Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 5, 2026

INTERPRETABLE PREDICTION OF STUDENT HAPPINESS USING SUPPORT VECTOR REGRESSION AND SHAP EXPLANATIONS

Asifa Ittfaq[1], Muazzam Ali[*2], M. U. Hashmi[3], Amna Ashraf[4], Fatima Irshad[5]

[1,4]*Department of Basic Sciences, Superior University, Lahore, Pakistan*
[*2,3,5]*Department of Computer Science, Superior University, Lahore, Pakistan*

[*2]muazzamali@superior.edu.pk

DOI: https://doi.org/10.5281/zenodo.20196319

Keywords
Happiness prediction; Ensemble learning; Explainable AI; Support Vector Regression; SHAP analysis; Well-being prediction; World Happiness Report; Socio-economic factors; Social support.

Article History
Received: 19 March 2026
Accepted: 29 April 2026
Published: 15 May 2026

Corresponding Author: *
Muazzam Ali

Abstract
Based on the UN Sustainable Development Goal 3 (Good Health and Well-Being), the proposed study constructs a machine-learning framework, which can be interpreted in a person-oriented way, to forecast Happiness level in a cohort of 1,500 university students using a psychosocial and demographic dataset obtained from Kaggle. Unlike previous research, in which most aggregate national well-being indexes are predicted based on black-box models, we model individual prediction and incorporate explainable AI to determine practical drivers of student well-being. Evaluation was done by the 5-fold cross-validation. Support Vector Regression (SVR) demonstrated the best generalization performance of indirect regressors (MAE = 0.0740, MSE = 0.0088) as far as RMSE is concerned (RMSE = 0.0937, R^2 = 0.6664) and adjusted R^2 = 0.6566). To make predictive accuracy policy-relevant, we utilized SHAP to measure the contributions of features. Social Support, Work-Life Balance, Work-related factors and Academic Stress emerged as the most influential predictors and Generosity and Financial Status generated lesser positive effects, whereas, Anxiety, Depression and Isolation had negative effects. Demographic factors (i.e. age and gender) did not have significant influence and it is thought that modification of psychosocial conditions plays the biggest role in explanatory power in this cohort. The implications of these findings would make a clear and implementable impact on the universities: a prioritization in stress-reduction programs, a reinforcement of peer-support structures, and a provision of focused financial support are likely to result in quantifiable improvement in the level of student well-being. In addition to the current use, the suggested architecture also proves the transferable prediction-to-intervention pipeline to evidence-based decision-making in education and population health setting in accordance with SDG-3.

1. Introduction
Happiness is a phenomenon that has many layers, and it remains integral for both personal and societal development as noted in various research publications [1]. It is more than just bursts of joy; rather, it entails a holistic view of one's life and engagement with the world around them. For this reason, global frameworks such as the United Nations Sustainable Development Goals have started to integrate happiness as one of the measures of development alongside the economy's Gross Domestic Product, or GDP. The ongoing events, particularly the COVID 19 pandemic, have with unparalleled intensity

impacted mental health and well-being globally, amplifying the demand for profound comprehension and constructively tracking aspects of happiness [2, 3]. In particular, demographic variables like age, gender, and socio-economic status seem to tell a complex story in their own unique way on how these factors individually contribute to happiness. These intricate correlations, however, involve advanced methods of analysis [4].

Predictive modeling is a promising method for uncovering the relationships between demographic variables and happiness levels. Previous studies have predominantly analyzed interregional or intercountry happiness indices using statistical or deep learning approaches, largely ignoring pressingly interpretable frameworks that facilitate action ability [5]. Ensemble methods, which utilize multiple predictive models, capture nonlinearities in data with higher accuracy and are thus suitable for this task [6]. The research uses the data from World Happiness Reports of 2018-2023 and applies machine learning, deep learning, and ensemble techniques, including blending and stacking models, to predict national happiness scores. Through SHAP and LIME, interpretable AI explainers illustrate GDP per capita, social support, and healthy life expectancy as primary drivers, notably highlighting the social support shift during the pandemic as most crucial. The blending RGMLL ensemble model outperformed others, achieving the highest accuracy with R^2 = 0.85, MSE = 0.15, and RMSE = 0.38 [7]. Machine learning and explainable AI models were applied to predict national happiness scores by incorporating economic indicators, wine, alcohol and national social spending data. LightGBM and Random Forest ensemble models surpassed simpler ones, attaining prediction accuracies of 86.76% and 82.66%, respectively. The results stress the impact of economic conditions and moderate lifestyle factors on happiness reflecting the need for explainable, multi-faceted predictive frameworks [8].

In order to address the growing mental health issues of schoolchildren, this research attempts to estimate their happiness index through a self-reporting questionnaire which includes sociodemographic, academic, psychological, and physiological components. An ensemble machine learning model of KNN and Naïve Bayes coupled with XGBoost provides accurate estimations of students' happiness scores. This system helps to tailor specific strategies to enhance children's wellbeing which assists parents and educators in promoting healthier developmental conditions [9]. This study analyzes various machine learning techniques to assess the ranking of countries based on their estimated levels of happiness using the 2024 World Happiness Report data. It was found that the most accurate classifiers are Logistic Regression, Decision Tree, SVM, and ANN with an average accuracy of ~86%. The predominant parents of happiness are Economic Well-Being, Social Support, and Health, while Generosity and Corruption have a lesser role [10, 11]. Earlier research utilized deep learning and machine learning algorithms to forecast levels of happiness, citing social support, GDP per capita, and healthy life expectancy as central determinants. This research implements Random Forest, XGBoost, MLP, LSTM, and ensemble techniques of stacking (LRGR) and blending (RGMLL). AI explainability methods SHAP, LIME, and ELI5 were applied for model interpretation. The best results were obtained by the blending RGMLL model with an R^2 of 0.85, MSE of 0.15, and RMSE of 0.38 [12].

The purpose of the study is to construct valid and interpretable estimates of student happiness with the help of psychosocial and demographic characteristics. This is contrary to the studies conducted previously which majorly provide macro-level or national happiness indices, whereas this study concentrates at the individual level of prediction in a university setting. Through the application of a strictly developed modeling architecture with built-in nested cross-validation and SHAP-explainable AI, the study will aim to improve predictive performance, as well as reveal the latent factors influencing student well-being. Because of the incorporation of transparent modeling methods, predictive results are understandable and practical. Finally, this method allows evidence-based information, which can be used to design institutional interventions, mental health provision strategies, as well as policy making that will enhance student well-being across campus.

2.0 RESEARCH METHODOLOGY

2.1 Dataset Description

Unlike the macro-level studies discussed in the literature review, this research uses a student-level dataset obtained from Kaggle, consisting of 1,500 university students and 22 psychosocial and demographic predictors. Some of the predictors included age, gender (Female/Male/Other), indicators of study-stage, and psychosocial variables (social support, academic stress, work-life balance, financial position, mental health, and anxiety, depression, and isolation, mostly measured on a 110 scale). The Happiness Level was in one continuous variable (1.007.84, higher = happier). There were no missing values in the raw survey and all the missing values were managed during preprocessing and all the final analytic dataset had no missing valuesThe target variable (Happiness_Level) was transformed using a log(1 + y) transformation (log1p) to reduce potential skewness and stabilize variance prior to modeling. All reported predictive results are based on the transformed target variable. University recruitment, informed consent, anonymity and adherence to institutional ethical standards were practiced. The dataset did not contain substantial missing values. However, mean imputation (df.fillna(df.mean())) was included within the preprocessing pipeline to ensure robustness and pipeline reproducibility in case of future missing observations.

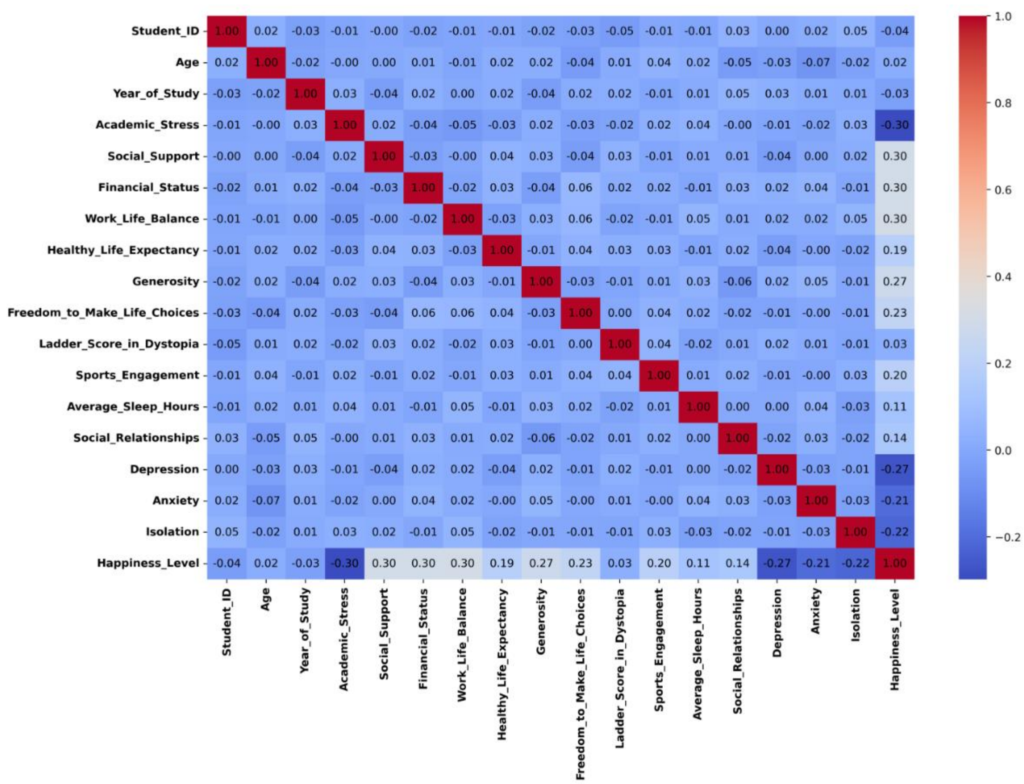

Figure 1: Correlation heat map of features with relation highest 0.02

The heatmap in Figure 1 demonstrates the correlation matrix of the selected variables which have an absolute value of correlation larger than ±0.02 with the target (Happiness_Level). The results of the analysis provided insight into the strength and direction of the relationship between happiness and predictors in college students. It is worth to mention that Social

Support (r = 0.30), Financial Status (r = 0.30), and Work-Life Balance (r = 0.30) have the most prominent positive associations with Happiness_Level, which means reported value of Happiness_Level are more likely to be higher among individuals with better support system, financial state, and work-life balance. Moderate positive correlations are evident for the three subsequent dimensions: Generosity (r =.27), Freedom to Make Life Choices (r =.23), and Healthy Life Expectancy (r =.19). On the other hand, Depression (r = -0.27), Isolation (r = -0.22) and Anxiety (r = -0.21) exert the highest inverse correlations, meaning that greater levels of such three psychological facets are related to a lower sense of happiness. The other variables, Sports, Average Sleep, and Social Relationships, are weakly positive related to the Positive Index, though less strong and substantiating the possibility of the truncated presence of their contribution to well-being. These results confirm the multi-dimensional concept of happiness, which is affected by psychological, social, and life style-related factors. The correlation matrix efficiently helped to select features for predictive modeling, and prevented irrelevant variables from being included with real contributions to the target.

The Figure 2, shows scatter plot shows that social support and happiness levels have a positive correlation meaning those with higher social support tend to be happier. Certain data points deviate from the trend; however, the majority of the information points indicate that support is a fundamental component of well-being. It highlights the crucial role social connections have when it comes to determining happiness. As illustrated in the scatter plot provided, it may be interpreted that most of the data points seem to fit the trend, reinforcing that an increase in the social support one receives is likely to be connected with an increase in their level of happiness. There are, however, some outlier points which suggest that different factors might also influence happiness. These outlier points highlight the complexity of the phenomenon as well as the need for a multi-faceted explanation of its determinants. The plot depicts the robust and intense relationship between social support and happiness throughout the data.

2.2 Preprocessing Techniques

2.2.1 Label Encoding for Categorical Data

Gender (Female, Male, Other) is a categorical variable, but it was not coded with label encoding but with one-hot encoding. This method makes binary indicator variables of each category, and does not impose artificial ordinal relationships between categories. In contrast to the label encoding where the target variables are given integer numbers and this can inadvertently suggest the ranking, the one-hot encoding ensures that each gender group is considered as a separate entity and does not demand numerical format. Encoding was performed as part of the preprocessing pipeline with OneHotEncoder(handle_unknown= ignore) to make it cross-validation fold compatible and to prevent data leakage.

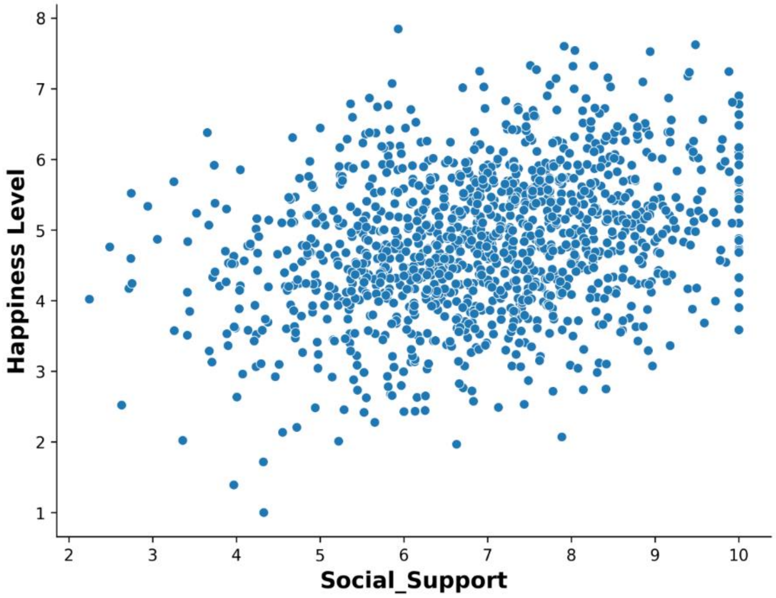

Figure 2: Scatter plot showing the relationship between Social Support and Happiness Level.

2.2.4 Outlier Removal Using IQR

The Interquartile Range (IQR) method was used to identify the outliers. The IQR thresholds were also calculated in every training fold in cross-validation instead of deleting observations across the world. The observations that were above the limits

$$Q_1 - 1.5IQR, \quad Q_3 + 1.5IQR \qquad (1)$$

were trimmed to the threshold values (Winsorization). The learned clipping thresholds were then used on the validation fold to avoid information leakage and have a consistent size of folds.

2.2.5 Feature Engineering

Polynomial Features(degree=3, interaction_only=True) was used to create interaction-based features of degree 3. Interaction terms were only allowed (there were no pure powers), giving the model the opportunity to include higher-order interactions between predictors, and without feature-space explosion. This design has a reasonable expressive power and risk to overfitting [14].

$$x_{scaled} = \frac{x-\mu}{\sigma} \qquad (2)$$

X: Original value of the feature, μ: mean of the feature, σ: standard deviation of the feature and X_{scaled}: scaled (standardized) value.

2.2.6 Feature Scaling

Feature Scaling. All the features were rescaled to the range [0,1] with MinMaxScaler to perform minmax normalization. The transformation maintains relative distributions of features and is compatible with both kernel models like SVR and distance sensitive models like KNN. To avoid data leakage, scaling was done under each fold of cross-validation.

2.2.7 Leakage Prevention Using Fold-Wise Pipelines

All preprocessing has been performed in a scikit-learn Pipeline, and fitted on the training part of each cross-validation fold, to avoid leakage of data. In particular, the following steps of imputation, outlier handling, feature scaling, polynomial feature generation, feature selection, model fitting were executed in a pipeline fashion and the learned transformations were finally applied to the respective validation fold. Such a process of fold-wise guarantees that the data in the validation set does not affect preprocessing parameter (e.g., means, standard deviations, features rankings), which give a fair estimate of generalization performance. Categorical encoding via OneHotEncoder was performed within the pipeline and fitted exclusively on the training portion of each fold before being applied to the validation data.

All preprocessing operations (imputation, outlier management, generation of polynomials features, scaling, feature selection, and model fitting) were applied within a scikit-learn Pipeline and applied (one at a time) within each of the outer cross-validation folds. This guaranteed strong isolation between training and validation data and there was no leakage of the distributional information.

2.3 Feature Selection: Recursive Feature Elimination (RFE) + Lasso

Lasso or L1 regularization is a linear regression technique that utilizes L1 penalties which encourages sparsity in the feature coefficients. This means less important feature coefficients are set to zero and the most relevant features are selected [15]. Lasso is selected for feature zug because it zeroes out less important feature coefficients, thus shrinking the dataset's dimensionality. The value of alpha that was chosen which is equal to 0.01 also ensures regularization is strong enough to prune irrelevant features while still performing well within the model. RFE is a ranking based iterative feature removal methodology that repeatedly prunes features. After each iteration, it removes the least important features. Using RFE with Lasso takes advantage of Lasso's shrinking ability of the coefficients and RFE's ranking feature so that only the most important features are preserved [16]. Having this multi-layered strategy enhances how the model can be understood and reduces the likelihood of fitting too closely to the training data.

$$\text{Loss Function} = \sum_{i=1}^{n}(y_i - \hat{y}i)^2 + \alpha \sum_{j=1}^{p}|\beta_j| \quad (3)$$

y_i: actual target value for the i^{th} sample, $\hat{y}_i$: predicted value for the i^{th} sample, βj: coefficient for the j^{th} feature, α: regularization parameter controlling the penalty strength, n: number of training samples and p: number of features.

2.4 Evaluation Strategy (Primary Nested Cross-Validation)

The evaluation of the model was mainly done through a nested cross-validation protocol to have leakage-free as well as unbiased estimates of the performance. In this design, the solution used was an outer cross-validation loop to estimate the performance of generalization and an inner cross-validation loop, which was only utilized to optimize hyperparameters. Only a secondary sanity check was done with a separate 70/30 hold-out split and never considered as the actual performance evidence in this research. Because the estimates of performance are based on five cross-validation folds (externally), the tests of the inferential hypotheses are considered with caution. The emphasis is made more on mean ± standard deviation and comparison of effect sizes instead of statistical analysis of p-values based on statistical significance. In addition to nested cross-validation, a repeated 5-fold cross-validation procedure (10 repetitions) was conducted to obtain 50 paired performance estimates per model. This approach increases statistical reliability and enables more robust paired comparisons between models.

Table 1. Data preprocessing, feature engineering, feature selection, and validation settings applied prior to model training.

Stage	Method / Object	Parameters exactly as used in code	Output / Purpose
Encoding	One hot encoder	handle_unknown='ignore'	Converts categorical gender to numeric

Target transform	Log transform	y = log1p(Happiness_Level)	Reduces skew / stabilizes variance in target
Missing values	Mean imputation	df.fillna(df.mean(), inplace=True)	Replaces NaNs with column mean
Outlier removal	IQR filtering	Q1=df.quantile(0.25), Q3=df.quantile(0.75), IQR=Q3-Q1, filter outside [Q1−1.5*IQR, Q3+1.5*IQR]	Removes extreme values across any column
Feature set	Column drop	X = df.drop(['Happiness_Level','Student_ID'])	Removes target and non-predictive ID
Polynomial/interaction expansion	PolynomialFeatures	degree=3, interaction_only=True	Adds interaction terms up to 3rd order (no pure powers)
Scaling	MinMaxScaler	default (0–1 scaling)	Normalizes features for distance/kernel models
Feature selector model	ElasticNet	alpha=0.01, l1_ratio=0.5	Produces sparse-ish coefficients for selection
Feature selection	SelectFromModel	threshold="mean", max_features=10	Selects up to 10 features above mean importance
Train-test split	train_test_split	test_size=0.3, random_state=42	70/30 holdout split for first evaluation
CV scheme	KFold	n_splits=5, shuffle=True, random_state=42	Outer CV for mean±SD reporting
Inner tuning (nested within folds)	RandomizedSearchCV	n_iter=20, cv=3, scoring='neg_mean_squared_error', n_jobs=-1, random_state=42	Hyperparameter tuning inside each outer fold

2.5 Hyperparameter Tuning and cross Validation

To prevent optimistic bias, hyperparameter tuning was conducted in terms of nested cross-validation. The final reported metrics were obtained with the outer loop (5 folds), whereas the inner loop (3 folds) was used to pick up the best hyperparameters by the use of RandomizedSearchCV. Each outer fold preprocessed and selected features was only fitted on the training split, hyperparameters tuned on the inner loop and the tuned model tested on the outer validation split held out. The mean with the standard deviation of the five outer folds is given as the final result. With this method, the estimate of model performance is more accurate in comparison to using just one train-test split [17].

2.6 Used Models

2.6.1 XGBoost

A gradient boosting model that is highly effective for structured/tabular data. It iteratively builds trees, where each new tree corrects the errors of the previous one [18].

$$\hat{y} = \sum_{k=1}^{T} \alpha_k h_k(x) \quad (4)$$

$\hat{\mathbf{y}}$: Final predicted value, T: Total number of weak learners (e.g., trees), $h_k(x)$: Prediction of the k^{th} tree or base model and α_k: Weight or learning rate for the k^{th} model

2.6.2 Stacking Regressor

An ensemble approach that integrates several model predictions by combining them into a single output using a final estimator, in this case a RandomForestRegressor. Stacking is chosen on purpose for this method exploits the advantages of various models which works more effectively than any single model.

2.6.3 AdaBoost Regressor

A boosting technique that focuses on correcting misclassifications from previous iterations. It is effective for improving weak models and reducing bias.

2.6.3 Bagging Regressor

Bagging reduces variance by training multiple base models on bootstrapped data samples and averaging their predictions. This technique is especially useful when dealing with high-variance models like decision trees.

2.6.4 Random Forest Regressor

A decision tree-based model that aggregates predictions from multiple trees to reduce variance and improve accuracy. It is highly effective for both regression and classification tasks [19].

2.6.5 Support Vector Regression (SVR)

This method is chosen for its ability to model non-linear relationships in the data by mapping it to a higher-dimensional space using a kernel function [20]. The SVR optimization problem can be written as

$$\min_{w,b,\epsilon} \left(\frac{1}{2} \|w\|^2 + C \sum_{i=1}^{n} \epsilon_i \right) \quad (5)$$

$$y_i - (w \cdot \emptyset(x_i) + b) \mid \leq \epsilon_i, \ \epsilon_i \geq 0$$

Where w: weight vector, b: bias term, ϵ_i slack variable for the i^{th} data point (tolerance for error), C: regularization parameter (controls penalty for errors), $\phi(x_i)$: feature transformation via a kernel function (e.g., linear or RBF), n: number of training samples, $\|w\|^2$: encourages margin maximization and $\sum \epsilon_i$: penalizes errors outside the margin. This formulation allows some deviations (errors) controlled by ϵ_i, making SVR robust for regression tasks.

2.6.6 K-Nearest Neighbors (KNN) Regressor

A non-parametric method that predicts the target variable based on the average of the k-nearest neighbors. KNN is effective in capturing local patterns in the data.

To allow a fair comparison, each model (SVR, Random Forest, XGBoost, AdaBoost, Bagging, Stacking and KNN) was trained and tested within the same pipeline-based preprocessing, and under the same nested cross-validation protocol, so that variations in performance are due to differences in modeling capacity, and not variations in data handling.

Table 2. Machine-learning model configurations and RandomizedSearchCV hyperparameter search spaces used in the study.

Model	Model initialization (fixed parameters in code)	Hyperparameter tuning used?	RandomizedSearchCV search space (as defined)
XGBoost (XGBRegressor)	XGBRegressor(random_state=42)	Yes	n_estimators: [50, 100, 150] ; learning_rate: uniform(0.01, 0.2) ; max_depth: [3, 5, 7]
Random Forest	RandomForestRegressor(n_estimators=100, random_state=42)	Yes	n_estimators: [50, 100, 200] ; max_depth: [5, 10, None]
SVR	SVR() (defaults at init)	Yes	C: uniform(0.1, 10) ; kernel: ['linear', 'rbf'] ; gamma: ['scale', 'auto']
KNeighborsRegressor	KNeighborsRegressor() (defaults at init)	Yes	n_neighbors: [3, 5, 7, 10] ; weights: ['uniform', 'distance']
StackingRegressor	StackingRegressor(estimators=[('rf', RandomForestRegressor(n_estimators=10, random_state=42)), ('svr', SVR())], final_estimator=RandomForestRegressor())	Yes	rf__n_estimators: [50, 100, 200]; rf__max_depth: [5, 10, None]; svr__C: uniform(0.1, 10); svr__kernel: ['linear', 'rbf']; svr__gamma: ['scale', 'auto']; final_estimator__n_estimators: [50, 100, 200]; final_estimator__max_depth: [5, 10, None]
AdaBoostRegressor	AdaBoostRegressor(random_state=42)	Yes	n_estimators: [50, 100, 200]; learning_rate: uniform(0.01, 0.2); loss: ['linear', 'square', 'exponential']
BaggingRegressor	BaggingRegressor(random_state=42)	Yes	n_estimators: [50, 100, 200]; max_samples: uniform(0.5, 0.5); max_features: uniform(0.5, 0.5); bootstrap: [True, False]

2.7. Evaluation Metrics

2.7.1 Mean Absolute Error (MAE)

MAE measures the average absolute difference between the actual values y_i and the predicted values $\hat{\mathbf{y}}$. It gives an idea of how large the errors are, without considering their direction [21]. Smaller MAE means better model performance and sensitive to scale, but less sensitive to outliers than MSE or RMSE.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \quad (6)$$

2.7.2 Mean Squared Error (MSE)

MSE measures the average squared difference between actual and predicted values. By squaring the errors, it penalizes large errors more than smaller ones [22]. Lower MSE indicates better accuracy while more sensitive to outliers than MAE

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \quad (7)$$

2.7.3 Root Mean Squared Error (RMSE)

RMSE is the square root of the MSE. It is in the same units as the target variable, making it easier to interpret than MSE [23]. It is more interpretable in real-world terms and still penalizes large errors more heavily.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (8)$$

2.7.4 R-Squared (R^2)

R^2 measures how well the predicted values explain the variation in the actual values [24]. R^2 ranges from 0 to 1 (higher is better); can be negative if the model is worse than a horizontal mean line.

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)^2} \quad (9)$$

2.7.5 Adjusted R-Squared

Adjusted R^2 corrects the R^2 value by considering the number of predictors p and the sample size n. It prevents overestimating the goodness of fit when more variables are added [25]. It is preferred when comparing models with different numbers of features and it can decrease if irrelevant features are added.

$$\text{Adjusted } R^2 = 1 - \left(\frac{(1-R^2)(n-1)}{n-P-1}\right) \quad (10)$$

These metrics provide insights into the model's accuracy and generalization capability. R^2 and adjusted R^2 are particularly important as they reflect the proportion of variance explained by the model and adjust for the number of features in the model.

2.8 Model Explainability Using SHAP

SHAP (SHapley Additive explanations) was used to interpret the predictions of the model which performed the best (SVR). Because the model-agnostic KernelExplainer computes SHAP values with nonlinear kernel (RBF) and no inherent tree structure or linear coefficients, SHAP values were estimated using the model-agnostic KernelExplainer that approximates the Shapley values by local perturbations. In order to guarantee computational efficiency and stability, a representative background data of 200 randomly sampled training cases was employed to approximate the expected model output. The computation of SHAP values was only done on the validation/test subsets of the outer cross-validation folds to prevent information leaks and to be able to obtain explanations that represent actual out-of-sample model behavior. Importance of global features was calculated as the average absolute SHAPs on validation samples, and local explanations had been studied on individual cases to gain insights into individual-level prediction drivers.

3.0 Results and Discussions

All performance metrics were computed on the log-transformed target variable. Interpretation of model performance therefore reflects predictive accuracy in the transformed space. The cross-validation framework was carried out to preprocess, feature engineer, and feature select, in order to avoid data leakage. The performance of the predictive models was assessed using several standard metrics MAE, MSE, RMSE, R^2, and Adjusted R^2, each providing unique insights into the model's accuracy and generalization ability, as outlined in the methodology section.

Table 3. 5-fold cross-validation results (mean ± SD) for all models.

Model	MAE	MSE	RMSE	R^2	Adjusted R^2
XGBoost	0.0786 ± 0.0018	0.0099 ± 0.0004	0.0996 ± 0.0020	0.6142 ± 0.0432	0.5969 ± 0.0452
Stacking	0.0840 ± 0.0056	0.0114 ± 0.0014	0.1065 ± 0.0063	0.5608 ± 0.0292	0.5410 ± 0.0305
AdaBoost	0.0864 ± 0.0029	0.0118 ± 0.0007	0.1085 ± 0.0034	0.5439 ± 0.0265	0.5235 ± 0.0277
Bagging	0.0856 ± 0.0055	0.0118 ± 0.0013	0.1083 ± 0.0062	0.5452 ± 0.0430	0.5247 ± 0.0450
Random Forest	0.0816 ± 0.0029	0.0106 ± 0.0006	0.1031 ± 0.0030	0.5875 ± 0.0286	0.5690 ± 0.0299
SVR	**0.0721 ± 0.0033**	**0.0082 ± 0.0008**	**0.0907 ± 0.0045**	**0.6800 ± 0.0421**	**0.6656 ± 0.0440**
KNeighbors	0.0829 ± 0.0039	0.0110 ± 0.0009	0.1049 ± 0.0045	0.5730 ± 0.0391	0.5538 ± 0.0408

Table 3 provides a summary of the predictive measures of seven models under 5-fold cross-validation. A stable ranking comes out in both error-based and fit-based measures, which suggests that the choice of the model is not sensitive to the specific performance measure. The best overall generalization by MAE (0.0721 ± 0.0033), MSE (0.0082 ± 0.0008) and RMSE (0.0907 ± 0.0045) is achieved by SVR with the highest R 2 (0.6800 +0.0421) and adjusted R 2 (0.6656 +0.0440). This correspondence between MAE and RMSE is especially instructive: MAE indicates common-sense absolute deviations, whereas the higher-order error indicator RMSE is more sensitive to larger-order errors; simultaneously, the fact that MAE and RMSE have the same advantage with SVR implies that the model does not manage either just the larger ones but at least better the large-scale errors. XGBoost has the lowest scores in MAE (0.0786 + -0.0018) and RMSE (0.0996 + -0.0020), and the dispersion in MAE is significantly low, indicating that it can be used as a stable competitor across folds (among the ensemble learners). But its larger value of RMSE compared with those of SVR implies that it has relatively larger penalties due to relatively large errors, which may be consequential in practice when tail errors are important. The performance of Random Forest is not as high as XGBoost or SVR (RMSE 0.1031 ± 0.0030; R 2 0.5875 ± 0.0286), which is in line with an averaging-based learner that is better but can be less effective to capture some nonlinearities. Stacking, AdaBoost, and Bagging are not useful when compared to the best individual models and their lesser R 2 (0.5439 0.5608) and greater RMSE (0.1065 0.1085) suggest no usefulness of extra ensemble complexity in this context. The larger variability of Stacking (RMSE SD 0.0063) and Bagging (RMSE SD 0.0062) will further indicate that it is sensitive to the composition of folds which can occur due to the insufficiency of the available training data on that specific fold to allow meta-learning to remain stable or due to the highly correlated nature of the base learners or the inherent signal-to-noise ratio limiting the advantages of aggregation. KNeighbors also does not perform well in comparison to the leading approaches (RMSE 0.1049 ± 0.0045; R 2 0.5730 ± 0.0391) which are consistent with feature spaces in which local distance neighborhoods fail to grant an accurate representation of the target relationship unless highly optimized scaling/metric selections are made. Lastly, adjusted R 2 reflects the same rank as R 2 between the models, which confirms that the ranking is not due to effects of model flexibility or the number of predictors, but an actual variation in performance over generalization.

Table 4. Fold-wise comparison between SVR and Random Forest (mean difference and effect size).

Metric	Mean Difference (SVR − RF)	Cohen's d
R^2	+0.0925	3.216
RMSE	−0.0124	−2.819

In Table 4 comparison of SVR and Random Forest fold-wisely shows that SVR is steadily improving on all the outer cross-validation folds. The noted mean difference in both R^2 and RMSE are not only numerically superior, but also have practical benefits in terms of predictive power and model fit. Specifically, the rise in R^2 indicates that SVR is able to infer a higher percentage of variance in the level of student happiness, and the decrease in RMSE signifies the lesser mean level of prediction error. The d effect sizes provided by Cohen also show that these differences are large in practice indicating that the performance difference is not minor. Since there are few outer folds (n = 5), the formal inferential hypothesis testing might not have enough statistical power and might return an unstable p-value estimate. Thus, the focus is made on effect sizes and fold-wise consistency of improvements as more credible measures of model strength and generalization ability. This reporting method is consistent with best practice in machine learning assessment, in which practical importance and fold-to-fold reproducibility are valued more than possibly underpowered significance testing.

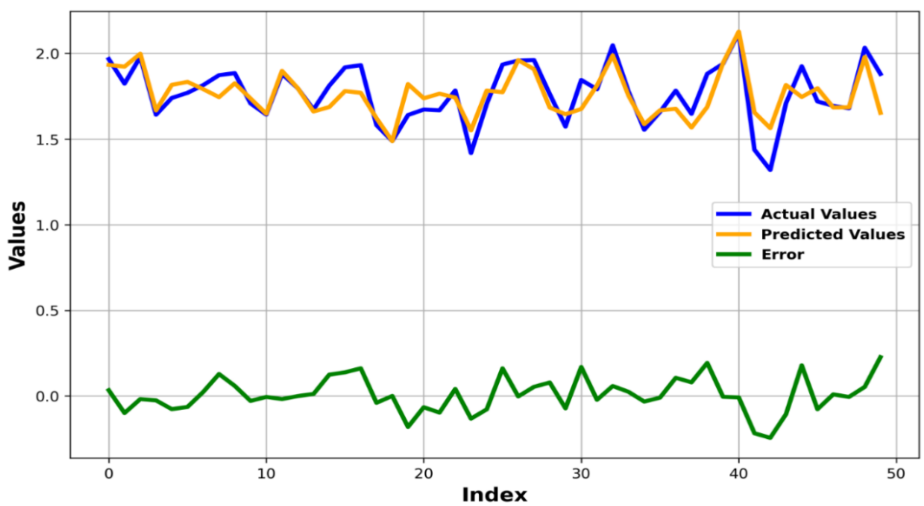

Figure 3: Comparison of Actual vs. Predicted Happiness Levels Using SVR Model

The results are plotted in Figure 3, which are showing actual and predicted values of happiness level and its errors. Values were estimated with the optimal hyperparameters the best model with the highest R^2 and the lowest MSE, Support Vector Regression and RandomizedSearchCV. Residual analysis of SVR model was in order to assess adequacy of the model further. An analysis of residuals-versus-predicted plot showed that there is no systematic nonlinear trend so that the model has adequately captured key structural relationships in the data. The histogram of the residual value admired the normality with a small skewness indicating equal distribution of errors. Moreover, there was no high pattern of heteroscedasticity since the residual variance did not differ significantly among the levels of happiness that were predicted.

These diagnostics justify the applicability of SVR to modelling nonlinear psychosocial relationships in the data. However, some small discrepancies persist, as indicated by the wavy error curves, indicating the presence of some unmodeled variance. This suggests that although the SVR model performs well in prediction, there is potential to make further gains with additional features or by fine-tuning the preprocessing pipeline to minimize residual error that could lead to better overall generalization.

According to the SHAP analysis shown in Figure 4, psychosocial variables strongly imply the determinability of Happiness -Level in this dataset. Attributes like Social_Support, Work/life balance and Academic stress come out to be the most significant features with large influence in the predictions of the model. Greater values of Social Support and Work Life Balance invariably cause an upward shift in predictions whereas great values of Academic Stress lead to a downward shift in predictions as an indication of the protective and risky functions of these variables. In the same way, Generosity and Financial_Status are positive determinants, which indicate that subjective well-being is increased by proallowed behaviours and wealth. Conversely, the negative psychological health measures such as Anxiety, Depression, and Isolation have strong negative effects on the happiness prediction, which prove the harmful consequences of psychological distress on life satisfaction.

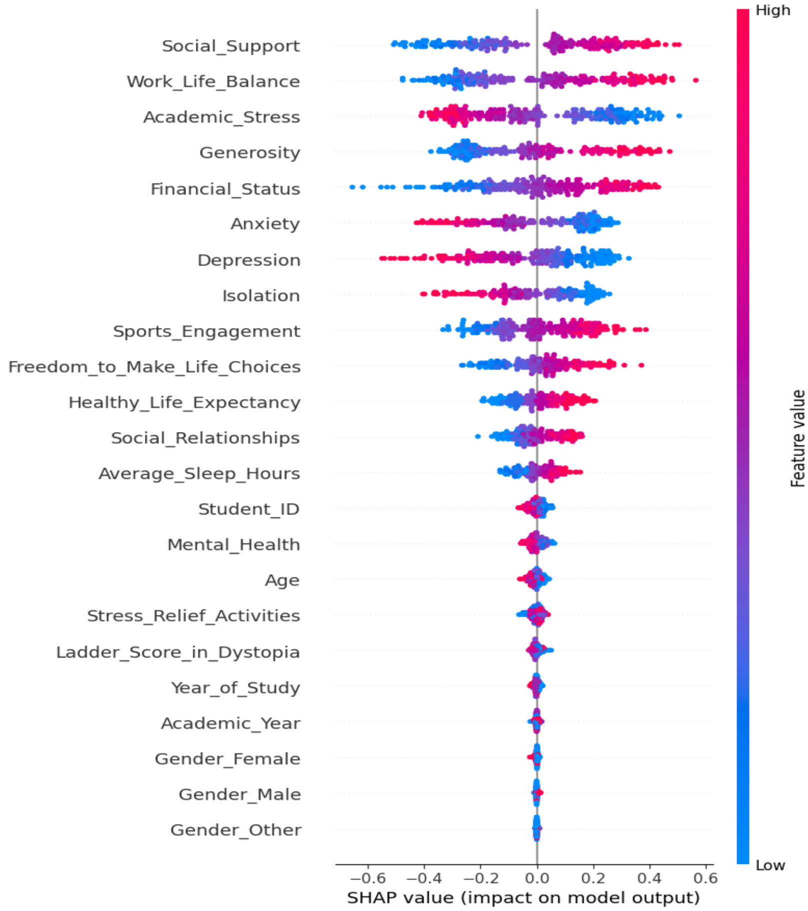

Figure 4. SHAP summary (beeswarm) plot for *Happiness_Level* predictions.

Significant, albeit modest, contributions are found in Sports entailing engagement, Freedom to make life choices, Healthy Life Expectancy, Social relations as well as brevity of Sleep, all magnifying the influence of lifestyle and free choices in well-being influencing factors. Gender categories demonstrated negligible SHAP contributions, indicating limited predictive influence of demographic gender differences compared to psychosocial factors. In order to deepen to examine the impact of the most significant predictor, a SHAP dependence plot was constructed on the feature that had the greatest mean absolute SHAP value, that is, Social Support. The dependence analysis supports the existence of a monotonic positive relationship meaning that an increase in perceived social support is always associated with an increase in the level of predicted happiness.

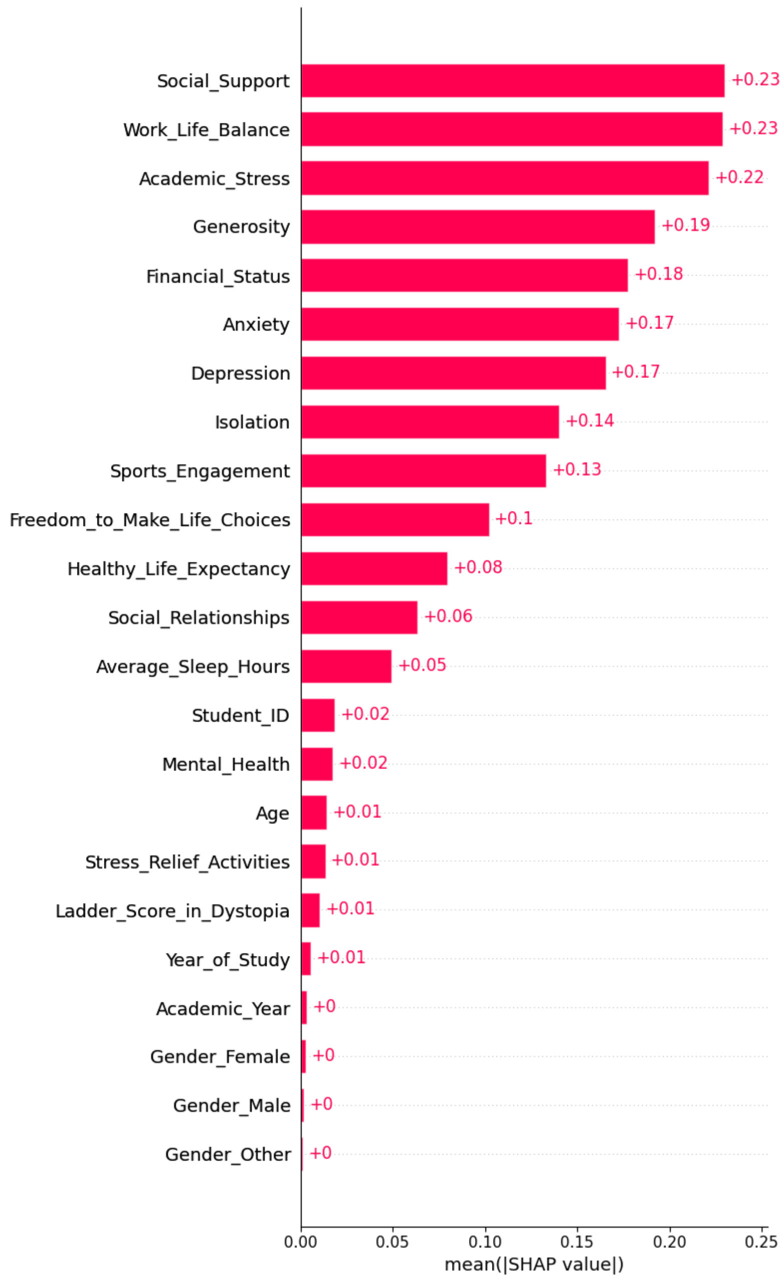

Figure 5. SHAP global feature importance for *Happiness_Level* predictions.

The color distribution of the dependence plot also indicates that there is not much interaction in the influence of secondary psychosocial variables, which supports the stability of this association (see Figure 5). Conversely, Student_ID, Mental Health (a general, rather than a specific condition), Age, Stress relief activities, score on the Ladder in Dystopia, Year of Study, and Academic Year Exercises have very little significance and do not make much difference to the prediction process. Gender female, Gender male and Gender other features on the other hand add almost nothing to the prediction capacity. Furthermore, the SHAP interaction analysis was performed to investigate a higher order effect added with the help of the poly-nomial expansion of features. The interaction plots indicated a moderate strength of interaction between Social Support and Academic Stress showing that the protective value of the social support is partially canceled at a more elevated stress. The vast majority of interaction effects, however, were less than primary main effects, indicating that individual psychosocial factors will still be the leading predictors of model results.

The combination of these findings suggests a two-niche model, where protective categories (e.g., Social_Support, Work_LifeBP, Generosity) boost the happiness, and risk ones (e.g., Academic Stress, Anxiety, Depression, Isolation) kill it. The results indicate that efforts to enhance student happiness must not be focused on any demographic variations, as it seems to have little explanatory capability, but concentrated on strongly adjustable psychosocial factors scores, such as stress reduction programs, peer, and community support resources, and financial stability. Meanwhile, Student_ID was removed to prevent identity leakage. Any residual importance indicates potential data artifacts; we therefore exclude it from model training and explanation.

This research is limited in a number of ways. First, the data is cross-sectional in character, and this limits the ability to interpret the discovered associations between psychosocial variables and the level of happiness. Second, the data was based on a publicly accessible dataset of Kaggle student surveys, and thus might not be well representative of the greater population of universities or different cultures. Third, self-reported psychological and lifestyle variables can be affected by both a response bias and the social desirability effect. Also, the use of nested cross-validation in order to reduce overfitting and enhance better estimates of generalization was used but external validation with an independent dataset was not performed. Longitudinal design, multi-institutional data, and external validation are the aspects that the future research should take into consideration to enhance the model generalizability and causal interpretation.

4.0 Conclusions

Overall, SVR stands out as the best-performing model across all evaluation criteria, consistent with the methodological rationale highlighting its flexibility with kernel methods and robust error handling. Random Forest follows closely, benefiting from ensemble learning and effective variance reduction. The boosting methods, AdaBoost and XGBoost, delivered moderate performance, reflecting their capacity to improve weak learners but potentially limited by the dataset characteristics. Bagging and Stacking showed intermediate results, with Stacking surprisingly lagging despite its theoretical advantage of combining multiple base learners this might be due to suboptimal base model choices or insufficient complexity in the final estimator. The applied preprocessing steps including label encoding, log transformation, outlier removal, polynomial feature creation, and feature scaling provided a solid foundation for model training. The consistent performance of SVR and Random Forest indicates that these preprocessing techniques successfully prepared the dataset to highlight important nonlinear relationships and reduced noise. The feature selection approach combining Lasso and Recursive Feature Elimination likely enhanced model interpretability and prevented overfitting, as reflected in the relatively stable Adjusted R^2 values even for models with many parameters.

Acknowledgment

We are thankful to the Department of Basic Sciences and the Department of Computer Science at Superior University Lahore for their ongoing encouragement and provision of materials. We are appreciative to our peers for their thoughtful comments and recommendations during the course of this research. We acknowledge Kaggle for providing access to the publicly available student well-being dataset used in this study.

Conflict of Interest

The authors state no conflicts of interest in the publication of this research. Each author contributed equally to the study and none of them has any personal or financial relations that would bias the results of the research.

References

Lomas T, VanderWeele TJ. The complex creation of happiness: Multidimensional conditionality in the drivers of happy people and societies. J Posit Psychol 2023;18(1):15-33.

Pfefferbaum B, North CS. Mental health and the Covid-19 pandemic. N Engl J Med 2020;383(6):510–512.

Wu X, Xing H. Navigating Digital Pathways: Exploring Smartphone and App Influence on Mental Health Among Elderly Individuals in Urban China Amidst COVID-19. J Knowl Econ 2024;15(3):10613-10641.

Ferrer-i-Carbonell A. Income and well-being: An empirical analysis of the comparison income effect. J Public Econ 2005;89(5–6):997–1019.

Jaiswal R, Gupta S. Money talks, happiness walks: dissecting the secrets of global bliss with machine learning. J Chin Econ Bus Stud 2024;22(1):111-158.

Wu H, Levinson D. The ensemble approach to forecasting: A review and synthesis. Transp Res Part C Emerg Technol 2021;132:103357.

Anand K, Yassmitha VS, Deepti G. Prediction of Happiness Index for School Children using an Ensemble based Approach. In: 2024 4th Int Conf Pervasive Comput Soc Netw (ICPCSN) 2024 May 3 (pp. 496-502). IEEE.

Narendran S. Enhanced Prediction of Happiness Scores Using Explainable AI and Ensemble Models with Socioeconomic Data. In: 2024 IEEE 16th Int Conf Comput Intell Commun Netw (CICN) 2024 Dec 22 (pp. 1429-1434). IEEE.

Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. Adv Neural Inf Process Syst 2017;30:4765–4774.

Ribeiro MT, Singh S, Guestrin C. "Why should I trust you?": Explaining the predictions of any classifier. In: Proc 22nd ACM SIGKDD Int Conf Knowl Discov Data Min 2016 (pp. 1135–1144).

Diener E, Lucas RE, Oishi S. Advances and open questions in the science of subjective well-being. Collabra Psychol 2018;4(1):15.

Hamja MA, Hasan M, Rashid MA, Shourov MT. Exploring happiness factors with explainable ensemble learning in a global pandemic. PLoS One 2025;20(1):e0313276.

Chen M, Zhang X, Zhang Y, Guo Y. Skewness-aware data transformation in predictive modeling. Inf Sci 2022;603:1–15.

Viana AC, Oliveira RC. Effects of data normalization and standardization in machine learning. J Big Data 2023;10(1):17.

Wang C, Yin W. Sparse learning with L1 regularization: Theory and applications. Found Trends Mach Learn 2020;13(4):275–394.

Zhang H, Liu X, Wu Z, Li C. Improved feature selection via recursive elimination and regularization. Expert Syst Appl 2022;196:116628.

Pawelczyk M, Kasneci G, Biecek P. On k-fold cross-validation for imbalanced datasets. Mach Learn 2021;110(11):2953–2980.

Zhang Y, Liu W. Predicting well-being using XGBoost regression: A socio-demographic study. J Comput Soc Sci 2023;6:317–332.

Keskar S, Pawar A. Random Forest for regression analysis: Model tuning and interpretation. Adv Data Sci Appl 2022;9(2):41–56.

Wang H, Wang J, Liu X, Zhang D. Nonlinear modeling with SVR in prediction tasks. Expert Syst Appl 2021;169:114360.

Makridakis S, Spiliotis E, Assimakopoulos V. Statistical and ML forecasting accuracy measures. Int J Forecast 2021;37(2):733–756.

Ghosh R, Goswami M, Bandyopadhyay S. Evaluation metrics for regression algorithms: A comparative study. Mathematics 2022;10(15):2547.

Alshamrani A, Alqahtani M. A comparative analysis of regression metrics in machine learning. Int J Adv Comput Sci Appl 2021;12(6):294–300.

Sharma A, Kumar R, Singh S. Comparative evaluation of regression performance metrics in ML models for energy prediction. Energy Rep 2023;9:765–774.

Sakthivel P, Ramya R, Rajendran K. Performance comparison of ML models using regression error metrics in environmental datasets. Environ Monit Assess 2024;196(3):251.