Predicting Rainfall using Machine Learning Techniques

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

11-04-2021 / 14-04-2021

CITATION

Oswal, Nikhil (2021): Predicting Rainfall using Machine Learning Techniques. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.14398304.v1>

DOI

[10.36227/techrxiv.14398304.v1](https://doi.org/10.36227/techrxiv.14398304.v1)

Predicting Rainfall using Machine Learning Techniques

Nikhil Oswal

School of Electrical Engineering and Computer Science (EECS),
University of Ottawa,
Ottawa, Canada
noswa023@uottawa.ca

Abstract. Rainfall prediction is one of the challenging and uncertain tasks which has a significant impact on human society. Timely and accurate predictions can help to proactively reduce human and financial loss. This study presents a set of experiments which involve the use of prevalent machine learning techniques to build models to predict whether it is going to rain tomorrow or not based on weather data for that particular day in major cities of Australia. This comparative study is conducted concentrating on three aspects: modeling inputs, modeling methods, and pre-processing techniques. The results provide a comparison of various evaluation metrics of these machine learning techniques and their reliability to predict the rainfall by analyzing the weather data.

Keywords: rainfall prediction, oversampling, undersampling, classifiers, model evaluation

1 Introduction

Rainfall prediction remains a serious concern and has attracted the attention of governments, industries, risk management entities, as well as the scientific community. Rainfall is a climatic factor that affects many human activities like agricultural production, construction, power generation, forestry and tourism, among others [1]. To this extent, rainfall prediction is essential since this variable is the one with the highest correlation with adverse natural events such as landslides, flooding, mass movements and avalanches. These incidents have affected society for years [2]. Therefore, having an appropriate approach for rainfall prediction makes it possible to take preventive and mitigation measures for these natural phenomena [3].

To solve this uncertainty, we used various machine learning techniques and models to make accurate and timely predictions. These paper aims to provide end to end machine learning life cycle right from Data preprocessing to implementing models to evaluating them. Data Preprocessing steps include imputing missing values, feature transformation, encoding categorical features, feature scaling and feature selection. We implemented models such as Logistic Regression, Decision Tree, K Nearest Neighbour, Rule-based and Ensembles. For evaluation purpose,

we used Accuracy, Precision, Recall, F-Score and Area Under Curve as evaluation metrics. For our experiments, we train our classifiers using Australian weather data gathered from various weather stations in Australia.

The paper is organized as follows. First, we describe the data set under consideration in Section 2. The adopted methods and techniques are presented in Section 3, while the experiments and results are shown and discussed in Section 4. Finally, closing conclusions are drawn (Section 5).

2 Case Study

In this paper, the data set under consideration contains daily weather observations from numerous Australian weather stations. The target variable is Rain-Tomorrow which means: Did it rain the next day? Yes or No. The dataset is available at <https://www.kaggle.com/jsphyg/weather-dataset-rattle-package> and definitions are adapted from <http://www.bom.gov.au/climate/dwo/IDCJDW0000.shtml>.

The data set consists of 23 features and 142k instances. Below are the features.

Table 1. Data set Description

Feature	Description
Date	The date of observation
Location	The common name of the location of the weather station
MinTemp	The minimum temperature in degrees celsius
MaxTemp	The maximum temperature in degrees celsius
Rainfall	The amount of rainfall recorded for the day in mm
Evaporation	The so-called Class A pan evaporation (mm) in the 24 hours to 9am
Sunshine	The number of hours of bright sunshine in the day.
WindGustDir	The direction of the strongest wind gust in the 24 hours to midnight
WindGustSpeed	The speed (km/h) of the strongest wind gust in the 24 hours to midnight
WindDir9am	Direction of the wind at 9am
WindDir3pm	Direction of the wind at 3pm
WindSpeed9am	Wind speed (km/hr) averaged over 10 minutes prior to 9am
WindSpeed3pm	Wind speed (km/hr) averaged over 10 minutes prior to 3pm
Humidity9am	Humidity (percent) at 9am
Humidity3pm	Humidity (percent) at 3pm
Pressure9am	Atmospheric pressure (hpa) reduced to mean sea level at 9am
Pressure3pm	Atmospheric pressure (hpa) reduced to mean sea level at 3pm
Cloud9am	Fraction of sky obscured by cloud at 9am.
Cloud3pm	Fraction of sky obscured by cloud at 3pm.
Temp9am	Temperature (degrees C) at 9am
Temp3pm	Temperature (degrees C) at 3pm
RainToday	1 if precipitation exceeds 1mm, otherwise 0
RISK_MM	The amount of next day rain in mm.
RainTomorrow	The target variable. Did it rain tomorrow?

3 Methodology

In this paper, the overall architecture include four major components: Data Exploration and Analysis, Data Pre-processing, Model Implementation, and Model Evaluation, as shown in Fig. 1.

Fig. 1. Overall Architecture.

3.1 Data Exploration and Analysis

Exploratory Data Analysis is valuable to machine learning problems since it allows to get closer to the certainty that the future results will be valid, correctly interpreted, and applicable to the desired business contexts [4]. Such level of certainty can be achieved only after raw data is validated and checked for anomalies, ensuring that the data set was collected without errors. EDA also helps to find insights that were not evident or worth investigating to business stakeholders and researchers.

We performed EDA using two methods - **Univariate Visualization** which provides summary statistics for each field in the raw data set (figure 2) and **Pair-wise Correlation Matrix** which is performed to understand interactions between different fields in the data set (figure 3).

Table 2. Irrelevant Features

Feature	% of Null values
Sunshine	43%
Evaporation	48%
Cloud3pm	40%
Cloud9am	38%

Fig. 2. Univariate Visualization.

We have other features with null values too which we will be imputing in our preprocessing steps. If we look the distribution of our target variable, it is clear that we have a class imbalance problem with number of positive instances - 110316 and number of negative instances - 31877.

Fig. 3. Heat Map.

The correlation matrix depicts that the features - MaxTemp, Pressure9am, Pressure3pm, Temp3pm and Temp9am are negatively correlated with target variable. Hence, we can drop this features in our feature selection step later.

3.2 Data Preprocessing

Data preprocessing is a data mining technique that involves transforming raw data into an understandable format. Real-world data is often incomplete, inconsistent, and/or lacking in certain behaviors or trends, and is likely to contain many errors. We have carried below preprocessing steps.

Missing Values: As per our EDA step, we learned that we have few instances with null values. Hence, this becomes one of the important step. To impute the missing values, we will group our instances based on the location and date and thereby replace the null values by there respective mean values.

Feature Expansion: Date feature can be expanded to Day, Month and Year and then these newly created features can be further used for other preprocessing steps.

Categorical Values: Categorical feature is one that has two or more categories, but there is no intrinsic ordering to the categories. We have a few categorical features - WindGustDir, WindDir9am, WindDir3pm with 16 unique values. Now it gets complicated for machines to understand texts and process them, rather than numbers, since the models are based on mathematical equations and calculations. Therefore, we have to encode the categorical data. We here tried two different techniques.

- **Dummy Variables:** A Dummy variable is an artificial variable created to represent an attribute with two or more distinct categories/levels [5]. However, as we have 16 unique values, our one feature will now get transformed to 16 new features which in turn results in **curse of dimensionality**. For each instance, we will have a feature with 1 value and rest 15 features with 0 values.

Example: Categorical Encoding of feature - **windDir3pm** using Dummy Variables

MinTemp	Rainfall	WindGustDir	WindGustSpeed	WindDir9am	WindDir3pm
13.4	0.6	W	44.0	W	WNW
7.4	0.0	WNW	44.0	NNW	WSW

Fig. 4. Sample Instance.

(WindDir3pm, S)	(WindDir3pm, SE)	(WindDir3pm, SSE)	(WindDir3pm, SSW)	(WindDir3pm, SW)	(WindDir3pm, W)	(WindDir3pm, WNW)	(WindDir3pm, WSW)
0	0	0	0	0	0	1	0
0	0	0	0	0	0	0	1

Fig. 5. Dummy Variables.

- **Feature Hashing:** Feature hashing scheme is another useful feature engineering scheme for dealing with large scale categorical features. In this scheme, a hash function is typically used with the number of encoded features pre-set (as a vector of pre-defined length) such that the hashed values of the features are used as indices in this pre-defined vector and values are updated accordingly [6].
Example: Categorical Encoding of feature - **windDir3pm** using Feature Hashing

MinTemp	Rainfall	WindGustDir	WindGustSpeed	WindDir9am	WindDir3pm
13.4	0.6	W	44.0	W	WNW
7.4	0.0	WNW	44.0	NNW	WSW

Fig. 6. Sample Instance.

WindDir9am_0	WindDir9am_1	WindDir9am_2	WindDir9am_3	WindDir3pm_0
-1.0	0.0	0.0	0.0	-2.0
-1.0	0.0	0.0	-2.0	-3.0

Fig. 7. Feature Hashing.

Feature Scaling: Our data set contains features with highly varying magnitudes and range. But since, most of the machine learning algorithms use Euclidean distance between two data points in their computations, this is a problem. The features with high magnitudes will weigh in a lot more in the distance calculations than features with low magnitudes. To suppress this effect, we need to bring all features to the same level of magnitudes. This can be achieved by scaling. We did this using scikit learn's min-max scalar and brought all the features in the range of 0 to 1 [7].

...	WindGustSpeed	WindSpeed9am	WindSpeed3pm	Humidity9am	Humidity3pm	RainToday	RainTomorrow
...	44.0	20.0	24.0	71.0	22.0	0.0	0
...	44.0	4.0	22.0	44.0	25.0	0.0	0

Fig. 8. Sample Instance before scaling.

.	WindGustSpeed	WindSpeed9am	WindSpeed3pm	Humidity9am	Humidity3pm	RainToday	RainTomorrow
...	0.289062	0.211765	0.258824	0.71	0.22	0.0	0.0
...	0.289062	0.023529	0.235294	0.44	0.25	0.0	0.0

Fig. 9. Sample Instance after scaling.

Feature Selection Feature Selection is the process where you automatically or manually select those features which contribute most to our prediction variable or output. Having irrelevant features in data can decrease the accuracy of the models and make the model learn based on irrelevant features. Feature selection helps to reduce over fitting, improves accuracy and reduces training time.

We used two techniques to perform this activity and got the same results.

- **Univariate Selection:** Statistical tests can be used to select those features that have the strongest relationship with the output variable. The scikit-learn library provides the SelectKBest class that can be used with a suite of different statistical tests to select a specific number of features. We used chi-squared statistical test for non-negative features to select 5 of the best features from our data set [8] [9].
- **Correlation Matrix with Heatmap:** Correlation states how the features are related to each other or the target variable. Correlation can be positive (increase in one value of feature increases the value of the target variable) or negative (increase in one value of feature decreases the value of the target variable). Heatmap makes it easy to identify which features are most related to the target variable, we plotted heatmap of correlated features using the seaborn library (figure 3) [9].

Handling Class Imbalance We learned in our EDA step that our data set is highly imbalanced. Imbalanced data results in biased results as our model doesn't learn much about the minority class. We performed two experiments one with oversampled data and another with undersampled data.

- **Undersampling:** We used Imblearn's random under sampler library to eliminate instances of majority class [10]. This elimination is based on distance so that there is minimum loss of information (figure 10)
- **Oversampling:** We used Imblearn's SMOTE technique to generate synthetic instances for minority class [10]. A subset of data is taken from the minority class as an example and then new synthetic similar instances are created. (figure 11)

3.3 Models

We chose different classifiers each belonging to different model family (such as Linear classifier, Tree-based, Distance-based, Rule-based and Ensemble). All the

Fig. 10. Undersampling.

Fig. 11. Oversampling.

classifiers were implemented using **scikit-learn** except for Decision table which was implemented using **weka**.

The following classification algorithms have been used to build prediction models to perform the experiments:

Logistic Regression is a classification algorithm used to predict a binary outcome (1 / 0, Yes / No, True / False) given a set of independent variables. To represent binary / categorical outcome, we use dummy variables. We can also think of logistic regression as a special case of linear regression when the outcome variable is categorical, where we are using log of odds as dependent variable. In simple words, it predicts the probability of occurrence of an event by fitting data to a logit function. Hence, this makes Logistic Regression a better fit as ours is a binary classification problem.

Decision Tree have a natural if then else construction that makes it fit easily into a programmatic structure. They also are well suited to categorization problems where attributes or features are systematically checked to determine a final category. It works for both categorical and continuous input and output variables. In this technique, we split the population or sample into two or more homogeneous sets (or sub-populations) based on most significant splitter / differentiator in input variables. This characteristics of Decision Tree makes it a good fit for our problem as our target variable is binary categorical variable.

K - Nearest Neighbour is a non-parametric and lazy learning algorithm. Non-parametric means there is no assumption for underlying data distribution. In other words, the model structure is determined from the dataset. Lazy algorithm means it does not need any training data points for model generation. All training data used in the testing phase. KNN performs better with a lower number of features than a large number of features. We can say that when the number of features increases than it requires more data. Increase in dimension also leads to the problem of overfitting. However, we have performed feature selection which helps to reduce dimension and hence KNN looks a good candidate for our problem.

Our Model's configuration: We tried various values of n ranging from 3 to 30 and learned that the model performs best with n as **25, 27** and **29**.

Decision table provides a handy and compact way to represent complex business logic. In a decision table, business logic is well divided into conditions, actions (decisions) and rules for representing the various components that form the business logic. [11] This was implemented using Weka.

Random Forest is a supervised ensemble learning algorithm. Ensemble means that it takes a bunch of weak learners and have them work together to form one

strong predictor. Here, we have a collection of decision trees, known as Forest. To classify a new object based on attributes, each tree gives a classification and we say the tree votes for that class. The forest chooses the classification having the most votes (over all the trees in the forest).

Our Model's configuration: number of weak learners = 100, maximum depth of each tree = 4

AdaBoost fits a sequence of weak learners on different weighted training data. It starts by predicting original data set and gives equal weight to each observation. If prediction is incorrect using the first learner, then it gives higher weight to observations which have been predicted incorrectly. Being an iterative process, it continues to add learner(s) until a limit is reached in the number of models or accuracy.

Our Model's configuration: number of weak learners = 50

Gradient Boosting Here, many models are trained sequentially. Each new model gradually minimizes the loss function ($y = ax + b + e$, where e is the error term) of the whole system using Gradient Descent method. The learning method consecutively fits new models to give a more accurate estimate of the response variable. The main idea behind this algorithm is to construct new base learners which can be optimally correlated with negative gradient of the loss function, relevant to the whole ensemble.

Our Model's configuration: number of weak learners = 100, learning rate = [0.05, 0.1, 0.25], maximum features = 2, maximum depth = 2

3.4 Evaluation

For evaluating our classifiers we used below evaluation metrics [12].

Accuracy is the ratio of number of correct predictions to the total number of input samples. It works well only if there are equal number of samples belonging to each class. As we have, imbalanced data, we will also consider other metrics.

Area Under Curve(AUC) is used for binary classification problem. AUC of a classifier is equal to the probability that the classifier will rank a randomly chosen positive example higher than a randomly chosen negative example

Precision is the number of correct positive results divided by the number of positive results predicted by the classifier.

Recall is the number of correct positive results divided by the number of all relevant samples (all samples that should have been identified as positive).

F1 Score is the Harmonic Mean between precision and recall. The range for F1 Score is $[0, 1]$. It tells how precise our classifier is (how many instances it classifies correctly), as well as how robust it is (it does not miss a significant number of instances). High precision but lower recall, gives you an extremely accurate, but it then misses a large number of instances that are difficult to classify. The greater the F1 Score, the better is the performance of our model.

Confusion Matrix gives us a matrix as output and describes the complete performance of the model. It focuses on **True Positives** - the cases in which we predicted YES and the actual output was also YES; **True Negatives** - the cases in which we predicted NO and the actual output was NO; **False Positives** - the cases in which we predicted YES and the actual output was NO; **False Negatives** - the cases in which we predicted NO and the actual output was YES.

Stratified k-fold As our data is imbalanced, we trained our models using a stratified k-fold approach where the data is divided into k folds each of equal proportion of positives and negatives. These metrics would be more reliable and less biased. We used the value of k as 10.

Statistical testing For the purpose of comparing the performance of different classifiers, we performed paired t testing among the top three classifiers.

4 Experiments and Results

For all the experiments and development of classifiers, we used Python 3 and Google colab's Jupyter Notebook. We used libraries such as Skit Learn, Matplotlib, Seaborn, Pandas, Numpy and Imblearn. We used weka for implementing Decision Table.

We carried experiments with different input data; one with the original dataset, then with the undersampled dataset and last one with the oversampled dataset. We splitted out dataset in ratio of 75:25 for training and testing purpose.

Experiment 1 - Original Dataset: Post all the preprocessing steps (as mentioned above in the Methodology section), we ran all the implemented classifiers each one with the same input data (Shape: 92037 x 4). Figure 12 depicts two considered metrics (10-skfold Accuracy and Area Under Curve) for all the classifiers.

Accuracy wise Gradient Boosting with a learning rate of 0.25 performed best, coverage wise Random Forest and Decision Tree performed worsts.

Fig. 12. Experiment 1.

Experiment 2 - Undersampled Dataset: Post all the preprocessing steps (as mentioned above in the Methodology section) including the undersampling step, we ran all the implemented classifiers each one with the same input data (Shape: 54274 x 4). Figure 13 depicts two considered metrics (10-skfold Accuracy and Area Under Curve) for all the classifiers.

Fig. 13. Experiment 2.

Accuracy and coverage wise Logistic Regression performed best and Decision Tree performed worsts.

Experiment 3 - Oversampled Dataset: Post all the preprocessing steps (as mentioned above in the Methodology section) including the oversampling step, we ran all the implemented classifiers each one with the same input data (Shape: 191160 x 4). Figure 14 depicts two considered metrics (10-skfold Accuracy and Area Under Curve) for all the classifiers.

Fig. 14. Experiment 3.

Accuracy and coverage wise Decision Tree performed best and Logistic Regression performed worsts.

We have varying range of results with respect to different input data and different classifiers. Other metrics are followed in appendix.

5 Discussion

With the issues with our original dataset, we learned many things considering all the preprocessing steps that we carried to rectify them. The first important thing we learned is the importance of knowing your data. While imputing the missing value, we grouped two other features and calculated the mean instead of directly calculating the mean for all the instances. This way our imputed values were closer to the correct information. Another thing we learned is about the leaky features. While exploring our data, we came to that one of our feature (RiskMM) was used for generating the target variable and hence it made no sense to use this feature for predictions.

We learned about the curse of dimensionality while dealing with categorical variables which we solved using feature hashing. We also learned two techniques for performing feature selection - univariate selection and correlation heat map. We also explore undersampling and oversampling techniques while handling the class imbalance problem.

With the experiments that we carried using different data, we also came to know that in a few cases we have achieved higher accuracy (Decision Tree) clearly implying the classic case of overfitting. We also observed that the performance of classifiers varied with different input data. To count a few, Logistic Regression performed best with undersampled data whereas it performed worst with over-sampled data; same goes with KNN, it performed best with oversampled data and worst with undersampled data. Hence we can say that the input data has a very important role here. Ensembles to be precise Gradient Boosting performed pretty consistently in all the experiments.

6 Conclusion and Future Work

In this paper, we explored and applied several preprocessing steps and learned their impact on the overall performance of our classifiers. We also carried a comparative study of all the classifiers with different input data and observed how the input data can affect the model predictions.

We can conclude that Australian weather is uncertain and there is no such correlation among rainfall and the respective region and time. We figured certain patterns and relationships among data which helped in determining important features. Refer to the appendix section.

As we have a huge amount of data, we can apply Deep Learning models such as Multilayer Perceptron, Convolutional Neural Network, and others. It would be great to perform a comparative study between the Machine learning classifiers and Deep learning models.

References

1. World Health Organization: Climate Change and Human Health: Risks and Responses. World Health Organization, January 2003
2. Alcántara-Ayala, I.: Geomorphology, natural hazards, vulnerability and prevention of natural disasters in developing countries. *Geomorphology* 47(24), 107124 (2002)
3. Nicholls, N.: Atmospheric and climatic hazards: Improved monitoring and prediction for disaster mitigation. *Natural Hazards* 23(23), 137155 (2001)
4. [Online] InDataLabs, Exploratory Data Analysis: the Best way to Start a Data Science Project. Available: <https://medium.com/@InDataLabs/why-start-a-data-science-project-with-exploratory-data-analysis-f90c0efcbe49>
5. [Online] Pandas Documentation. Available: https://pandas.pydata.org/pandas-docs/stable/reference/api/pandas.get_dummies.html
6. [Online] Skit-Learn Documentation Available: https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.FeatureHasher.html
7. [Online] Skit-Learn Documentation Available: <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MinMaxScaler.html>
8. [Online] Skit Learn Documentation Available: https://scikit-learn.org/stable/modules/generated/sklearn.feature_selection.SelectKBest.html

9. [Online] Raheel Shaikh, Feature Selection Techniques in Machine Learning with Python Available: <https://towardsdatascience.com/feature-selection-techniques-in-machine-learning-with-python-f24e7da3f36e>
10. [Online] Imbalanced-learn Documentation Available: <https://imbalanced-learn.readthedocs.io/en/stable/introduction.html>
11. V. Veeralakshmi and D. Ramyachitra, Ripple Down Rule learner (RIDOR) Classifier for IRIS Dataset. Issues, vol 1, p. 79-85.
12. [Online] Aditya Mishra, Metrics to Evaluate your Machine Learning Algorithm Available: <https://towardsdatascience.com/metrics-to-evaluate-your-machine-learning-algorithm-f10ba6e38234>

Appendix

Knowledge / Pattern discovered from our dataset.

Fig. 15. Location and Rainfall.

With consideration of domain knowledge, we assumed that rainfall will be strongly related to location and season (date). However, we learned that there is no such pattern with our data. Humidity has a strong correlation with the target variable. However, with further exploration we observed that Humidity at 9 AM and at 3 PM are strongly related and hence we will only consider Humidity at 3 PM. We also observed that Rainfall today and tomorrow are related and hence

Fig. 16. Date (Month) and Rainfall.

Fig. 17. Humidity at 9 AM and Rainfall.

Fig. 18. Humidity at 3 PM and Rainfall..

Fig. 19. Rain today and Rainfall.

is a good feature for our predictions. Refer figure 15, 16, 17, 18 and 19 depicting this patterns.

Below are the evaluation metrics for all the experiments carried.

```

=== Evaluation on test split ===

Time taken to test model on test split: 0.98 seconds

=== Summary ===

Correctly Classified Instances      38771      79.5319 %
Incorrectly Classified Instances    9978      20.4681 %
Kappa statistic                    0.5905
Mean absolute error                 0.394
Root mean squared error             0.4446
Relative absolute error             78.7901 %
Root relative squared error         88.9195 %
Total Number of Instances          48749

=== Detailed Accuracy By Class ===

                TP Rate  FP Rate  Precision  Recall  F-Measure  MCC      ROC Area  PRC Area  Class
                0.890   0.300   0.749     0.890   0.813     0.601   0.832    0.827    No
                0.700   0.110   0.864     0.700   0.774     0.601   0.832    0.807    Yes
Weighted Avg.   0.795   0.205   0.806     0.795   0.793     0.601   0.832    0.817

=== Confusion Matrix ===

      a    b  <-- classified as
21727 2688 | a = No
 7290 17044 | b = Yes

```

Fig. 20. Decision Table using Weka.

	Gradient Boosting - 0.25	KNN - 25	Logistic Regression			Alpha = 0.05	
					GB - KNN	GB - LR	KNN - LR
1	0.739498895	0.733235077	0.737103906				
2	0.782608696	0.780397937	0.783161385	t-Value	-4.781971	-0.12177	2.721574
3	0.775607959	0.776529108	0.770081061	p-value	0.001	0.90576	0.02355
4	0.768607222	0.763080324	0.768975682				
5	0.756448047	0.746683861	0.747789241				
6	0.77210759	0.767686072	0.771002211		Significant	No	Significant
7	0.789240973	0.775792189	0.798636699		Difference	Significant	Difference
8	0.765941762	0.753225212	0.762992997			Difference	
9	0.718761519	0.711205308	0.724474751				
10	0.753040914	0.739402875	0.755621084				

Fig. 21. Paired t-test.

	Accuracy	Time Taken	NoRain_Precision	NoRain_f1-score	NoRain_recall	Rain_Precision	Rain_f1-score	Rain_recall	Sensitivity	Specificity
Gradient Boosting - 0.25	0.84447704	2.036684752	0.849403919	0.906922871	0.972797656	0.803773585	0.527423548	0.392482034	39.24820343	97.27976564
Gradient Boosting - 0.05	0.843050972	2.175215721	0.845746127	0.906482156	0.976616447	0.818955043	0.512158055	0.372581537	37.25815368	97.6616447
Gradient Boosting - 0.1	0.843580654	2.118630409	0.848328697	0.90646851	0.973163842	0.803749044	0.522571819	0.387138382	38.71383822	97.31638418
Logistic Regression	0.837092568	0.171946526	0.858378871	0.900513555	0.946998242	0.707486137	0.550620392	0.450691787	45.06917869	94.69982416
Ensemble AdaBoost	0.848993562	2.64432621	0.851922901	0.909631797	0.97572714	0.824773414	0.540995789	0.40250645	40.25064504	97.57271396
Random Forest	0.843778013	4.453097582	0.84515313	0.907063557	0.978761247	0.831114809	0.510344828	0.368227055	36.82270549	97.87612471
KNN - 29	0.849808492	4.521082401	0.853996513	0.90989978	0.973634652	0.816593886	0.549057989	0.41356432	41.35643199	97.36346516
KNN - 27	0.850378942	4.545867205	0.85461058	0.910202485	0.973530027	0.817064353	0.551757813	0.416513085	41.65130851	97.35300272
KNN - 25	0.85005297	4.532132864	0.854492008	0.909998043	0.973216154	0.815162455	0.551000488	0.416144489	41.61444895	97.3216154
Decision Tree	0.830576155	0.278856039	0.846474567	0.897842858	0.955848504	0.714479026	0.503937008	0.389237007	38.9237007	95.58485039

Fig. 22. Experiment 1 Evaluation Metrics.

	Accuracy	Time Taken	NoRain_Precision	NoRain_f1-score	NoRain_recall	Rain_Precision	Rain_f1-score	Rain_recall	Sensitivity	Specificity
Gradient Boosting - 0.25	0.84447704	2.036684752	0.849403919	0.906922871	0.972797656	0.803773585	0.527423548	0.392482034	39.24820343	97.27976564
Gradient Boosting - 0.05	0.843050972	2.175215721	0.845746127	0.906482156	0.976616447	0.818955043	0.512158055	0.372581537	37.25815368	97.6616447
Gradient Boosting - 0.1	0.843580654	2.118630409	0.848328697	0.90646851	0.973163842	0.803749044	0.522571819	0.387138382	38.71383822	97.31638418
Logistic Regression	0.837092568	0.171946526	0.858378871	0.900513555	0.946998242	0.707486137	0.550620392	0.450691787	45.06917869	94.69982416
Ensemble AdaBoost	0.848993562	2.64432621	0.851922901	0.909631797	0.97572714	0.824773414	0.540995789	0.40250645	40.25064504	97.57271396
Random Forest	0.843778013	4.453097582	0.84515313	0.907063557	0.978761247	0.831114809	0.510344828	0.368227055	36.82270549	97.87612471
KNN - 29	0.849808492	4.521082401	0.853996513	0.90989978	0.973634652	0.816593886	0.549057989	0.41356432	41.35643199	97.36346516
KNN - 27	0.850378942	4.545867205	0.85461058	0.910202485	0.973530027	0.817064353	0.551757813	0.416513085	41.65130851	97.35300272
KNN - 25	0.85005297	4.532132864	0.854492008	0.909998043	0.973216154	0.815162455	0.551000488	0.416144489	41.61444895	97.3216154
Decision Tree	0.830576155	0.278856039	0.846474567	0.897842858	0.955848504	0.714479026	0.503937008	0.389237007	38.9237007	95.58485039

Fig. 23. Experiment 2 Evaluation Metrics.

	Accuracy	Time Taken	NoRain_Precision	NoRain_f1-score	NoRain_recall	Rain_Precision	Rain_f1-score	Rain_recall	Sensitivity	Specificity
Gradient Boosting - 0.25	0.84447704	2.036684752	0.849403919	0.906922871	0.972797656	0.803773585	0.527423548	0.392482034	39.24820343	97.27976564
Gradient Boosting - 0.05	0.843050972	2.175215721	0.845746127	0.906482156	0.976616447	0.818955043	0.512158055	0.372581537	37.25815368	97.6616447
Gradient Boosting - 0.1	0.843580654	2.118630409	0.848328697	0.90646851	0.973163842	0.803749044	0.522571819	0.387138382	38.71383822	97.31638418
Logistic Regression	0.837092568	0.171946526	0.858378871	0.900513555	0.946998242	0.707486137	0.550620392	0.450691787	45.06917869	94.69982416
Ensemble AdaBoost	0.848993562	2.64432621	0.851922901	0.909631797	0.97572714	0.824773414	0.540995789	0.40250645	40.25064504	97.57271396
Random Forest	0.843778013	4.453097582	0.84515313	0.907063557	0.978761247	0.831114809	0.510344828	0.368227055	36.82270549	97.87612471
KNN - 29	0.849808492	4.521082401	0.853996513	0.90989978	0.973634652	0.816593886	0.549057989	0.41356432	41.35643199	97.36346516
KNN - 27	0.850378942	4.545867205	0.85461058	0.910202485	0.973530027	0.817064353	0.551757813	0.416513085	41.65130851	97.35300272
KNN - 25	0.85005297	4.532132864	0.854492008	0.909998043	0.973216154	0.815162455	0.551000488	0.416144489	41.61444895	97.3216154
Decision Tree	0.830576155	0.278856039	0.846474567	0.897842858	0.955848504	0.714479026	0.503937008	0.389237007	38.9237007	95.58485039

Fig. 24. Experiment 3 Evaluation Metrics.