Explainable AI for Climate-Smart Housing: A Multi-Hazard Recommendation System Using XGBoost and Cloud Technologies

Kowshik Padala, Rahul Thota, Teja Sai Sathwik P, Dhanush B, and Divya Udayan J*

Department of Computer Science and Engineering,
Amrita School of Computing, Amrita Vishwa Vidyapeetham, Amritapuri, India
am.en.u4cse22245@am.amrita.edu, am.en.u4cse22257@am.amrita.edu,
am.en.u4cse22271@am.amrita.edu, am.en.u4cse22501@am.amrita.edu,
divyaudayanj@am.amrita.edu

Abstract. This paper presents a robust, end-to-end AI-powered Climate-Smart Housing Recommendation System that evaluates global housing safety risks associated with floods, heatwaves, and droughts. Leveraging Google Earth Engine for multisource geospatial data acquisition and Google BigQuery for scalable processing, the system generates a comprehensive Climate Risk Index for candidate locations. We train and deploy an optimized XGBoost classifier, which achieves high predictive accuracy with an Area Under the Curve (AUC) of 0.92, proving effective even with significant class imbalance. To enhance trust and transparency, the model's predictions are explained using SHAP (SHapley Additive exPlanations), allowing stakeholders to understand the contribution of each risk factor. The entire pipeline is containerized and deployed on Google Cloud, supporting real-time, scalable inference. Users interact with a modern web-based interface built on Next.js, featuring interactive maps and safe zone recommendations. Furthermore, a conversational Gemini-powered AI assistant is integrated for user-friendly decision support. Our solution directly advances SDG 11 (Sustainable Cities and Communities) and SDG 13 (Climate Action), offering actionable insights for urban planners, policymakers, and individuals committed to building safer communities.

Keywords: Climate Risk · Multi-Hazard Assessment · Google Earth Engine · XGBoost · SHAP · Explainable AI · Gemini · Sustainable Housing · Geospatial AI · Cloud Deployment

1 Introduction

Climate change is widely recognized as one of the greatest challenges facing humanity in the 21st century. Increasingly frequent and severe climate hazards, including floods, heatwaves, and droughts, threaten not only public health and

* Corresponding author: divyaudayanj@am.amrita.edu

safety, but also the long-term sustainability of urban housing and infrastructure [1]. The United Nations projects that almost 70% of the world's population will reside in urban areas by 2050, increasing exposure and vulnerability to climate-induced disasters. In this context, ensuring climate-smart, safe housing has become a paramount concern for urban planners, governments, and individuals around the world.

Traditional approaches to housing safety and urban planning have relied on static historical datasets, broad zoning regulations, or generic risk categorizations. Such methods often fail to account for the rapidly changing and spatially heterogeneous nature of modern climate risks. Consequently, there is an urgent need for advanced dynamic solutions that can process real-time environmental data and deliver actionable location-specific recommendations.

Recent breakthroughs in artificial intelligence (AI), cloud computing, and geospatial data analytics are transforming the landscape of climate risk assessment. The emergence of scalable cloud native platforms enables the integration of multi-source satellite data and large-scale environmental datasets, facilitating the generation of actionable insights with unprecedented resolution and speed [2]. These advances support a shift from static, retrospective analyses to interactive, AI-driven systems capable of personalizing risk predictions and recommendations for diverse housing scenarios.

However, despite this technological progress, significant gaps remain. Many existing tools are limited in spatial coverage, lack real-time analytics, or do not provide user-friendly, interpretable interfaces. To effectively support stakeholders, from policy makers to everyday citizens, there is a clear need for comprehensive solutions that combine real-time data, explainable machine learning, and intuitive user experiences.

To address these challenges, this paper presents an AI-powered Climate-Smart Housing Recommendation System designed to deliver real-time, global scale safety assessments and recommendations for housing under the threat of floods, heatwaves, and droughts. The system combines advanced geospatial data pipelines, scalable machine learning models, and a user-centric web interface to empower decision-makers at every level. By supporting Sustainable Development Goals (SDGs) 11 (Sustainable Cities and Communities) and 13 (Climate Action), this work directly contributes to building safer and more climate-resilient urban environments around the world.

2 Related Work

Recent advances in climate risk analytics, explainable machine learning, and cloud-native platforms have enabled a new generation of solutions for urban resilience and sustainable housing. Here, we critically examine influential studies, emphasizing their core contributions, limitations, and how our approach advances the field.

Explainability in AI has gained significant attention in various domains. For instance, Harishankar et al. [16] introduced an explainable hybrid learning model for Indian food image classification, highlighting the importance of transparent AI in understanding model decisions. Similarly, George and Mohan Kumar [17] applied SHAP analysis to enhance CO_2 emission modeling, demonstrating how explainable AI can provide critical insights into environmental predictions and carbon risk modeling.

Roy et al. [8] proposed explainable AI (XAI) for building energy retrofitting, demonstrating the value of interpretability. Their work, however, is focused on energy efficiency, not climate hazard risk. We extend their principle by deploying SHAP, a game-theoretically robust XAI method, to provide not just general model insights but also precise, instance-level risk attributions which are critical for user trust in a safety-focused application.

Several studies address regional or single-hazard risks. The SaferPlaces Consortium [10] developed a platform for flood risk mapping in Europe, while Lee et al. [9] applied permutation feature importance for pluvial flood risk in South Korea. While permutation importance is useful for identifying critical features, it does not explain the magnitude or direction of a feature's impact on individual predictions. Our use of SHAP directly overcomes this, offering richer, more actionable explanations.

Other works like Springer et al. [7] and Alqahtani et al. [6] show strength in data fusion but do not provide real-time, global, or user-facing recommendation systems. These systems highlight the challenge of creating a platform that is at once analytically robust, scalable, and user-centric. Our work addresses this by choosing XGBoost for its superior performance and scalability on structured, tabular data, which allows us to model complex, non-linear interactions between multiple hazard inputs effectively.

The value of user-centric design is underscored by works from Raghu Vamsi et al. [11], Tamil Priya and Udayan [12], and Devika and Udayan [13], which, despite focusing on human-computer interaction, inspire the conversational and intuitive interface of our platform. Furthermore, Anand et al. [18] explored techno-social synergy for disaster resilience in coastal communities, emphasizing community-centered design and sustainable approaches, which aligns with our platform's goal of fostering disaster preparedness through accessible information and a conversational AI assistant. Finally, works by Akash et al. [14] and the survey by Doe et al. [15] confirm the need for operational, end-to-end platforms that move beyond reactive or regionally constrained analytics.

Summary of Gaps and Our Contribution: While these studies mark substantial advances, they are often limited by a single-hazard focus, regional constraints, or a lack of real-time, actionable engagement. Furthermore, many proposed models based on deep learning or federated approaches present scalability challenges due to computational or communication overhead. Our platform directly addresses these gaps by unifying cloud-scale geospatial analytics (GEE and BigQuery) with an optimal combination of interpretable, high-performance machine learning. Specifically, we selected **XGBoost** for its proven scalability and exceptional accuracy in handling complex, tabular geospatial data, and integrated **SHAP** to provide the transparent, trustworthy, and feature-level explanations that are essential for a globally scalable and user-centric housing risk recommendation system.

3 Dataset

The proposed Climate-Smart Housing Recommendation System leverages a multi-source, global dataset that captures key climate and geographic risk factors relevant to housing safety. All data sources are chosen for their reliability, spatial coverage, and compatibility with large-scale geospatial analysis through Google Earth Engine.

- **Flood Risk:** Data is sourced from the JRC Global Surface Water dataset, which provides global-scale information on the occurrence, extent, and temporal dynamics of surface water bodies. This dataset enables identification of flood-prone regions based on historical flood events and changes in water coverage.
- **Heat Risk:** The system uses MODIS Land Surface Temperature (LST) data, collected via NASA's Terra and Aqua satellites. This dataset delivers regular, high-resolution surface temperature readings, making it suitable for detecting regional heat anomalies and identifying areas at risk of heatwaves.
- **Drought Risk:** Drought indices are obtained from the TerraClimate Palmer Drought Severity Index (PDSI), which measures meteorological and agricultural drought severity globally. PDSI integrates precipitation, temperature, and soil moisture, providing a comprehensive indicator for prolonged dry spells.
- **Ocean Disaster and Natural Hazard Risk:** To extend risk assessment beyond land, we incorporate ocean-based natural disaster datasets from GEE, which help identify hazards such as cyclones and sea-level anomalies that could threaten coastal housing.
- **Safe Zones and Elevation:** Aggregated housing data and global elevation models (such as SRTM DEM) are utilized to filter and identify locations that are naturally less vulnerable to flooding and other hazards due to favorable topography.

Each dataset is ingested and preprocessed via Google Earth Engine, allowing for spatial normalization and alignment. The datasets vary in their spatial and temporal resolution: for example, MODIS LST provides daily observations at 1km resolution, while JRC Surface Water offers yearly and monthly global maps. To ensure data integrity and minimize gaps, preprocessing steps include resampling, outlier removal, and masking of waterbodies and non-inhabitable regions. The combined dataset covers multiple years (typically 2000–2023) and enables both retrospective analysis and near real-time risk assessment. By fusing these data sources, the system generates a composite Climate Risk Index (CRI) for each candidate housing location, supporting robust, explainable predictions across diverse geographic contexts. The resulting dataset exhibits a natural class imbalance reflective of real-world conditions, with a larger proportion of global land area classified as safe compared to high-risk zones.

4 Methodology

The architecture of our Climate-Smart Housing Recommendation System is designed as a modular pipeline integrating Data, Processing, AI/ML, Interface, and Visualization layers. Each component contributes to a robust, scalable, and interpretable workflow for multi-hazard climate risk analysis and recommendation. The overall system architecture and cloud-native design, including Vertex AI model serving, are depicted in the System Design section (see Fig. 1). Here, we detail the functional roles and technical flow across each architectural layer.

4.1 Data Layer

Our system collects and integrates heterogeneous climate and geographic datasets from global sources. Key inputs include:

- **Flood risk:** JRC Global Surface Water occurrence, capturing historical and seasonal flood patterns.
- **Heat risk:** MODIS Land Surface Temperature (LST), monitoring heatwaves and surface temperature anomalies.
- **Drought risk:** TerraClimate Palmer Drought Severity Index, providing drought severity on a global grid.
- **Natural disaster context:** Global ocean shapefiles and elevation models to filter and mask non-inhabitable or high-risk coastal areas.

Data ingestion and preprocessing are managed with Google Earth Engine (GEE), which enables large-scale geospatial queries [3], [4]. It is important to note that while our system provides *real-time inference* upon a user's request, the underlying datasets are updated periodically according to their source (e.g., daily for MODIS LST, monthly or annually for others), not in real-time.

4.2 Processing Layer

Each climate risk variable is normalized to a common scale and combined into a composite Climate Risk Index (CRI), computed as:

$$CRI = w_1 \cdot F + w_2 \cdot H + w_3 \cdot D$$

where F, H, and D denote normalized flood, heat, and drought risk scores, and w_1, w_2, w_3 are tunable weights. These were set empirically for this study as $w_1 = 0.4$, $w_2 = 0.3$, and $w_3 = 0.3$ based on established hazard precedence. We acknowledge that a sensitivity analysis on these weights is a valuable direction for future work. Additional filtering is performed to exclude ocean regions and uninhabitable areas using shapefile overlays. Data is stored and queried in Google BigQuery, supporting low-latency queries for real-time inference and larger-scale batch analytics [5].

4.3 AI/ML Layer

We utilize an XGBoost classifier as the core predictive engine, chosen for its superior performance on imbalanced, tabular data and its ability to capture nonlinear interactions among hazards. The model is trained on a globally diverse dataset; our final test set, for instance, consists of **14,522 locations** with a significant class imbalance of **10,246 Safe (class 0)** to **4,276 Risky (class 1)**. To optimize performance, extensive hyperparameter tuning was conducted using grid search, with key parameters including `n_estimators`, `max_depth`, and `learning_rate`. For serving, the trained model is deployed to **Google Vertex AI** as a fully managed endpoint. A containerized FastAPI backend, orchestrated via Google Cloud Run, mediates user requests, prepares feature vectors, and interacts with the Vertex AI endpoint. This architecture decouples training from production and enables scalable, low-latency inference. To ensure transparency, a SHAP (SHapley Additive exPlanations) module is integrated into the backend, providing instance-level insights into the factors driving each risk prediction [9]. The insights gained from SHAP analysis are crucial for understanding how features, such as those related to flood, heat, and drought, contribute to the overall climate risk, as similarly demonstrated in carbon emission modeling by George and Mohan Kumar [17].

4.4 Interface Layer

The frontend is developed with Next.js and Tailwind CSS for responsive design, and Folium/Leaflet for interactive maps. Users can input addresses, use live geolocation, or click on the map to receive instant risk assessments and recommendations. The interface displays dynamic safety scoring, nearest safe zone suggestions, and a contextual breakdown of predicted risk. We also integrate a chatbot powered by Gemini Pro (Google's LLM) to answer user questions, assist navigation, and explain model results in plain language. This integration aligns with the broader goal of community-centered design and disaster preparedness, as discussed by Anand et al. [18].

4.5 Visualization and Analytics Layer

Comprehensive visualization is provided through interactive Folium-based maps, global and regional heatmaps, and analytics dashboards built with Matplotlib. The system supports real-time exploration of safe and risky zones, pie charts summarizing global safety distribution, and downloadable reports for both individual users and planners. Visualizations also include feature importance plots and SHAP summaries to support model interpretability for end-users and policymakers, an approach that builds on advancements in explainable AI for complex systems, such as the hybrid learning model proposed by Harishankar et al. [16].

5 System Design

The revised architecture of our Climate-Smart Housing Recommendation System (see Fig. 1) is designed as a modular, cloud-native pipeline that integrates robust geospatial data processing, scalable AI/ML inference, and explainable analytics within a single unified workflow. The system is engineered for global scalability, transparency, and responsiveness, ensuring reliable climate risk analysis and real-time recommendations for diverse end-users.

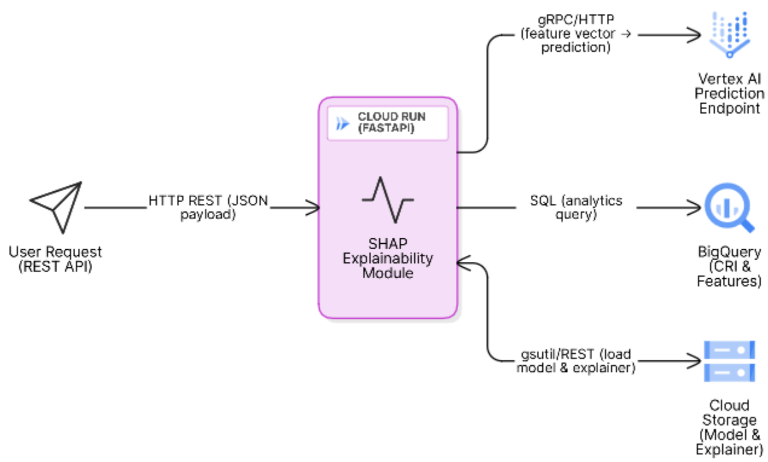

Fig. 1. Overall System Architecture: A modular pipeline from user request to data ingestion, cloud processing, ML prediction, and explainable feedback.

5.1 Modular System Pipeline

- **Data Ingestion Layer:** Multi-source climate and geographic datasets—including JRC Global Surface Water, MODIS Land Surface Temperature, and TerraClimate Palmer Drought Severity Index—are ingested and preprocessed through Google Earth Engine (GEE). Data is normalized and aligned, with features engineered for downstream analysis.
- **Processing and Storage:** Preprocessed data and derived features (such as the Climate Risk Index) are stored in Google BigQuery, facilitating scalable queries and on-demand analytics. Additional features and model artifacts are managed securely in Google Cloud Storage for traceability and version control.
- **Analytics and Visualization:** BigQuery supports advanced spatial and temporal analytics, while interactive dashboards and maps (built with Folium and Matplotlib) allow users and policymakers to visualize risk distributions and safe/risky zones at various geographic scales.

- **User Interface and Feedback:** The frontend web application supports real-time location queries, safe zone recommendations, and chatbot guidance. End-users receive both numerical risk scores and visual explanations, facilitating informed decision-making.

5.2 Vertex AI Integration for Model Serving

A critical innovation of the system is the integration of Google Vertex AI for model deployment and inference. The backend sends prepared feature vectors to the Vertex AI endpoint, which returns risk predictions in real time. This approach ensures reliability, scalability, and ease of maintenance, with built-in support for automated versioning, model monitoring, and secure access management. By decoupling model training (performed offline) from serving (managed on Vertex AI), the platform supports continuous improvement and safe model upgrades.

5.3 Explainability and Transparency

For each user request, the backend computes SHAP-based explanations, quantifying how individual features (such as drought index, flood occurrence, heat anomaly, latitude, and longitude) contribute to the model's output. These explanations are delivered alongside predictions, building trust and supporting regulatory compliance in critical, safety-related applications. This is in line with recent advancements in explainable AI, such as the SHAP analysis for CO_2 emission modeling by George and Mohan Kumar [17] and the explainable hybrid learning models discussed by Harishankar et al. [16].

In summary, the system design unites advanced geospatial processing, managed AI services, and explainable analytics within a secure, modular, and cloud-native architecture, maximizing usability, transparency, and real-world impact for global climate-smart housing recommendations.

6 Results and Discussion

The proposed XGBoost-based classifier for climate-smart housing safety was evaluated on a large, multi-continental dataset comprising both safe and risky locations. To ensure robustness and generalizability, we split the dataset into training and test sets with stratification on the class label due to inherent class imbalance.

6.1 Quantitative Evaluation

The performance of the model is best reflected by its confusion matrix (see Fig. 2), ROC curve (Fig. 3), and a detailed classification report (Table 2). As observed in the confusion matrix, the model achieves high true positive and true negative rates, but, like all real-world climate risk classifiers, displays a tradeoff in false positives due to overlapping risk signals and class imbalance. The ROC curve illustrates a strong separation capability, with an Area Under the Curve (AUC) of 0.92—demonstrating the model's robustness across risk thresholds.

Baseline Model Comparison To further validate our choice of XGBoost, we compared its performance against two standard baseline classifiers: Logistic Regression and a Support Vector Machine (SVM). As shown in Table 1, while simpler models like SVM can achieve perfect recall for the 'Risky' class, this comes at the cost of extremely low precision (0.16), leading to an impractical F1-score (0.27). This indicates that the SVM model generates an excessive number of false positives, rendering it unsuitable for a real-world recommendation system. In contrast, our XGBoost model (Table 2) provides a far superior balance, achieving a high recall of 0.99 for the risky class while maintaining a strong precision of 0.67. The resulting F1-score of 0.80 confirms that XGBoost is significantly more effective at accurately identifying high-risk locations without overwhelming users with false alarms, making it the clear choice for this safety-critical application.

Table 1. Baseline Model Comparison for the 'Risky' Class.

Model	**Precision**	**Recall**	**F1-score**	**AUC**
Logistic Regression	0.17	0.87	0.29	0.92
Support Vector Machine (SVM)	0.16	1.00	0.27	0.86

Table 2. Classification Report (XGBoost).

Class	Precision	Recall	F1-score	Support
0 (Safe)	1.00	0.80	0.88	10246
1 (Risky)	0.67	0.99	0.80	4276
Avg/Total	**0.90**	**0.85**	**0.86**	**14522**

6.2 Model Explainability and Feature Attribution

To address the critical need for interpretability in safety-critical systems, we employ SHAP (SHapley Additive exPlanations) analysis. The summary plot in Fig. 4 visualizes the contribution of each input feature to model predictions, revealing that drought, flood, and heat indices are the primary determinants of risk classification, while geographic coordinates (longitude, latitude) add nuanced regional context. This aligns with the importance of explainability in various AI applications, as demonstrated by Harishankar et al. [16] in their work on image classification and George and Mohan Kumar [17] in their carbon emission modeling.

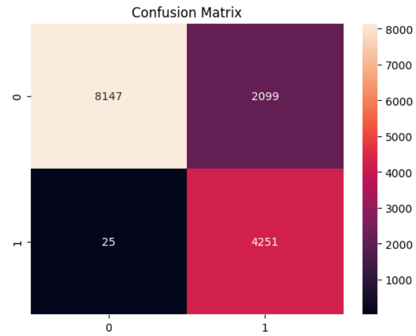

Fig. 2. Confusion Matrix for the XGBoost model.

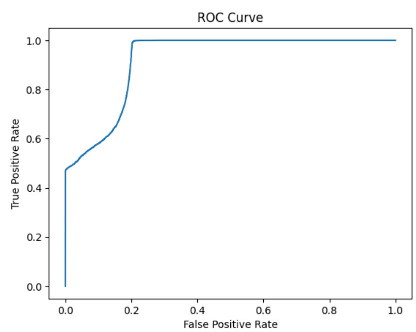

Fig. 3. ROC Curve for the XGBoost model (AUC = 0.92).

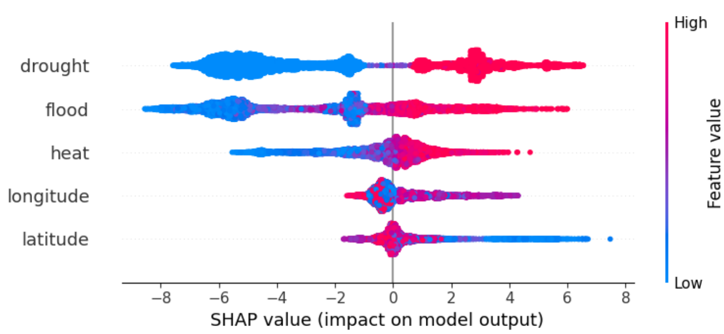

Fig. 4. SHAP Summary Plot: Feature Importance and Effect on Risk Prediction. Drought, flood, and heat are dominant features.

6.3 Discussion

Our results demonstrate that deploying a cloud-native, XGBoost-based model with SHAP explanations provides robust, interpretable safety guidance at a global scale. This approach complements XAI frameworks developed for critical infrastructure by Roy et al. [8], by making explainability a core feature of end-user climate decision support. Moreover, the end-to-end, cloud-integrated architecture aligns with the recommendations of Doe et al. [15] for scalable, real-world machine learning deployment in climate risk domains. By providing both actionable predictions and transparent explanations, our system overcomes a major limitation noted in regional, single-hazard platforms such as SaferPlaces [10], enabling personalized, multi-hazard risk assessment for both policymakers and the general public.

The strong model performance and feature explainability—combined with modular cloud deployment—illustrate the practical viability of large-scale, real-time climate risk assessment and recommendation. Further, our user-centric approach draws on insights from Raghu Vamsi et al. [11], Tamil Priya and Udayan [12], and Devika and Udayan [13] to provide interpretable results and conversational guidance, paving the way for even more adaptive, emotion-aware, and trustworthy climate resilience platforms in the future. The emphasis on community-centered design and conversational AI, as discussed by Anand et al. [18], further strengthens the real-world applicability of our platform for disaster preparedness.

7 Conclusion

This work introduces an AI-powered Climate-Smart Housing Recommendation System that combines cloud-native geospatial analytics, scalable machine learning, and explainable AI to deliver real-time, actionable risk insights. By integrating multi-source climate data via Google Earth Engine and BigQuery, and deploying an optimized XGBoost model on Vertex AI, our system translates large-scale environmental data into personalized, interpretable housing recommendations.

A standout feature is the SHAP-based explainability module, which enhances transparency in model decisions—critical for high-stakes applications. This aligns with advances in explainable AI for domains such as food classification [16] and carbon modeling [17]. The interactive web interface and Gemini-powered assistant [18] further support user engagement and disaster preparedness.

Limitations include reliance on periodically updated datasets, potential biases in sparse regions, and the costs of maintaining global-scale infrastructure. Future work will focus on expanding hazard coverage, enabling multi-class prediction, and validating usability through pilot deployments with planners and policy stakeholders.

Ultimately, this platform aims to become a scalable, trustworthy tool for supporting climate resilience and sustainable housing decisions worldwide.

References

1. Reichstein, M., et al.: Early warning of complex climate risk with integrated artificial intelligence. Nature Communications, 16(1), 2564 (2025)
2. Razavi-Termeh, S.V., et al.: Improving flood-prone areas mapping using geospatial artificial intelligence (GeoAI): A non-parametric algorithm enhanced by math-based metaheuristic algorithms. Journal of Environmental Management, 375, 124238 (2025)
3. Nandalal, K.D.W., et al.: Google Earth Engine for Large-scale Drought Monitoring. Remote Sensing (2021)
4. Sellami, E.M., Rhinane, H.: Google Earth Engine and Machine Learning for Flash Flood Exposure Mapping–Case Study: Tetouan, Morocco. Geosciences, 14(6), 152 (2024)
5. Google Cloud: Google Earth Engine for planetary-scale environmental monitoring. Google Cloud Blog (2022)
6. Alqahtani, D., et al.: Optimizing Residential Construction Site Selection in Mountainous Regions Using Geospatial Data and eXplainable AI. Sustainability, 16(10), 4235 (2024)
7. Springer, A., et al.: Biophysical Drivers using ML and Google Earth Engine. Springer Nature (2025)
8. Roy, R., et al.: Explainable AI for Building Energy Retrofitting. arXiv preprint arXiv:2504.06055 (2025)
9. Lee, E., et al.: Explainable AI-based risk assessment for pluvial floods over South Korea. Journal of Environmental Management, 385, 125640 (2025)
10. SaferPlaces Consortium: SaferPlaces: AI-driven Urban Flood Planning Platform. Project Report (2023)
11. Raghu Vamsi, U., et al.: Speech Emotion Recognition (SER) using Multilayer Perceptron and Deep Learning Techniques. IEEE Access, 27(5), 386–394 (2021)
12. Tamil Priya, D., Divya Udayan, J.: Transfer learning techniques for emotion classification on visual features of images in the deep learning network. International Journal of Speech Technology, 23(2), 361–372 (2020)
13. Devika, R., Divya Udayan, J.: Lightweight Deep Learning Facial Expression Recognition Model. In: ICT Systems and Sustainability: Proceedings of ICT4SD 2022, pp. 493–503 (2022)
14. Akash, et al.: Drone-based Climate Resilience Housing Mapping. arXiv preprint arXiv:2312.10306 (2023)
15. Doe, J., et al.: Machine Learning Potential for Climate Risk Assessment. ScienceDirect (2021)
16. Harishankar, A., Anitha, M.S., Vaishnavi, V.: An Explainable Hybrid Learning Model for Indian Food Image Classification. In: Proc. ICCCNT, IEEE (2024)
17. George, M., Mohan Kumar, A.: Carbon Prognosticator: A Triple Ensemble Regressor and SHAP Analysis for Enhanced CO Emission Modeling. In: Proc. ICUIS, IEEE (2024)
18. Anand, M., Vinayaka, P., Krishnan, S.: Techno-Social Synergy for Disaster Resilience in Coastal Communities: A Sustainable Approach. In: Proc. ICDCN, ACM (2023)