterization.

Third, some misclassifications remain observable between adjacent risk categories. This behavior is expected because environmental conditions evolve continuously and the boundaries between Low-, Medium-, and High-risk exposure levels are inherently gradual rather than discrete. Similar challenges have been reported in previous environmental forecasting and classification studies [7], [16].

Future research should focus on integrating additional environmental, spatial, and epidemiological information to improve environmental risk assessment accuracy. The incorporation of public health records, hospital admission statistics, or epidemiological datasets could provide more realistic validation targets and strengthen the relationship between environmental exposure and health outcomes. Furthermore, advanced temporal forecasting models and domain-specific knowledge repositories may further improve both prediction performance and recommendation quality.

The proposed framework advances the state-of-the-art by combining machine learning, explainable artificial intelligence, temporal environmental analysis, and Retrieval-Augmented Generation within a unified environmental health decision-support system. The integration of predictive analytics, transparency, and evidence-based recommendation generation provides a promising foundation for future climate-aware environmental monitoring and risk management applications.

VI. Conclusion

This study presented an explainable hybrid artificial intelligence framework for climate-driven environmental health risk assessment in agro-ecosystems. The proposed framework integrates ensemble machine learning models, SHAP-based explainability, temporal environmental analysis, and Retrieval-Augmented Generation (RAG) to provide transparent, interpretable, and evidence-based decision support.

Experimental results demonstrated the effectiveness of the proposed approach. Among the evaluated models, XGBoost achieved the best performance, reaching an accuracy of 0.88 and an AUC of 0.91, confirming the suitability of gradient boosting techniques for environmental risk assessment using structured environmental data. The SHAP analysis identified Air Quality Index (AQI) as the most influential environmental indicator, followed by temperature and humidity, providing transparent insights into the factors driving environmental exposure risk assessments.

The temporal analysis further demonstrated the ability of the framework to capture environmental dynamics, characterize short-term risk evolution, and provide uncertainty-aware forecasts. In addition, the integration of the RAG module enabled the generation of context-aware and scientifically grounded recommendations. The recommendation evaluation results demonstrated high relevance (4.42/5), consistency (4.31/5), and scientific validity (4.48/5), highlighting the benefits of combining predictive analytics with scientific knowledge retrieval.

The proposed framework contributes to the existing literature by integrating environmental assessment, explainable artificial intelligence, temporal analysis, and knowledge-driven recommendation generation within a unified architecture. Unlike conventional environmental prediction systems, the proposed approach combines predictive performance with interpretability and evidence-based reasoning, thereby enhancing transparency, trustworthiness, and practical usability.

Nevertheless, several limitations should be acknowledged. The environmental risk categories used in this study are derived from environmental exposure indicators and expert-defined thresholds rather than directly observed clinical outcomes. Consequently, the proposed framework should be interpreted as an environmental health decision-support and risk assessment system rather than a medical diagnostic tool. Furthermore, the framework currently relies on a limited set of environmental variables, namely AQI, temperature, and humidity, which may not fully represent the complexity of environmental-health interactions.

Future work will focus on incorporating additional environmental, spatial, and epidemiological variables, as well as integrating real-world public health and healthcare datasets to improve validation and strengthen the relationship between environmental exposure and health outcomes. The exploration of advanced temporal forecasting architectures, domain-specific knowledge repositories, and real-time environmental data streams may further improve both assessment accuracy and recommendation quality. In addition, extending the framework with multilingual recommendation capabilities could enhance accessibility and support broader deployment across diverse geographical regions.

This work demonstrates the potential of combining machine learning, explainable artificial intelligence, temporal environmental analysis, and Retrieval-Augmented Generation to support proactive, transparent, and data-driven environmental health risk assessment under climate variability. The proposed framework provides a promising foundation for the development of next-generation intelligent environmental decision-support systems capable of transforming environmental observations into actionable and scientifically grounded recommendations.

REFERENCES

[1] J. Lelieveld, J. S. Evans, M. Fnais, D. Giannadaki, and A. Pozzer, "The contribution of outdoor air pollution sources to premature mortality on a global scale," *Nature*, vol. 525, no. 7569, pp. 367–371, 2015, doi: 10.1038/nature15371.

[2] R. Burnett *et al.*, "Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter," *Proceedings of the National Academy of Sciences*, vol. 115, no. 38, pp. 9592–9597, 2018, doi: 10.1073/pnas.1803222115.

[3] Q. Di *et al.*, "Air pollution and mortality in the Medicare population," *New England Journal of Medicine*, vol. 376, no. 26, pp. 2513–2522, 2017, doi: 10.1056/NEJMoa1702747.

[4] X. Li, L. Peng, Y. Hu *et al.*, "Deep learning architecture for air quality predictions," *Environmental Science and Pollution Research*, vol. 23, pp. 22408–22417, 2016, doi: 10.1007/s11356-016-7812-9.

[5] W. Mao, W. Wang, L. Jiao, S. Zhao, and A. Liu, "Modeling air quality prediction using a deep learning approach: Method optimization and evaluation," *Sustainable Cities and Society*, vol. 65, p. 102567, 2021, doi: 10.1016/j.scs.2020.102567.

[6] G. K. Kang, J. Z. Gao, S. Chiao, S. Lu, and G. Xie, "Air quality prediction: Big data and machine learning approaches," *International Journal of Environmental Science and Development*, vol. 9, no. 1, pp. 8–16, 2018.

[7] Q. Liao, M. Zhu, L. Wu, X. Pan, X. Tang, and Z. Wang, "Deep learning for air quality forecasts: A review," *Current Pollution Reports*, vol. 6, no. 4, pp. 399–409, 2020.

[8] Z. Zhang, S. Zhang, C. Chen, and J. Yuan, "A systematic survey of air quality prediction based on deep learning," *Alexandria Engineering Journal*, vol. 93, pp. 128–141, 2024.

[9] A. Bekkar, B. Hssina, S. Douzi, and K. Douzi, "Air-pollution prediction in smart city using deep learning approach," *Journal of Big Data*, vol. 8, no. 1, p. 161, 2021, doi: 10.1186/s40537-021-00548-z.

[10] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016, doi: 10.1145/2939672.2939785.

[11] G. Ke *et al.*, "LightGBM: A highly efficient gradient boosting decision tree," in *Advances in Neural Information Processing Systems*, 2017.

[12] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.

[13] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017.

[14] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021, doi: 10.1016/j.ijforecast.2021.03.012.

[15] P. Lewis *et al.*, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Advances in Neural Information Processing Systems*, 2020.

[16] M. Karmoude *et al.*, "Machine learning for air quality prediction and data analysis: Review on recent advancements, challenges, and outlooks," *Science of The Total Environment*, p. 180593, 2025, doi: 10.1016/j.scitotenv.2025.180593.

[17] N. Fatima *et al.*, "Time series AQI forecasting using Kalman-integrated Bi-GRU and Chi-square divergence optimization," *Scientific Reports*, vol. 15, p. 12422, 2025, doi: 10.1038/s41598-025-12422-8.

[18] T. Tasioulis *et al.*, "The quest for the best explanation: Comparing models and XAI methods in air quality modeling tasks," *Applied Sciences*, vol. 15, no. 13, p. 7390, 2025, doi: 10.3390/app15137390.

[19] X. Zhou *et al.*, "Urban air quality management: PM2.5 hourly forecasting with POA–VMD and LSTM," *Processes*, vol. 13, no. 8, p. 2482, 2025, doi: 10.3390/pr13082482.

[20] B. Peng *et al.*, "Graph retrieval-augmented generation: A survey," *ACM Computing Surveys*, 2024, doi: 10.1145/3777378.

[21] Y. Ma *et al.*, "A survey of retrieval-augmented generation for large language models," in *IEEE International Conference on Information Technologies and Applications*, 2025, doi: 10.1109/ICTBAI68361.2025.00008.

[22] M. Bouni *et al.*, "Integrated IoT approaches for crop recommendation and yield prediction using machine learning," *IoT*, vol. 5, no. 4, p. 28, 2024, doi: 10.3390/iot5040028.

[23] D. Rajesh Kumar *et al.*, "On-device AI for climate-resilient farming with intelligent crop yield prediction using lightweight models on smart agricultural devices," *Scientific Reports*, vol. 15, p. 16014, 2025, doi: 10.1038/s41598-025-16014-4.

[24] F. M. Talaat *et al.*, "CardioRiskNet: A hybrid AI-based model for explainable risk prediction and prognosis in cardiovascular disease," *Bioengineering*, vol. 11, no. 8, p. 822, 2024, doi: 10.3390/bioengineering11080822.

[25] V. Gkouliaveras *et al.*, "Effects of climate change on health and health systems: A systematic review of preparedness, resilience, and challenges," *International Journal of Environmental Research and Public Health*, vol. 22, no. 2, p. 232, 2025, doi: 10.3390/ijerph22020232.

[26] C. De Guzman *et al.*, "Recent trends in research on the health impacts of climate change in the Philippines (2018–2023): A scoping review," *Frontiers in Climate*, vol. 7, p. 1594793, 2025, doi: 10.3389/fclim.2025.1594793.

[27] A. Samara *et al.*, "Impact of climate change and environmental adversities on maternal and fetal health," *Frontiers in Global Women's Health*, vol. 6, p. 1483938, 2025, doi: 10.3389/fgwh.2025.1483938.

[28] L. Gonah *et al.*, "The confluence of global climate change crisis, human health and health inequalities: A scoping review," *Global Policy*, vol. 17, 2025, doi: 10.52872/001c.137872.

[29] M. Shoaib *et al.*, "Air pollution and climate change: A dual threat to ecosystems and human health," *Journal of Environmental Challenges*, vol. 1, no. 1, 2025, doi: 10.62019/7w0pmn06.

[30] S. Bansal *et al.*, "Machine learning algorithms for predicting air quality index: A case study in urban and industrial zones," *Pollution Engineering*, vol. 13, no. 2, p. 354, 2025, doi: 10.21533/pen.v13.i2.354.

[31] R. Thangarajan *et al.*, "An AI based approach for smart air quality monitoring," in *IEEE International Conference on Contemporary Computing and Communication Technologies*, 2024, doi: 10.1109/ICCCNT61001.2024.10725962.

[32] A. Chinnaraju, "Explainable AI (XAI) for trustworthy and transparent decision-making: A theoretical framework for AI interpretability," *World Journal of Advanced Engineering Technology and Sciences*, 2025.

[33] C. S. Reddy and M. Annamalai, "Explainable AI for diabetes risk assessment: Enhancing clinical decision support with interpretable machine learning models," in *IEEE International Conference on Contemporary Computing and Communication Technologies*, 2025, doi: 10.1109/ICSCSA66339.2025.11171123.

[34] I. A. Aderemi *et al.*, "Explainable AI for water quality monitoring: A systematic review of transparency, interpretability, and trust," in *IEEE Symposium on Reliability*, 2025, doi: 10.1109/SR.2025.3595500.

[35] F. El Guma *et al.*, "Early-warning prediction of visceral leishmaniasis mortality using a multivariate STL–deep learning hybrid approach on 20 years of monthly time series," *Frontiers in Public Health*, vol. 14, p. 1754966, 2026, doi: 10.3389/fpubh.2026.1754966.

[36] M. Hamzah *et al.*, "Emerging trends in statistical, machine learning, and deep learning models for air quality prediction: A bibliometric analysis," in *IEEE International Conference on Computing Technologies and Innovation*, 2025, doi: 10.1109/ICCTEIE66144.2025.11341833.

[37] S. Rafi *et al.*, “AI document search chatbot using retrieval-augmented generation (RAG),” in *IEEE International Conference on Communication, Power and Computing Technologies*, 2026, doi: 10.1109/I5CPS67958.2026.11452596.

[38] R. Yousefi Maragheh *et al.*, “ARAG: Agentic retrieval augmented generation for personalized recommendation,” *arXiv preprint arXiv:2506.21931*, 2025, doi: 10.48550/arXiv.2506.21931.

[39] Z. Li *et al.*, “TutorLLM: Customizing learning recommendations with knowledge tracing and retrieval-augmented generation,” *arXiv preprint arXiv:2502.15709*, 2025, doi: 10.48550/arXiv.2502.15709.

[40] Y. G. Nikam *et al.*, “AI and IoT integration in smart farming: Real-time monitoring and predictive analytics for sustainable agriculture,” in *IEEE International Conference on Unmanned Aerial Systems*, 2025, doi: 10.1109/WorldSUAS66815.2025.11199045.

[41] Z. Han *et al.*, “An improved hybrid GC-LSTM framework for hourly nowcasting of ground-level NO2 concentrations over Beijing-Tianjin–Hebei region,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, 2025, doi: 10.1109/TGRS.2024.3514158.

[42] Y. Chudasama *et al.*, “Toward interpretable hybrid AI: Integrating knowledge graphs and symbolic reasoning in medicine,” *IEEE Access*, 2025, doi: 10.1109/ACCESS.2025.3529133.

[43] P. Roy *et al.*, “Interpretable artificial intelligence for cervical cancer risk analysis leveraging stacking ensemble and expert knowledge,” *SAGE Open Nursing*, vol. 11, 2025, doi: 10.1177/20552076251327945.

[44] A. Alshraideh *et al.*, “Hybrid AI framework for the early detection of heart failure: Integrating traditional machine learning and generative language models with clinical data,” *Cureus*, vol. 17, no. 6, 2025, doi: 10.7759/cureus.85638.

[45] P. K. Gupta *et al.*, “Data protection in enhanced explainable digital twins: Insights from hybrid AI,” in *IEEE International Conference on Communication Technologies*, 2025, doi: 10.1109/CICTN64563.2025.10932525.