
Imbalance-XGBoost: leveraging weighted and focal losses for binary label-imbalanced classification with XGBoost

A Preprint

Chen Wang

Department of Computer Science
Rutgers University - New Brunswick
Piscataway, NJ 08854, USA
chen.wang.cs@rutgers.edu

Chengyuan Deng

Department of Computer Science
Rutgers University - New Brunswick
Piscataway, NJ 08854, USA
charles.deng@rutgers.edu

Suzhen Wang

Department of Health Statistics
Weifang Medical University
Weifang, Shandong, 261053 China
wangsz@wfmcc.edu.cn

August, 2021

ABSTRACT

The paper presents *Imbalance-XGBoost*, a Python package that combines the powerful XGBoost software with weighted and focal losses to tackle binary label-imbalanced classification tasks. Though a small-scale program in terms of size, the package is, to the best of our knowledge, the first of its kind which provides an integrated implementation for the two loss functions on XGBoost and brings a general-purpose extension to XGBoost for label-imbalanced scenarios. In this paper, the design and usage of the package are discussed and illustrated with examples. Furthermore, as the first- and second-order derivatives of the loss functions are essential for the implementations, the algebraic derivation is discussed and it can be deemed as a separate contribution. The performances of the methods implemented in the package are extensively evaluated on Parkinson's disease classification dataset, and multiple competitive performances are presented with the ROC and Precision-Recall (PR) curves. To further assert the superiority of the methods, the performances on four other benchmark datasets from the UCI machine learning repository are additionally reported. Given the scalable nature of XGBoost, the package has great potentials to be broadly applied to real-life binary classification tasks, which are usually of large-scale and label-imbalanced.

Keywords Imbalanced Classification · XGBoost · Python Package

1 Introduction

XGBoost is an advanced Gradient Tree Boosting-based software that can efficiently handle large-scale Machine Learning tasks ([1]). Merited by its performance superiority and affordable time and memory complexities, it has been widely applied to a variety of research fields since been proposed, ranging from cancer diagnosis ([2]) and medical record analysis ([3]) to credit risk assessment ([4]) and metagenomics ([5]). Also, because of its easy-to-use Python interface and explainable nature, it has become the de facto method-of-the-first-choice for a majority of data science problems. For instance, in the famous *Kaggle* (<https://www.kaggle.com/competitions>) competitions, many winning teams built their models based on XGBoost and expressed positive views on the method and its variations ([1, 6, 7]). It could be tentatively predicted that in the near future, XGBoost and its variations will remain one of the most-applied methods in the data science community.

On the other hand, although XGBoost has achieved considerable success on both regression and classification problems, its performance often becomes subtle when a situation of label-imbalanced classification is brought up. The problem of classification under label-imbalanced situations exists commonly in practice, and most accuracy-driven 'vanilla' Machine Learning methods suffer performance decline if the ratio between labels becomes biased. For example, in cancer diagnosis, it is common to have 95% of the patients without cancer and only 5% with it. If the model simply predict everyone as 'no cancer', then the accuracy is 95%, which is remarkably high. However, missing to spot any cancer patient can lead to fatal consequences.

There have been mixed reports on the capabilities of XGBoost in handling label-imbalanced data. For example, [8] demonstrated through their experiments that XGBoost can outperform other methods on skewed datasets, while the figures in [9] suggested that vanilla XGBoost must be combined with other ensembling methods to achieve satisfactory results. Given the popularity of XGBoost and the fact that label-skewed data is, unfortunately, commonly encountered in practice, this performance decay will leave significant negative effects on related research and applications.

This paper introduces *imbalance-XGBoost*, an XGBoost-based Python package addressing the label-imbalanced issue in the binary label regime by implementing weighted (cross-entropy) and focal losses on the boosting machine. Weighted cross-entropy loss is one of the simplest algorithm-level cost-sensitive methods ([10]) for learning imbalanced data. It follows the straightforward idea to increase the penalization of the misclassification of a certain class, and it has been widely applied to adjust vanilla machine learning algorithms to the label-imbalanced domain ([11]). In contrast, focal loss ([12]) is a relatively novel method originated from research in object detection. The idea of the method is to add a $(1 - y_j)^\gamma$ factor to the cross-entropy function (where y_j is the prediction of $p(x = j)$), and this will reduce the importance of the well-classified data points. Comparing with weighted cross-entropy, focal loss enjoys a more robust parameter configuration as the method will work in our favor as long as $\gamma > 0$.

To the best of our knowledge, there has not been significant publication discussing the implementation of the two losses on XGBoost previously. Existing studies on XGBoost under label-imbalanced scenarios usually adopt data-level approaches such as re-sampling ([13]) and/or cost-sensitive loss with non-trainable *a priori* modifications ([14]). [15] mentioned weighted XGBoost in their work, but details regarding the implementation are not presented. A major challenge in applying the two loss functions to XGBoost is that to approximate the incremental learning objective, first- and second-order derivatives of the loss functions with respect to the predictions must be presented (One can refer to section 3 for more details on this). And an algebraic contribution of this paper is the derivations and implementations of the derivatives that enable the two losses to be run with XGBoost.

The package is written in Python with hard dependency on XGBoost, Numpy ([16]), and Scikit-learn ([17]). The losses are integrated into the XGBoost system by the customized loss framework of the software, provided the derived expressions of the derivatives. Since the major methods in the program are included in the dependency graph, the core part of the package is of small scale, with only a few hundred lines of codes. Nevertheless, the function derivatives and implementations and the significance in practical applications make the work non-trivial. The software has been released on Github¹ and PyPI², and it has started to attract users with considerable downloads.

The rest of the paper goes as follows. Section 2 introduces the package from the perspectives of toolkit designers and users; Section 3 provides the theoretical foundation of the second-order approximation of gradient boosting trees used in XGBoost and the first- and second-order derivatives of the two losses; Some related studies are surveyed and discussed in section 4, then the experimental performances of the package are empirically examined on five imbalanced binary classification datasets in section 5; And finally, section 6 gives a general conclusion of the paper.

2 Design and Usage of Imbalance-XGBoost

2.1 Code Design

Though the XGBoost method has implementations in multiple languages, Python is picked as the language-of-choice for its wide recognition and application in data science. The codes follow the standard of *PEP8*, and the project has been designed as open-source with codes on the Github page. We strive to keep the program consistent with 'standard' practices in Python-based data science. The input data is designated as *Numpy array* ([16]), but by explicitly adding `np.array()` conversion, data types compatible with *Numpy array* (e.g. *Pandas Dataframe* ([18])) will also work on the package. As a small project, the usage of it can be clearly presented with the Readme file, and there is no additional documentation required.

¹<https://github.com/jhwjhw0123/Imbalance-XGBoost>

²<https://pypi.org/project/imbalance-xgboost/>

The overall program is consist of three classes: one main class `imbalance_xgboost`, which contains the method the users will be applying, and two customized-loss classes, `Weight_Binary_Cross_Entropy` and `Focal_Binary_Loss`, on which the imbalanced losses are based. The loss functions are designed as separate classes for the convenience of parameter tuning, and they are not supposed to be called by the users. When initializing an `imbalance_xgboost` object, keyword `special_objective` will be recorded by the `__init__()` method. Then, when executing the `fit()` function, the corresponding loss function object will be instantiated with the given parameter, and the built-in `xgb.train()` method of XGBoost will be able to fit the model based on the customized loss function. Figure 1 illustrates the overall structure of the program.

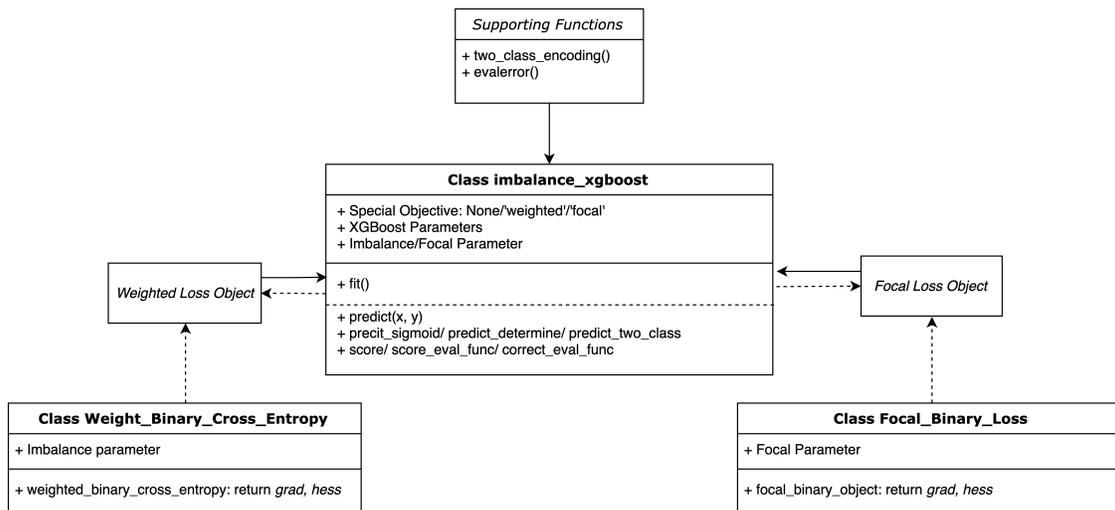


Figure 1: The Overall Structure of the Program

The package is designed to be an estimator class of the *Scikit-learn* toolkit. This scheme enables the model and parameter selection methods in *Scikit-learn*, such as `GridsearchCV()` and `RandomizedSearchCV()`, to be directly applied to find the best parameters for imbalanced XGBoosts. In practical data science tasks, this feature is crucial as the optimal models rely heavily on parameter tuning and selection. Also, estimator in *Scikit-learn* can be combined with other estimators (transformers) by integrating them to a `Pipeline` object ([17]). This allows the weighed- and focal-XGBoost to be easily combined with other pre-processing methods, such as resampling, to produce more robust results. Section 2.2 will provide more details for the package to tune parameters and perform cross-validation with *Scikit-learn*.

2.2 Model Optimization and Evaluation with Scikit-Learn

Similar to common classifiers in *Scikit-learn*, the best classifier/parameter can be obtained through exhaustive or random search with the functions `GridsearchCV()`. It is noticeable that after fitting the model, it is possible to retrieve the 'plain' booster by accessing member `opt_booster.booster`, and the object will be a XGBoost class (instead of Imbalance-XGBoost class). This makes it possible for the user to train the model on a machine where *Imbalance-XGBoost* is available, save the model as 'plain' XGBoost, and run on another machine where only the original XGBoost package is installed.

Likewise, the cross-validation evaluation of *Imbalance-XGBoost* is similar to the process on an 'ordinary' classifier, and one can simply follow the documentation of *Scikit-Learn* to perform it. Notice that if one intends to perform a combination of parameter selection and cross-validation, it is recommended that a new booster with random initialization and selected parameters should be instantiated after the parameter selection, instead of directly using the fitted optimal booster from `GridsearchCV()`. The reason for doing so is that for the iterative training process of XGBoost, a booster trained from a randomized state is essential for a fair evaluation.

2.3 Built-in Evaluation Score

There are three evaluation functions in the package. The overriding `score()` function serves the purpose to evaluate prediction accuracy under the format of predictions, which are pre-sigmoid values (in range $[-\text{inf}, +\text{inf}]$) by default, by wrapping the sigmoid transformation and accuracy checking together. In comparison, function `score_eval_func()`

is the method to return metrics other than accuracy. In label-imbalanced binary classification, accuracy cannot reliably reveal the performance quality on its own as the metric can be 'tricked' by predicting all the instances as the majority class. This type of prediction will lead to high accuracy, yet the classifier actually does nothing. Thus, metrics taking 'preciseness' into accounts, such as precision, recall, F_1 score and Matthew's Correlation Coefficient (MCC) ([19]), are often applied for the scenario ([20]). Function `score_eval_func()` provides implementations for all the metrics mentioned above, and it can be overloaded by specifying the partial argument 'mode'(which can be accomplish by `functools.partial()`).

3 Theories and Derivatives

In this section, the mathematical foundations and derivations for the loss functions are discussed. For a high-level introduction, since XGBoost adopts an additive learning scheme with a second-order approximation, the first-order derivative (short-handed as 'gradient') and second-order derivative (noted as 'hessian' although somehow a misnomer) of the loss functions with respect to the prediction are required for fitting the model. To illustrate a clear mechanism, the section will first review the second-order approximation of additive tree boosting in section 3.1. Subsequently, the derivatives of gradients and Hessians of the weighted and focal losses will be discussed in section 3.2 and 3.3, respectively.

The notations used in this section will be as follows. We use m to denote the number of data and n for the number of features. The 'raw prediction' before the sigmoid function will be denoted as z_i , and the probabilistic prediction will be $\hat{y}_i = \sigma(z_i)$, where $\sigma(\cdot)$ represents the Sigmoid function. It is important to keep in mind that there is a discrepancy between the notations of this paper and the original XGBoost literature ([1]), as the \hat{y} in their analysis is denoted as z here. y_i is used to denote the true label, and α and γ are used for the parameters for the two loss functions, respectively.

3.1 Second-order Approximation of Gradient Boosting Tree

Denoting the input as \mathbf{x}_i and the target as y_i , according to [1], the additive learning objective used in practice is:

$$\mathcal{L}^{(t)} = \sum_{i=1}^n l(y_i, z_i^{(t-1)} + f_t(\mathbf{x}_i)) + \Omega(f_t) \quad (1)$$

Equation 1 is a *regularized objective*, where $l(\cdot, \cdot)$ denotes the loss between the targets (y_i) and predictions (z_i), and $\Omega(\cdot)$ denotes the complexity of the model. The tree is fitted in an additive manner with $z_i^{(t)} = z_i^{(t-1)} + f_t(\mathbf{x}_i)$, and t stands for the t -th iteration of the training process. Notice that the notations used above are slightly different from the original literature. Our goal is to find the $f_t(\cdot)$ of the t -th iteration to optimize objective 1; To achieve this, applying a second-order Taylor expansion on equation 1, one will get:

$$\begin{aligned} \mathcal{L}^{(t)} &\approx \sum_{i=1}^n [l(y_i, z_i^{(t-1)}) + g_i f_t(\mathbf{x}_i) + \frac{1}{2} h_i (f_t(\mathbf{x}_i))^2] + \Omega(f_t) \\ &\propto \sum_{i=1}^n [g_i f_t(\mathbf{x}_i) + \frac{1}{2} h_i [f_t(\mathbf{x}_i)]^2] + \Omega(f_t) \end{aligned} \quad (2)$$

The last line comes from the fact that the $l(y_i, z_i^{(t-1)})$ term can be removed from the learning objective as it is unrelated to the fitting of the model in the t -th iteration. In equation 2, there are $g_i = \frac{\partial L}{\partial z_i}$ and $h_i = \frac{\partial^2 L}{\partial z_i^2}$, which are the 'gradient' and 'hessian' terms mentioned before. Notice that both g_i and h_i are scalars, as individual boosting trees only deal with binary problems. Multi-class classification tasks are usually processed by an ensemble of binary classification trees (so-called *one-vs-all* scheme, see [21, 22]).

Since XGBoost does not provide automatic differentiation, the hand-derived derivatives will be essential. Meanwhile, the derived expressions have further potentials to be applied into other machine learning tasks. Therefore, the derivatives are discussed in sections 3.2 and 3.3. For both loss functions, sigmoid is selected as activation, and the following basic property of sigmoid will be consistently used in the derivatives:

$$\begin{aligned} \frac{\partial \hat{y}}{\partial z} &= \frac{\partial \sigma(z)}{\partial z} \\ &= \sigma(z)(1 - \sigma(z)) \\ &= \hat{y}(1 - \hat{y}) \end{aligned} \quad (3)$$

3.2 Weighted Cross-entropy Loss

The weighted cross-entropy loss for binary classification can be denoted as follows:

$$L_w = - \sum_{i=1}^m (\alpha y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)) \quad (4)$$

where α indicates the 'imbalance parameter'. Intuitively, if α is greater than 1, extra loss will be counted on 'classifying 1 as 0'; On the other hand, if α is less than 1, the loss function will weight relatively more on whether data points with label 0 are correctly identified.

The first order derivative is presented as follows:

$$\frac{\partial L_w}{\partial z_i} = -\alpha^{y_i} (y_i - \hat{y}_i) \quad (5)$$

The derivative is similar with the $\partial L/\partial z$ term for ordinary cross-entropy loss. A notable difference is that a α^{y_i} term is added to control the present of the parameter.

Taking derivative with respect to z_i again, one will get the second-order derivative:

$$\frac{\partial^2 L_w}{\partial^2 z_i} = \alpha^{y_i} (1 - \hat{y}_i)(\hat{y}_i) \quad (6)$$

After plugging equation 3 to the derivation.

3.3 Focal Loss

According to [12], the binary focal loss can be denoted as:

$$L_f = - \sum_{i=1}^m y_i (1 - \hat{y}_i)^\gamma \log(\hat{y}_i) + (1 - y_i) \hat{y}_i^\gamma \log(1 - \hat{y}_i) \quad (7)$$

As one can observe, if one sets $\gamma = 0$, the equation will become ordinary cross-entropy loss. Taking equation 3 into consideration, the first derivative of the focal loss can be denoted as:

$$\begin{aligned} \frac{\partial L_f}{\partial z_i} = & \gamma [(\hat{y}_i + y_i - 1)(y_i + (-1)^{y_i} \hat{y}_i)^\gamma \log(1 - y_i - (-1)^{y_i} \hat{y}_i)] \\ & + (-1)^{y_i} (y_i + (-1)^{y_i} \hat{y}_i)^{\gamma+1} \end{aligned} \quad (8)$$

And if γ is set to 0 in equation 8, the derivative will be the same as cross-entropy loss. The equation follows a clear structure, but it is still lengthy. To simplify the expression, one can set the following short-hand variables:

$$\begin{cases} \eta_1 = \hat{y}_i(1 - \hat{y}_i) \\ \eta_2 = y_i + (-1)^{y_i} \hat{y}_i \\ \eta_3 = \hat{y}_i + y_i - 1 \\ \eta_4 = 1 - y_i - (-1)^{y_i} \hat{y}_i \end{cases} \quad (9)$$

Plugging these representations into equation 8, the first-order derivative can be denoted as:

$$\frac{\partial L_f}{\partial z_i} = \gamma \eta_3 \eta_2^\gamma \log(\eta_4) + (-1)^{y_i} \eta_2^{\gamma+1} \quad (10)$$

Finally, taking derivatives with respect to z_i , and combining with equation 3 and 10, one can get the second-order derivative ('hessian'), which can be denoted as:

$$\begin{aligned} \frac{\partial^2 L_f}{\partial z_i^2} = & \eta_1 \left\{ \gamma [(\eta_2^\gamma + \gamma (-1)^{y_i} \eta_3 \eta_2^{\gamma-1}) \log(\eta_4) - \frac{(-1)^{y_i} \eta_3 \eta_2^\gamma}{\eta_4}] \right. \\ & \left. + (\gamma + 1) \eta_2^\gamma \right\} \end{aligned} \quad (11)$$

Again, if $\gamma = 0$, the second-order derivative becomes $\eta_1 = \hat{y}_i(1 - \hat{y}_i)$, which matches the formula of the second-order derivative of ordinary cross-entropy.

4 Related Work

The paper is built on the foundation of the original papers of XGBoost ([1]) and focal loss ([12]), and the methodology to program customized loss function is provided in the software's Github page¹. XGBoost is based on the algorithm

¹https://github.com/dmlc/xgboost/blob/master/demo/guide-python/custom_objective.py

of gradient tree boosting ([23]), and this method has been deemed as a powerful Machine Learning technique long before the XGBoost was born ([24]). Besides XGBoost, there are other implementations of gradient boosting, such as pGBRT ([25]), LightGBM ([26]), and CatBoost ([27]). Some of the implementations have additional features and are able to outperform XGBoost on some specific problems, but XGBoost remains the method-of-the-first-choice in the data science community at large. As for the recently proposed focal loss, studies related to it are usually affiliated with Neural Networks and Deep Learning ([28, 29, 30]). The loss function is usually applied in an end-to-end manner with automatic differentiation, and to the best of our knowledge, there has not been any notable publication comprehensively discussing the derivatives of the loss function (despite the first-order derivative was briefly discussed in another form in the original paper).

Previous applications of XGBoost in label-imbalanced scenarios focus mostly on data-level algorithms. For example, [13] applies several commonly-used data resampling methods before using XGBoost for label-imbalanced breast cancer classification, and [31] utilized a more advanced under-sampling method called BalanceCascade ([32]) with XGBoost for credit scoring. Among the limited number of publications discussing algorithm-level modification for XGBoost in imbalanced classification, [14] used a *a priori* modification of the sigmoid activation to achieve a better result, but the loss function was unchanged. As it has been mentioned in section 1, [15] is by far the only implementation explicitly applied weighted function to XGBoost to best of our knowledge. It is noticeable that a Tensorflow-based gradient boosting implementation called *Tf Boosted Trees* ([33]) is able to run with the loss functions without the derivatives provided in this paper as it has an automatic differentiation mechanism. Nevertheless, it is a less popular package without supports of large-scale Machine Learning and compatibility with *Scikit-learn* toolkit.

As a common issue frequently encountered in practice, label-imbalanced classification has been intensively studied by researchers and there are multiple existing software programs designed to handle the problem. For a great example, [34] provides an integrated Python package called *Imbalanced-learn* for data-level resampling for imbalanced classification, and it has similar counterparts in the regime of other programming languages, such as *ROSE* in R ([35]). It is worth noting that the *Imbalanced-learn* package can be considered as an extension of *Scikit-learn*, and the Machine Learning toolkit itself also provides elementary methods to deal with label-imbalanced problems ([17]). Other software programs concerning label-imbalanced classification include popular Data Mining toolkits, such as KEEL ([36]) and WEKA ([37]). In addition, [38] provides a software containing a set of algorithms specifically for multi-class label-imbalanced problems, serving as one of the most recent studies on this topic.

5 Experiments

In this section, experimental results of the proposed Imbalance-XGBoost package are demonstrated in two parts: (1) Binary classification on Parkinson Disease data, a recently-collected imbalance dataset; (2) Binary classification on four benchmark datasets (with imbalance ratio) published in UCI Machine Learning Repository ([39]): *ecoli*(9:1), *arrhythmia*(17:1), *car_eval_4*(26:1) and *ozone*(42:1). To provide an observation of the performance improvement particularly by two loss functions, we compare the results of Weighted-XGBoost and Focal-XGBoost with vanilla XGBoost using the same parameter set, and the performances are also compared against SVM and Logistic regression.

5.1 Parkinson Disease Dataset and Setup

The Parkinson’s Disease(PD) classification data¹ is a recently introduced dataset ([40]) with 757 features categorized into 7 specific groups (two originally separate groups, Bandwidth and Formant, are merged in experiments), the data was gathered from 188 Parkinson’s disease patients and 64 healthy individuals at the Department of Neurology in Cerrahpaşa Faculty of Medicine, Istanbul University. Each individual corresponds to 3 records, and due to the differences between the number of participants of the two sides, a label imbalance ratio of 188:64 (roughly 3:1) emerges.

The original literature [40] demonstrated the classification performances with seven conventional classification algorithms and two ensemble approaches. The experiments in the original literature were conducted with leave-one-object-out cross validation. To keep consistent with the original system, the same cross validation setup is applied in this paper. Furthermore, the results in [40] illustrate a high accuracy and relatively lower F_1 score, indicating that the classifiers failed to tell the two classes clearly and likely achieved the performance by overwhelming predicting the majority class. This is an unfavourable behavior in label-imbalanced classification, and one advantage of *Imbalance-XGBoost* is that it does not suffer from this problem. This characteristic can be further verified by the ROC and Precision-Recall (PR) curves in figures 2-4.

To conduct the leave-one-object-out cross validation, the correctness collection function mentioned in section 2.3 is applied. By collecting results of True_Positive (TP), True_Negative (TN), False_Positive (FP), and False_Negative (FN),

¹available publicly, url: <https://archive.ics.uci.edu/ml/datasets/Parkinson%27s+Disease+Classification>

the confusion metric can be obtained and accuracy and F_1 score are computed accordingly. The records are evaluated in a *per-record* manner, which means the 3 records of one object (patient/healthy individual) will be evaluated individually, and 3 counts of the correctness will be added.

5.2 Classification Results and Discussion

Accuracy and F_1 score of the test set with 6 sets of features are presented in Table 1, where Best in [40] indicates the best performance of accuracy and F_1 score retrieved from the paper.

Table 1: Classification performance on individual features

	Baseline features		MFCC	
	Accuracy	F_1 score	Accuracy	F_1 score
Best in [40]	0.79	0.75	0.84	0.83
XGBoost	0.76	0.84	0.80	0.87
Weighted-XGBoost	0.76	0.85	0.80	0.87
Focal-XGBoost	0.76	0.85	0.82	0.89
	Wavelet Features		Bandwidth + Formant	
	Accuracy	F_1 score	Accuracy	F_1 score
Best in [40]	0.78	0.74	0.77	0.72
XGBoost	0.72	0.82	0.71	0.83
Weighted-XGBoost	0.75	0.85	0.74	0.85
Focal-XGBoost	0.75	0.85	0.75	0.85
	Intensity Based		Vocal Fold-Based	
	Accuracy	F_1 score	Accuracy	F_1 score
Best in [40]	0.77	0.74	0.77	0.74
XGBoost	0.73	0.83	0.72	0.82
Weighted-XGBoost	0.75	0.85	0.75	0.84
Focal-XGBoost	0.75	0.85	0.76	0.85

Without exception, although a slight declination of accuracy could be observed in weighted-XGBoost and focal-XGBoost, both classifiers generate a significantly higher F_1 score. The increase of F_1 score and the decrements of accuracy suggest that the previous-obtained higher accuracy is a consequence of overlooking minority class. Furthermore, for almost all the feature groups, the highest F_1 score is obtained by focal-XGBoost. This observation can be explained from an algorithmic perspective that focal loss is more robust to parameters, while weighted loss is prone to the effect of sub-optimal parameters even if parameter search is applied.

[40] also applied a classifier with 50 top-ranked features selected by mRMR (minimum Redundancy-Maximum Relevance). The feature selection technique is based on the principle of maximizing the joint dependency of top ranking variables on the targeted one by reducing the redundancy among them ([41, 42]). For a comparison purpose, this paper employs the same technique with provided Python interface¹, and produces a subset of top-50 features to run with *Imbalance-XGBoost*. The performance of weighted- and focal-XGBoost on the top-50 features can be observed in table 2.

Table 2: Classification performance on top 50 features

	Top 50 Features	
	Accuracy	F_1 score
Best in [40]	0.86	0.84
XGBoost	0.83	0.89
Weighted-XGBoost	0.82	0.88
Focal-XGBoost	0.84	0.90

Consistent with the performance on individual groups of features, focal-XGBoost classifier has the highest F_1 score, slightly better than weighted-XGBoost. Both weighted- and focal-XGBoost outperform best classifier in [40] by a large

¹<https://github.com/fbrundu/pymrmr>

margin, and since the top-50 feature can be regarded as a 'master subset', the superiority of the methods implemented in imbalance-XGBoost can be further corroborated.

To provide further insights for the predictions of weight- and focal-XGBoosts, the ROC and Precision-Recall (PR) curves of the two boosting trees and two other methods (SVM and Logistic Regression, which are used in [40]) are illustrated in figures 2-4. Limited by the space, only the feature sets of Baseline, MFCC and Top 50 are adopted. From the figures, it can be observed that the XGBoost-based models consistently provide superior performance in terms of AUC and PR curve. This further validate our earlier assertion that the higher accuracy of the method in [40] comes from mis-classifying the minority instances into the majority class.

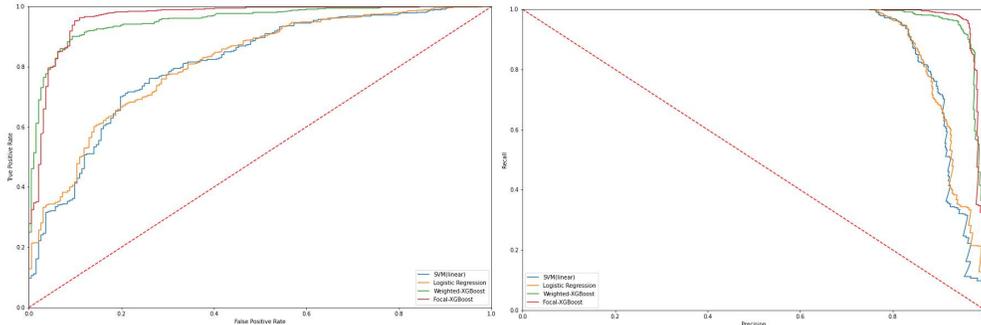


Figure 2: ROC and Precision-Recall Curve of Baseline Feature

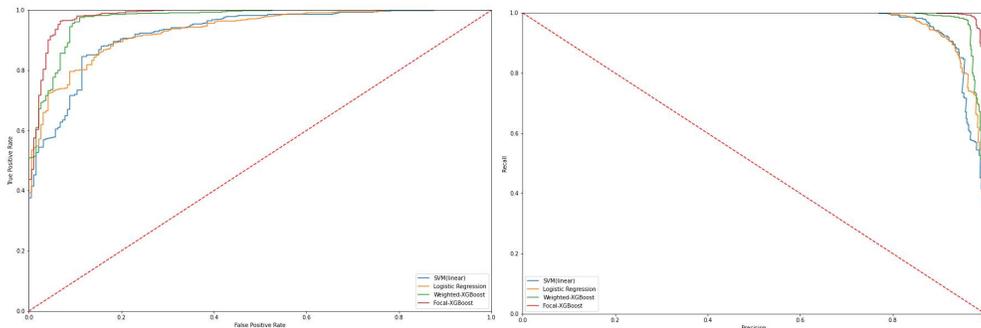


Figure 3: ROC and Precision-Recall Curve of MFCC Feature

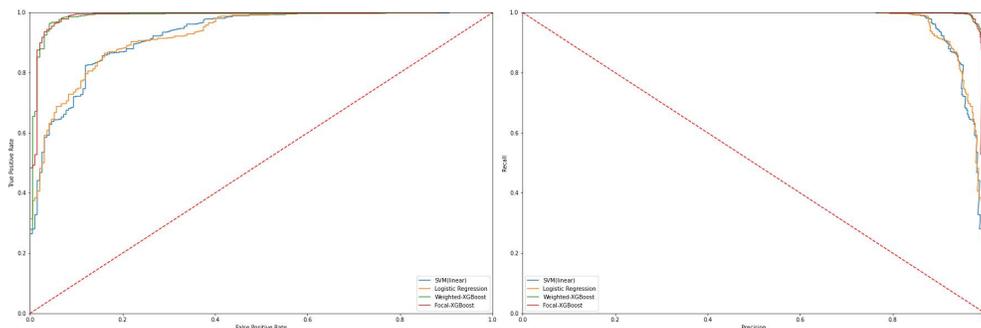


Figure 4: ROC and Precision-Recall Curve of Top50 Features

5.3 Classification for Benchmark Imbalanced Datasets

To further evaluate the performance of Imbalance-XGBoost, we introduce four other imbalanced datasets for binary classification from UCI Machine Learning Repository: ecoli, arrhythmia, car_eval_4 and ozone. Notice that the Parkinson Disease dataset has an imbalance ratio of 3:1, while these four datasets have higher ratios progressively as 9:1 (ecoli), 17:1 (arrhythmia), 26:1 (car_eval_4), 42:1 (ozone). Slightly different from the experiments on Parkinson

Dataset, we apply K-Fold(K=5) cross validation for Grid Search instead of leave-one-out, on the grounds that these four datasets include a considerable number of instances and K-Fold is commonly used in practice.

Our focus here is to illustrate the effectiveness of the weighted and focal loss functions, so we compare weighted- and focal-XGBoosts against its vanilla counterpart. The classification results are presented in table 3, with accuracy, F_1 score and MCC(Matthews Correlation Coefficient) as evaluation metrics.

Table 3: Classification performance on benchmark imbalanced datasets

	ecoli (9 : 1)			arrhythmia (17 : 1)		
	Accuracy	F_1 score	MCC	Accuracy	F_1 score	MCC
XGBoost	0.926	0.627	0.586	0.958	0.627	0.605
Weighted-XGBoost	0.935	0.645	0.616	0.967	0.681	0.665
Focal-XGBoost	0.935	0.645	0.616	0.960	0.679	0.662
	car_eval_4 (26 : 1)			ozone (42 : 1)		
	Accuracy	F_1 score	MCC	Accuracy	F_1 score	MCC
XGBoost	0.904	0.239	0.217	0.967	0.144	0.154
Weighted-XGBoost	0.959	0.268	0.266	0.959	0.268	0.266
Focal-XGBoost	0.911	0.287	0.273	0.964	0.180	0.173

From the table, it can be found that the classification accuracy of weighted- and focal-XGBoost is at least as competitive as XGBoost. In terms of F_1 score and MCC, which are considered the most crucial metrics for imbalanced classification, weighted- and focal-XGBoost outperforms vanilla XGBoost by a large margin on all datasets. We also remark that as the imbalance ratio goes up, the improvements on F_1 score and MCC for weighted- and focal-XGBoosts become more significant.

6 Conclusion

This paper presents a novel Python-based package, namely *Imbalance-XGBoost*, for binary label-imbalanced classification with XGBoost. The package implemented weighted cross-entropy and focal loss functions on XGBoost, and it is fully compatible with the popular *Scikit-learn* package in Python. The design and usage of the package are introduced, and the discussion of methods provide a clear and comprehensive user guidance. The theories and derivatives essential to the package are further discussed, and experiments based on five imbalance classification datasets are conducted with competitive performances illustrated. Overall, the package demonstrated in this paper successfully combines XGBoost with popular label-imbalance-robust loss functions and provides one of the most competitive performances up to date.

In summary, this paper has made three main contributions. Firstly, the paper has introduced a novel package that leverages the power of weighted and focal loss function for XGBoost, and it has huge potentials to be applied to a variety of real-life binary classification problems. Secondly, the paper has studied the theoretical foundations of the second-order approximation of XGBoost and has provided essential derivatives for the loss functions to be applied. The derivatives can also be applied to other fields in Machine Learning, and the equations in the merged form are convenient to be vectorized. And finally, the paper has offered new advanced performances on the five benchmark datasets, and the emphasis of the imbalanced nature provides new a perspective to study them.

References

- [1] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794. ACM, 2016.
- [2] Ching-Wei Wang, Yu-Ching Lee, Evelyne Calista, Fan Zhou, Hongtu Zhu, Ryohei Suzuki, Daisuke Komura, Shumpei Ishikawa, and Shih-Ping Cheng. A benchmark for comparing precision medicine methods in thyroid cancer diagnosis using tissue microarrays. *Bioinformatics*, 34(10):1767–1773, 2017.
- [3] Chen Wang, Suzhen Wang, Fuyan Shi, and Zaixiang Wang. Robust propensity score computation method based on machine learning with label-corrupted data. *arXiv preprint arXiv:1801.03132*, 2018.
- [4] Yung-Chia Chang, Kuei-Hu Chang, and Guan-Jhih Wu. Application of extreme gradient boosting trees in the construction of credit risk assessment models for financial institutions. *Applied Soft Computing*, 73:914–920, 2018.

- [5] Jyotsna Talreja Wassan, Haiying Wang, Fiona Browne, and Huiru Zheng. A comprehensive study on predicting functional role of metagenomes using machine learning methods. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 16(3):751–763, 2019.
- [6] Kamil Belkhaty Abou Omar. Xgboost and lgbm for porto seguro’s kaggle challenge: A comparison. *Semester Project, ETH Zurich*, 2018.
- [7] Didrik Nielsen. Tree boosting with xgboost-why does xgboost win ”every” machine learning competition? Master’s thesis, Norwegian University of Science and Technology, 2016.
- [8] Zhixun Zhao, Hui Peng, Chaowang Lan, Yi Zheng, Liang Fang, and Jinyan Li. Imbalance learning for the prediction of n 6-methylation sites in mrnas. *BMC genomics*, 19(1):574, 2018.
- [9] Ruisen Luo, Songyi Dian, Chen Wang, Peng Cheng, Zuodong Tang, YanMei Yu, and Shixiong Wang. Bagging of xgboost classifiers with random under-sampling and tomesk link for noisy label-imbalanced data. In *IOP Conference Series: 3rd International Conference on Automation, Control and Robotics Engineering (CACRE 2018)*, volume 428, page 012004. IOP Publishing, 2018.
- [10] Yanmin Sun, Andrew KC Wong, and Mohamed S Kamel. Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence*, 23(04):687–719, 2009.
- [11] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016.
- [12] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [13] Md Faisal Kabir and Simone Ludwig. Classification of breast cancer risk factors using several resampling approaches. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1243–1248. IEEE, 2018.
- [14] Yufei Xia, Chuanzhe Liu, and Nana Liu. Cost-sensitive boosted tree for loan evaluation in peer-to-peer lending. *Electronic Commerce Research and Applications*, 24:30–49, 2017.
- [15] Wenbin Chen, Kun Fu, Jiawei Zuo, Xinwei Zheng, Tinglei Huang, and Wenjuan Ren. Radar emitter classification for large data set based on weighted-xgboost. *IET Radar, Sonar & Navigation*, 11(8):1203–1207, 2017.
- [16] Stefan Van Der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22, 2011.
- [17] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
- [18] Wes McKinney. pandas: a foundational python library for data analysis and statistics. *Python for High Performance and Scientific Computing*, 14, 2011.
- [19] Brian W Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451, 1975.
- [20] David Martin Powers. Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. *Journal of Machine Learning Technologies*, 2011.
- [21] Erin L Allwein, Robert E Schapire, and Yoram Singer. Reducing multiclass to binary: A unifying approach for margin classifiers. *Journal of machine learning research*, 1(Dec):113–141, 2000.
- [22] Günther Eibl and Karl-Peter Pfeiffer. Multiclass boosting for weak classifiers. *Journal of Machine Learning Research*, 6(Feb):189–210, 2005.
- [23] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 29(5):1189–1232, 2001.
- [24] Alexey Natekin and Alois Knoll. Gradient boosting machines, a tutorial. *Frontiers in neurorobotics*, 7:21, 2013.
- [25] Stephen Tyree, Kilian Q Weinberger, Kunal Agrawal, and Jennifer Paykin. Parallel boosted regression trees for web search ranking. In *Proceedings of the 20th international conference on World wide web*, pages 387–396. ACM, 2011.
- [26] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, pages 3146–3154, 2017.

- [27] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. Catboost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, pages 6638–6648, 2018.
- [28] Xiaoliang Wang, Peng Cheng, Xinchuan Liu, and Benedict Uzochukwu. Focal loss dense detector for vehicle surveillance. In *2018 International Conference on Intelligent Systems and Computer Vision (ISCV)*, pages 1–5. IEEE, 2018.
- [29] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raska. Deepglobe 2018: A challenge to parse the earth through satellite images. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 172–181. IEEE, 2018.
- [30] Tianqi Zhang, Li-Ying Hao, and Ge Guo. A feature enriching object detection framework with weak segmentation loss. *Neurocomputing*, 335:72–80, 2019.
- [31] Hongliang He, Wenyu Zhang, and Shuai Zhang. A novel ensemble method for credit scoring: Adaption of different imbalance ratios. *Expert Systems with Applications*, 98:105–117, 2018.
- [32] Xu-Ying Liu, Jianxin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550, 2008.
- [33] Natalia Ponomareva, Soroush Radpour, Gilbert Hendry, Salem Haykal, Thomas Colthurst, Petr Mitrichev, and Alexander Grushetsky. Tf boosted trees: A scalable tensorflow based framework for gradient boosting. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 423–427. Springer, 2017.
- [34] Guillaume Lemaître, Fernando Nogueira, and Christos K Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *The Journal of Machine Learning Research*, 18(1):559–563, 2017.
- [35] Nicola Lunardon, Giovanna Menardi, and Nicola Torelli. Rose: A package for binary imbalanced learning. *R journal*, 6(1), 2014.
- [36] Jesús Alcalá-Fdez, Alberto Fernández, Julián Luengo, Joaquín Derrac, Salvador García, Luciano Sánchez, and Francisco Herrera. Keel data-mining software tool: data set repository, integration of algorithms and experimental analysis framework. *Journal of Multiple-Valued Logic & Soft Computing*, 17:255—287, 2011.
- [37] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. The weka data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.
- [38] Chongsheng Zhang, Jingjun Bi, Shixin Xu, Enislay Ramentol, Gaojuan Fan, Baojun Qiao, and Hamido Fujita. Multi-imbalance: An open-source software for multi-class imbalance learning. *Knowledge-Based Systems*, 174:137–143, 2019.
- [39] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [40] C Okan Sakar, Gorkem Serbes, Aysegul Gunduz, Hunkar C Tunc, Hatice Nizam, Betul Erdogan Sakar, Melih Tutuncu, Tarkan Aydin, M Erdem Isenkul, and Hulya Apaydin. A comparative analysis of speech signal processing algorithms for parkinson’s disease classification and the use of the tunable q-factor wavelet transform. *Applied Soft Computing*, 74:255–263, 2019.
- [41] Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 27(8):1226–1238, 2005.
- [42] Salvador García, Julián Luengo, and Francisco Herrera. Tutorial on practical tips of the most influential data preprocessing algorithms in data mining. *Knowledge-Based Systems*, 98:1–29, 2016.

Supplementary Materials

Two code listings regarding the usage of the *Imbalance-XGBoost* package are illustrated here. Details of the explanations of the codes can be found in section 2.

```

from imxgboost.imbalance_xgb import imbalance_xgboost as imb_xgb
# weighted XGBoost
xgboster_weighted = imb_xgb(special_objective='weighted')
xgboster_weighted.fit(X, y, imbalance_alpha=2.0)
# focal XGBoost
xgboster_focal = imb_xgb(special_objective='focal')
xgboster_focal.fit(X, y, focal_gamma=2.0)

```

Code Listing 1: Basic Usage of Imbalance-XGBoost

```

from imxgboost.imbalance_xgb import imbalance_xgboost as imb_xgb
from sklearn.model_selection import GridSearchCV, LeaveOneOut, cross_validate
# focal XGBoost
xgboster_focal = imb_xgb(special_objective='focal')
# cross-validation booster
CV_focal_booster = GridSearchCV(xgboster_focal, {"focal_gamma": [1.0, 1.5, 2.0, 2.5, 3.0]})
# fit the booster
CV_focal_booster.fit(X, y)
# retrieve the best model and parameter
opt_focal_booster = CV_focal_booster.best_estimator_
opt_focal_parameter = CV_focal_booster.best_params_
# instantiate an imbalance-xgboost instance for
# cross-validation
xgboost_focal_opt = imb_xgb(special_objective='focal', **xgboost_opt_param)
# initialize a leave-one splitter
loo_splitter = LeaveOneOut()
# Leave-One cross validation
loo_info_dict = cross_validate(xgboost_focal_opt, X=x, y=y, cv=loo_splitter)

```

Code Listing 2: Parameter Tuning and Model Evaluation of Imbalance-XGBoost with SK-learn