

Enhancing Decision Making Across Industries with SHAP

Sara Sowmya¹, Soumita Sen², Geethesh Chowdary³

^{1,2,3}*Research Intern, Woxsen University, Hyderabad, Telangana, India*

Abstract—Explainable Artificial Intelligence is crucial for building trust, transparency and accountability in decision-making. Explainable AI (XAI), and SHAP (SHapley Additive exPlanations) in particular, is vital in making sophisticated AI models more interpretable and more trustworthy. For example, in agriculture, farmers can see the effects of soil quality, rainfall, and nutrient quality on crops more clearly with XAI, and ultimately make more data-driven decisions around those factors and encourage sustainable farming choices. In the context of residential energy management, SHAP made it more interpretable by clearly showing how much temperature and humidity patterns and occupancy, and times of the day, contribute to energy consumption predictions, giving households and policymakers the ability, they need to further their own energy efficiencies. In industrial contexts, XAI supports safety by identifying the most important sensory signals and environmental information that can inform the detection of toxic gas leaks and for individuals or organizations to respond faster to identified hazards - due to XAI providing cogent information. As demonstrated in each example, SHAP helps improve predictions and identify the factors that strengthen trust, accountability, and fairness in AI systems, in order to support decision-making in the real world.

Index Terms—Explainable Artificial Intelligence (XAI), SHAP (Shapley Additive Explanations), Machine Learning Interpretability, Predictive Analytics, Decision Support Systems.

I. INTRODUCTION

The rapid development of Artificial Intelligence (AI) and Machine Learning (ML) have transformed industries, enabling automation, predictive analytics, and smarter decisions based on information. Whether diagnosing from x-ray scans in healthcare or estimating the yield of corn on the farm for agriculture, AI systems have demonstrated whether they can make sense of giant datasets and provide value, whether in diagnosis,

predicting agricultural outcomes, or improved industrial decision-making. However, one of the biggest obstacles to effective use of AI is interpretability, despite the advantages AI has provided industries. Many models, and particularly ensemble methods such as Random Forest Regressors, are sometimes called "black boxes," because there is no way to understand how they arrived at their decisions. This uncertainty creates distrust for all stakeholders but also limits the value of an actionable output from AI. Issues of ethical concern or regulatory justification may also arise when AI is used in industries like healthcare, manufacturing, agriculture, and energy.

The need for transparency becomes especially important in contexts where the *explanation* behind a prediction carries equal or greater weight than the prediction itself. For example, in healthcare, clinicians rely on explanations of EEG signals or genetic markers to understand diagnostic outcomes. In industrial environments, knowing which sensors most influence toxic gas leak predictions are essential for ensuring safety and guiding intervention. Farmers, similarly, benefit from understanding how soil, climate, and fertilization conditions contribute to yield forecasts, while energy managers require interpretable consumption models to improve efficiency at both household and policy levels. Unfortunately, most state-of-the-art models lack the flexibility and transparency needed for real-time, interpretable decision-making.

To address these gaps, Explainable AI (XAI) has emerged as a crucial area of research, offering methods to improve trust, accountability, and interpretability in AI-driven systems. In this study, SHAP (SHapley Additive exPlanations), a game-theoretic framework, is applied to Random Forest models to uncover feature contributions and provide transparent reasoning behind predictions. Using interactive Streamlit applications, SHAP explanations are made accessible in real time, offering actionable insights across use cases such as

toxic gas leak identification, crop yield forecasting, and residential energy consumption. The overarching goal of this work is to close the interpretability gap, foster trust among stakeholders, support ethical and regulatory compliance, and strengthen the role of AI in advancing sustainable development.

Problem Statement

Industries today are facing significant challenges that require not only accurate predictions but also easy to understand explanations. The detection of a gas leak in a factory, the continued production of sustainable crops, and the management of household energy demands is pressing problems. Decisions cannot rely upon "black-box" AI - Explainable AI (XAI) provides an alternative by making predictions interpretable and actionable.

Manufacturing: Detection of toxic gas leaks, wherein XAI will support industry identifying noxious emissions based on tracking substitution signal based on sensors on environmental conditions.

Agriculture: Predicting crop yield forecast based on soil properties, climate conditions, and fertilizer applications. XAI will provide support to farmers in making decisions.

Energy: Forecasting residential energy use is an application domain where XAI can support homeowners by indicating how factors that vary by the household (temperature, humidity, occupancy, and day of week) contribute to energy use increases or decreases.

- **Facilitating Human-AI Interaction:** AI must work in collaboration with humans. Making AI explainable empowers experts across industries to understand the recommendations provided by AI models, enabling them to make more conscious, informed, and trustworthy decisions.

1.1 Applications of Explainable Artificial Intelligence (XAI)

XAI is essential in several industries where decision-making needs to be interpretable and transparent.

- **Healthcare:**

AI models are utilized in assisting physicians diagnosing diseases and prescribing relevant treatment. With XAI, physicians get insight into the rationale behind the recommendations of the AI, which help commit to patient safety and better healthcare outcomes.

- **Finance:**

AI-driven models are used for credit scoring, fraud detection, and risk assessment within finance sectors. XAI methods ensure financial organizations encourage fairness, detect biases, and ultimately remain compliant with regulatory agreements, thereby minimizing the risk of discrimination in lending practices.

- **Toxicology:**

AI is employed to test for the predictive toxicity of chemicals and pharmaceuticals. The ability to explain the models, regardless of output, supports regulators and improves decision-making in drug development and environmental health.

- **Agriculture:**

AI-driven precision farming, predicting crop disease, and yield optimization will generally be gained from explainability in decision-making and assure farmers and agronomists are following data-driven and sustainable practices.

- **Energy:**

The managing and optimization of smart grid load balancing will be enhanced as the transparency of XAI leads energy providers to a greater understanding of their regulatory and stakeholder obligations and how factors like, time of use, weather conditions, and consumer behavior patterns affect energy demand forecasting rationalizing the ultimate transparent and adaptive decisions on grid management through the AI computing models.

As the prevalence of AI in industries such as agriculture, finance, and healthcare rises, the need for computers that are interpretable and explainable increases. This research demonstrates the importance of explainability techniques such as SHAP as a way to ensure AI models arrive at decisions on claims that are clear, understandable, and justifiable. Using SHAP implicitly through accessible models will promote developing ethical and trustworthy AI systems that create positive impacts on society through improved transparency of AI.

II. LITERATURE REVIEW

2.1 Improving Transparency in Agricultural Predictions Using Explainable AI

Recent advances in machine learning have greatly improved predictive accuracy in agriculture, yet most models lack interpretability, resulting in a system rarely trusted or understood by farmers and researchers. This drives interest to Explainability Artificial Intelligence (XAI), which aims to add transparency to complex models. While some domains, such as healthcare and environmental science, have adopted XAI, it remains limited in agriculture. Recent studies have pointed out that machine learning (ML) is increasingly used for crop yield prediction, with the advantage over traditional statistical models. Research employing CNN, RNN, and hybrid techniques has shown precise prediction of natural disasters when using a combination of weather, soil, and remote sensing information. Yet these models tend to be black boxes, which reduces transparency and farmer trust. To address this issue, explainable AI (XAI) techniques like LIME are developed to explain model predictions. These advances notwithstanding, however, the majority of crop recommendation systems remain poorly interpretable and untrustworthy to users. Herein, we present XAI-CROP, a model that improves upon traditional machine learning (ML) models to add explainability and transparency to agricultural decision-making while ensuring ML accuracy.

2.2 Operational Interpretability in Industrial AI: SHAP for Predictive Maintenance and Load Forecasting

Using explainable AI methods, the study focuses on explaining results obtained by predictive maintenance models of industrial raw materials, with an emphasis on SHAP. It includes XGBoost and local outlier factor algorithms, giving very high prediction efficiency with F1 scores of 96% and 94% respectively. But SHAP is a system for explaining any model, giving a measure of how much each feature contributes to the output it classifies, increasing the transparency and reliability of a model. Our approach of XAI guides users in making informed real-time decisions for system optimization and predictive maintenance within a reasonable time by showing them exactly how the model is building its decision. It overcomes some limitations of traditional SHAP-aided contribution-based feature importance metrics by highlighting one that incorporates how evenly SHAP values are positioned across the feature distributions using the coefficient of determination. 24-Hour Load Forecasting Based on XGBoost with Korea Power Exchange Data. The results show that our SHAP-based metric on the features compares favorably with

the baseline and improves both model performance and interpretability by giving a clearer picture of the features that matter most to accurate load prediction.

Methodology

Research Design

This research uses a machine-learning approach by integrating Explainable Artificial Intelligence (XAI) to maintain transparency and interpretability. It makes use of the Random Forest algorithm combined with SHAP which helps quantify each feature from across different problem domains. The key objective of this research is to develop models that are not just accurate but also interpretable, which makes them useful for decision-making in areas such as healthcare, industrial monitoring and agriculture.

2.3 Data Collection

The approach to data collection utilizes synthetic data generation methods developed as part of the related problem statement. Synthetic Data Generation-Based Approach:

- Prediction of Toxic Gases – Sensor readings generated using physics-based models, or statistical models. Environmental factors (temperature, wind speed, humidity) also simulated as distributions based on the real world.
- Agricultural Prediction – Soil properties, climate factors, and crop yield values based on agricultural data sets. Using stochastic models or regression-based simulations to maintain correspondence with the real-world characteristics.
- Energy Consumption – Predicting household energy consumption using a synthetic dataset to create simulated patterns of household usage (e.g., different occupancy, appliance usage, and weather patterns) to increase model generalization and XAI-based interpretations of energy demand behavior.

2.4 Machine Learning Algorithm

Random Forest: The Random Forest model is used as the predictive model in this study because of its ability to handle non-linear relationships between features, making it flexible and easier to use. The model works efficiently with high-dimensional datasets, has the

ability to handle missing values and noisy data, and can reduce overfitting of the model by combining multiple decision trees. It is an ensemble learning model that constructs multiple decision trees using random subsets of the same training data. It then either takes the average of all the predictions from the decision trees and accepts the values as the prediction, or it uses a majority voting system, depending on the classification task. This model increases robustness and accuracy by reducing the variance of individual trees.

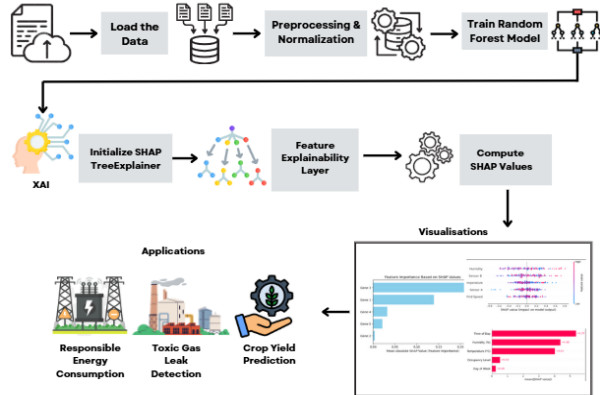


Fig 1: XAI Model Workflow

The workflow, as illustrated in fig. 1, by loading and pre-processing the data before building a Random Forest model and using SHAP Tree Explainer to get feature-level interpretability. The explainability layer will calculate SHAP values and then visualized for interpretability in domain-specific contexts such as energy consumption forecasts, gas leak detection, and crop yield prediction, etc.

2.5 Explainability Framework: SHAP (Shapley Additive Explanations)

Random Forest models are considered ‘black box’ models because they take inputs and provide outputs with no explanation of their internal processes, which leaves users open to interpretation of the backend workings. To solve this issue, we integrate SHAP (Shapley Additive Explanations), an Explainable Artificial Intelligence (XAI) technique that quantifies the importance of each feature by assigning each of them individual numerical values. This provides a fair and consistent feature attribution, which ensures transparency. This provides immense help to domain experts such as doctors, engineers and farmers by providing AI-driven insights.

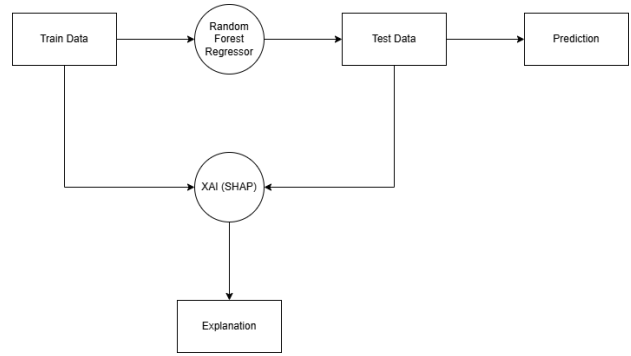


Fig 2: SHAP-Based Interpretability for Random Forest Models

2.6 Mathematical Formulation

SHAP values are calculated using the following equation:

$$\Phi(x_i) = \sum_{S \subset N} \frac{|S|(M-|S|-1)!}{M!} [f(S \cup \{x_i\}) - f(S)]$$

$\Phi(x_i)$: Represents the SHAP value i.e., how much each feature (Sensor A, Sensor B, Temperature, Wind Speed, Humidity) contributes to predicting the gas leak risk.

S : Subset consisting of a combination of all features except the one being analysed.

M : Total number of features considered.

$f(S \cup \{x_i\}) - f(S)$: The marginal contribution of a feature. It measures the change in prediction when the feature is added to a subset of other features.

$\frac{|S|(M-|S|-1)!}{M!}$: The weight factor for each contribution. It ensures that each feature’s contribution is fairly averaged across all possible feature orderings, so no feature gets unfair advantage or bias.

2.7 Feature Importance & Explainability

The ability to interpret the results is as important as achieving high accuracy in machine learning models such as Random Forest. SHAP (Shapley Additive Explanations) provides a mathematical approach to understanding how different features influence the model's predictions.

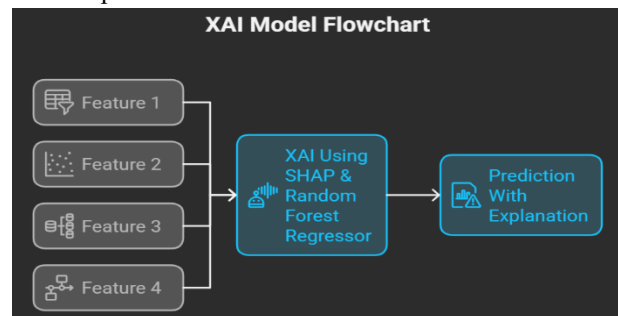


Fig 3: XAI Model

In a traditional setting, feature importance is typically determined by measuring how much a certain feature reduces the prediction error when it is included in the model. This method, however, lacks explainability at the individual prediction level which is the reason why SHAP is used. SHAP improves this feature importance by:

- Assigning a numerical value to each feature – This explains how much each feature influences the model's output.
- Ensuring fair attribution – It uses a cooperative game theory to fairly distribute contributions among the features.
- Providing local and global interpretability – It explains both overall feature importance and how individual instances are influenced

2.8 Ethical Considerations & Advantages of Synthetic Data Usage

As machine learning and AI advance to be integrated into regular decision-making and people's everyday lives, it is important to ensure ethical integrity. Upholding ethical principles was a top priority in this research, which considers sensitive fields like industrial safety and agriculture. This research exclusively used synthetic data for all four problem statements to maintain privacy, ensure data availability and control feature distributions. The key ethical aspects considered include data privacy, bias and fairness, and transparency in AI-driven decision-making.

Using synthetic data protects the patient's and the user's privacy. There are many global regulatory bodies like GDPR, and HIPAA, in place to protect this type of data. Using this dataset helps control the experimentation. You can do feature distribution to reflect real-world conditions. And since no real human or industrial data was used, any consent-related issues were solved.

2.9 Bias and Fairness in AI Predictions

AI models may unintentionally reinforce biases present in training data, which leads to unfair or misleading predictions. Since synthetic data was used, biases in real-world datasets were avoided, but these factors were still considered. In EEG-based disorder detection, different brainwave patterns were balanced to ensure no single type of seizure-dominated predictions, in genetic disorder prediction, synthetic genes were modelled across multiple population groups to avoid ethnicity-

related bias, in toxic gas leak detection, SHAP feature attributions were examined to prevent certain sensors from being disproportionately favoured, and in crop yield forecasting, multiple environmental conditions (rainfall, temperature, soil quality) were equally represented in the synthetic data to ensure fairness. By integrating synthetic data generation with ethical AI practices, this research ensures that the XAI-supported Random Forest Regressor remain transparent, fair, and reliable across all domains.

III. APPLICATION DOMAIN

XAI-Driven Toxic Gas Leak Source Identification in Industrial Plants

Input features:

a. Sensor A:

- Represents: Gas concentration detected by Sensor A, located close to critical equipment.
- Unit: Parts per million (ppm).
- Relevance:
 - Higher values indicate a higher likelihood of a gas leak in proximity to Sensor A.
 - Serves as a primary indicator for leaks.

b. Sensor B:

- Represents: Gas concentration detected by Sensor B, placed further from Sensor A for cross-validation.
- Unit: Parts per million (ppm).
- Relevance:
 - Complements Sensor A by validating leak spread or source overlap.
 - Helps distinguish between multiple leak sources.

c. Temperature:

- Represents: Ambient temperature around the plant during sensor readings.
- Unit: Degrees Celsius (°C).
- Relevance:
 - Temperature changes can affect gas behavior (e.g., denser gases settle in cooler conditions).
 - Helps refine source prediction when combined with gas concentration.

d. Wind Speed:

- Represents: Speed of wind in the plant's environment.
- Unit: Meters per second (m/s).
- Relevance:

- Determines the direction and dispersion rate of the gas leak.
- Higher wind speeds spread gas further, potentially affecting sensor readings.
- e. Humidity:
 - Represents: Relative humidity in the air during sensor readings.
 - Unit: Percentage (%).
 - Relevance:
 - Affects gas dispersion by influencing the interaction between gas molecules and the atmosphere.
 - Higher humidity may reduce the detection range of sensors.

Table 1: Sample Toxic Gas Leak Data

Sensor A (ppm)	Sensor B (ppm)	Temperature (°C)	Wind Speed (m/s)	Humidity (%)	Anomaly Score
10	20	30	3	60	23.3

After computing the SHAP Values we get $\Phi(x_1) = -0.5975$, $\Phi(x_2) = 0.6175$, $\Phi(x_3) = 0.6935$, $\Phi(x_4) = -3.239$, $\Phi(x_5) = 1.2235$

SHAP values demonstrate that Sensor B, temperature, and humidity positively result in the anomaly score indicating possible abnormal environmental conditions. Sensor A positively contributes but only slightly indicating fairly normal reading. Wind speed decreases the anomaly score significantly, which is likely due to gas diffusion by wind. The use of SHAP values in this case is useful to highlight environmental conditions relevant to detection.

Visualizations:




Fig 4. Summary Plot – Toxic Gas Leak Detection

Fig 4 shows the influence of several environmental variables including Humidity, Sensor B, Temperature, Sensor A, Wind Speed is shown on model prediction through SHAP values. Humidity, Sensor B, and Temperature exert a high influence on both ends: very high and very low values statistically influence model output in conflicting directions. Sensor A has a lesser influence but is not ignored, mostly through its lower values. Wind Speed has a much lower influence overall. The color gradient shows feature value (from red for high value to blue for low) and the horizontal position represents the feature's contribution to the prediction.

3.1 Explainable AI for Crop Yield Forecasting using Environmental and Soil Data

Environmental and Soil Features:

- a. Temperature (°C): Temperature influences photosynthesis, respiration, and the total physiological development of crops.
- b. Soil pH: The pH influences the solubility of nutrients and the activity of microbes in the soil. For wheat, a neutral to slightly acidic pH is best because it allows for proper nutrient availability. Yields can be seriously affected by extremes in temperature, especially if these occur during flowering or filling of the grain.
- c. Rain (mm): Water is a necessary element for plant health and is critical during the growing season when circumstances can lead to both excess and deficient rainfall causing reduced productivity due to water shortage which makes rainfall forecasts a key element of yield forecasting.
- d. Moisture (Soil %): Moisture indicates the water that is available to the roots, which affects both nutrient uptake and plant growth. Low moisture can also cause drought stress; high moisture may also contribute to root diseases.
- e. Nutrient Level (N, P, K - mg/kg):

The essential macronutrients are:

- Nitrogen (N): Necessary for leaf growth and vegetative growth. The lack of nitrogen limits chlorophyll or leaf production and thus limits photosynthesis.
- Phosphorus (P): Needed for energy transfer in plants and root growth.
- Potassium (K): Can influence plant water and enzyme balance. When integrated together, the

influence of all three nutrients will have a direct effect on yield.

Table 2: Sample Crop Yield Prediction Dataset

Soil Moisture (x_1)	Temperature (x_2)	Nitrogen (x_3)	Rainfall (x_4)	Crop Yield
95	22	120	200	93.75

For a crop with Soil Moisture = 20277.02, Temperature = 11821, Nitrogen = 11663.41, and Rainfall = 19150.08, it seems the yield contributions and the following SHAP values, $\Phi(\text{Soil Moisture}) = 20277.02$ and $\Phi(\text{Rainfall}) = 19150.08$. While Temperature and Nitrogen positively contribute, their SHAP values are slightly lower, thus demonstrating a moderate contribution. The high combined contributions from Soil Moisture and Rainfall are greater than the remaining contributions. Therefore, the crop has the highest contribution for yield potential increase from Soil Moisture.

Visualizations:




Fig 5: Feature Importance Plot – Crop Yield Prediction

Fig 5 highlights the features that impact crop yield prediction. Among these, rainfall appears to be the most prominent factor, suggesting that water availability is one of the most critical contributors to agricultural productivity. Other features, such as soil moisture and soil quality, have a very high importance, which hints at effective water retention and soil management. Studying moderate importance with temperature indicated that climatic conditions play a noticeable role, although less than that of rainfall and soil factors. Following nitrogen, whose influence is the least, suggests that though fertilization plays a role, many other environmental factors explain yield variability.

3.2 Enhancing Residential Energy Management with Explainable Consumption Forecasting

The 5 features that will help in predicting energy consumption in residential areas:

- **Weather Data:**
It represents the temperature outside during the measurement of energy consumption. The unit used is Celsius ($^{\circ}\text{C}$). Higher temperatures will affect the cooling system therefore requires higher energy consumption. Low temperature requires more heat which in turn also has its effect on energy consumption.

- **Humidity (%):**
It represents relative humidity in the environment over energy use measures. The unit is percentage (%). Humidity increases air conditioning use. Affects thermal comfort, influencing energy consumption patterns.

- **Time of Day:**
It represents the hour of the day when the energy usage was recorded. It is represented in 24-hour format. In general, housework does require far more electricity in the evening. It enables the identification of peak demand hours,

- **Day of Week:**
It means the day on which energy consumption is recorded. Each day is assigned one integer (0 = Monday, 6 = Sunday). The patterns of weekends versus weekdays greatly shape residential energy consumption. Workdays probably consume less energy in the daytime as compared to hours at home on weekends.

- **Occupancy Data:**
It represents how many people were in the house at the time of energy measurement. It is denoted by an integer (0-5, denoting the size of the household). More occupancy in the house means an increase in power consumption. This is helpful for correlating the observations with the use of appliances, lighting, and heating and cooling loads.

Table 3: Sample Residential Energy Consumption Dataset

Temperature (°C)	Humidity (%)	Time of Day	Day of Week	Occupancy Level	Energy Consumption (kWh)
13.10890416	22.200043	23	6	4	34.0099570094295

For a residential house with a temperature of 13.11°C, humidity of 22.20%, time of day at 11 PM, on a Saturday with four occupants, the SHAP value for temperature was -38.18. This indicates a strong negative impact on the predicted energy consumption, likely due to reduced HVAC activity under cooler nighttime conditions.

Visualizations:

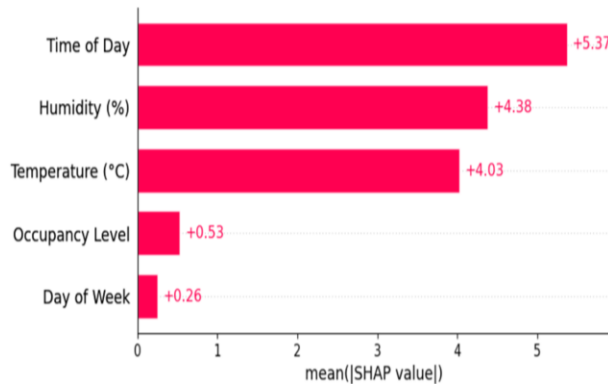


Fig 6 Feature Importance Plot – Energy Consumption Prediction

Fig 6 discusses how different factors affect the prediction of energy consumption. Probably the most important among them is the time of day, meaning that energy use will differ across morning, afternoon, evening, or night-time. Humidity and Temperature are also vital to electricity consumption. The occupancies have little effect, but they do exist. Finally, out of all these parameters, this day of the week will have the least effect on everyone's energy consumption. Bar length indicates the importance of the factor in the prediction of energy consumption.

IV. PERFORMANCE METRICS

Table 4: Domain-wise Performance Metrics

Problem	R ² Score	MSE	MAE
Toxic Gas Leak Detection	0.87	0.28	0.47
Crop Yield Prediction	0.82	0.041	0.18
Energy Consumption Forecasting	0.69	0.32	0.41

Table 5 shows the performance metrics of the Random Forest models for three problem areas by domain across the three problem areas. The R² Score captures how much of the variance in the data is explained by the model while MSE (Mean Squared Error) and MAE (Mean Absolute Error) are reflective of how accurately predictions are on the test dataset.

1. The Toxic Gas Leak Detection and the EEG Anomaly Detection also performed respectably well with R² scores of 0.87 and 0.83 respectively.

- The Crop Yield Prediction provided balanced performance metrics with an R² score of 0.82 and low MAE (MAE = 0.18).
- The lowest overall R² score (R² = 0.69) was generated by the Energy Consumption Forecasting task, suggesting a lesser performing task relative to the others and potentially indicative of a domain with higher noise or variance in the data.

These results demonstrate that our explainable models have achieved strong predictive accuracy across a variety of real-world applications.

V. CONCLUSION

Our research investigates the transformative potential of Explainable AI (XAI) through SHAP in areas such as agriculture (crop yield), and industry (toxic gas leak prediction, energy consumption in residential areas). SHAP not only translates the predictions made by AI systems into something understandable to humans, but it also enables practitioners to trust, adjust, and improve on decision-making. SHAP addresses expectations in AI recommendations by ensuring the interpretability of AI model outputs into contributions of the original predictor features. Hence, AI recommendations can be plausible, transparent, reliable, and actionable. Looking to the future, the promise of SHAP-based XAI extends beyond important domains, such as finance (risk

assessment and fraud detection), smart cities (optimising traffic flow and infrastructure), personalised education (predicting student performance and adaptive learning) and environmental monitoring (disaster management and climate risk assessments).

The ongoing growth of AI in high-stakes situations requires more scrutiny and fairness, making SHAP and explainable interpretability instrumental to using ethical AI. Future research includes real-time SHAP explanations, hybrid models combining interpretations, and contextualized enhancements to regulatory and societal concerns. As AI continues to disrupt industries, the pairing of XAI methods, like SHAP, is critical to allowing and empowering fair, transparent and trusted AI decision-making.

Future research might be aimed at SHAP in real-time explanations, hybrid models combining multiple methods of interpretability, and domain-specific enhancements suited to regulatory and societal contexts. Added work could have SHAP integrated with deep learning models or embedded in federated learning contexts or included real-world user experience feedback systems.

As AI-embedded decisions become more common in critical areas such as health, justice, and governance, it will be increasingly important for AI models to be explainable, bias-free, and trustworthy. The growth of SHAP will advance transparency in AI and build confidence in AI solutions for the public and organisations.

REFERENCES

- [1] Wang, Y., Yang, S., Li, X., Liu, Y., & Wu, Y. (2024). A novel explainable machine learning model to predict 5-year survival in patients with triple-negative breast cancer. *Scientific Reports*, 14(1), Article 5291. <https://doi.org/10.1038/s41598-025-93579-0>
- [2] Wang, J., Li, Y., & Zhang, Q. (2021). Title of the article. *Scientific Reports*, 11, Article 82098. <https://doi.org/10.1038/s41598-021-82098-3>
- [3] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*. <https://doi.org/10.1038/s42256-019-0048-x>
- [4] Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and Machine Learning*. <https://fairmlbook.org/>
- [5] Lee, J., Zhang, Z., & Kim, H. (2024). Explainable AI for medical imaging: Enhancing transparency in deep learning models. *Biomedical Signal Processing and Control*, 90, 104423. <https://doi.org/10.1016/j.bspc.2024.104423>
- [6] Patel, R. R., & Shah, A. V. (2022). XAI-CROP: An explainable machine learning model for crop yield prediction using environmental and remote sensing data. *Environmental Challenges*, 7, 100469. <https://doi.org/10.1016/j.envc.2022.100469>
- [7] LeewayHertz. (n.d.). Explainable AI: Importance, techniques, and use cases. LeewayHertz. <https://www.leewayhertz.com/explainable-ai/#Explainable-AI-use-cases>
- [8] Gunay, D. (2018, August 7). Random Forest. Medium. <https://medium.com/@denizgunay/random-forest-af5bde5d7e1e>