

# Unified Robust Boosting

ZHU WANG<sup>1,\*</sup>

<sup>1</sup>*Memphis, TN, Department of Preventive Medicine, The University of Tennessee Health Science Center, United States*

## Abstract

Boosting is a popular algorithm in supervised machine learning with wide applications in regression and classification problems. It combines weak learners, such as regression trees, to obtain accurate predictions. However, in the presence of outliers, traditional boosting may yield inferior results since the algorithm optimizes a convex loss function. Recent literature has proposed boosting algorithms that optimize robust nonconvex loss functions. Nevertheless, there is a lack of weighted estimation to indicate the outlier status of observations. This article introduces the iteratively reweighted boosting (IRBoost) algorithm, which combines robust loss optimization and weighted estimation. It can be conveniently constructed with existing software. The output includes weights as valuable diagnostics for the outlier status of observations. For practitioners interested in the boosting algorithm, the new method can be interpreted as a way to tune robust observation weights. IRBoost is implemented in the R package *irboost* and is demonstrated using publicly available data in generalized linear models, classification, and survival data analysis.

**Keywords** *boosting; CC-family; IRBoost; IRCO; machine learning; robust method*

## 1 Introduction

Boosting is a powerful supervised machine learning algorithm. As an ensemble method, boosting combines many weak learners to generate a strong prediction. Being a functional descent method, boosting finds wide applications in regression and classification problems. Friedman (2001) and Friedman et al. (2000) discussed boosting for a variety of convex loss functions. Boosting can be utilized to fit various models with different base learners, including linear least squares, smoothing splines, and regression trees (Bühlmann and Hothorn, 2007; Wang, 2018b).

The R package *mboost* implements robust boosting with the Huber loss for continuous responses (Hothorn et al., 2023). Likewise, the package *xgboost* (Chen et al., 2024) implements a pseudo-Huber loss function. Notably, these loss functions are convex. While convex loss functions have computational advantages and are commonly used due to their ease of optimization, they can lack robustness to leverage points in the predictor variables (Maronna et al., 2019) and provide poor approximations to certain loss functions like the 0-1 loss (Wu and Liu, 2007; Zhao et al., 2010; Park and Liu, 2011) in classification problems.

On the other hand, nonconvex loss functions, such as Tukey’s biweight loss, can be robust to both vertical outliers in the response variable and leverage points. Moreover, nonconvex loss functions can achieve better generalization accuracy.

Recent research has explored strategies to address outliers in boosting algorithms with

---

\* Email: [zwang145@uthsc.edu](mailto:zwang145@uthsc.edu).

nonconvex loss functions. For example, a study on boosting in the presence of outliers examined the role and efficiency of nonconvex loss functions for binary classification problems, aiming to adaptively handle outliers during the boosting process (Li and Bradic, 2018). Wang (2018a) and Wang (2018b) proposed robust functional gradient boosting for nonconvex loss functions in the context of regression and classification problems, implemented in the R package *bst* (Wang and Hothorn, 2023). These methods applied a majorization-minimization (MM) scheme, an extension of the popular expectation-maximization (EM) algorithm in statistics. These MM algorithms involve quadratic approximations of functions and difference-of-convex functions. However, they fail to generate weights as an indication of the outlier status of the observations. Small weights should be assigned to observations deviating from the underlying model.

In classical robust estimation, weights are derived from robust loss functions, such as the Huber loss. Recent progress has been made in generating weights from robust loss functions in more complex problems. Wang (2024b) innovatively proposed a new framework for robust estimation by reducing the weight of the observation that leads to a large loss. The author introduced a unified class of robust loss functions known as the concave-convex (CC) family. Additionally, the author proposed iteratively reweighted convex optimization (IRCO), a specialized application of the MM algorithm designed to minimize the loss functions within the CC family. The CC-family includes traditional robust loss functions such as the Huber loss, robust hinge loss for support vector machines, and robust exponential family for generalized linear models. IRCO can be conveniently implemented with existing methods and software.

In this article, we integrate IRCO and boosting into IRBoost for the CC-family. This functional optimization is more general than the parameter-based estimation in Wang (2024b). For instance, IRBoost permits a function space derived from regression trees. Unlike previous robust boosting methods, including Wang (2018a,b), the major novelty is that IRBoost provides a unified framework for a large class of robust loss functions to estimate weights and identify outliers.

We illustrate the proposed algorithm using the R package *irboost* (Wang, 2024a), applying it to various models within the robust exponential family, such as regression, logistic regression, and Poisson regression. Additionally, we demonstrate its application to robust survival regression with the accelerated failure time model. The package further includes implementations of IRBoost for gamma regression, Tweedie regression, hinge classification, and multinomial logistic regression.

## 2 Robust Boosting

### 2.1 CC-Family Function Estimation

To unify robust estimation, Wang (2024b) proposed the concave convex family with functions  $\Gamma = g \circ s$  satisfying the following conditions:

1.  $g$  is a nondecreasing closed concave function whose domain is the range of function  $s$ .
2.  $s$  is convex on  $\mathbb{R}$ .

Examples of concave component  $g$  are listed in Table 1. Note that the function `tcave` is not differentiable everywhere but subdifferentiable. The parameter  $\theta$  controls robustness level a model is allowed to have, and a smaller value leads to a more robust estimation. See Wang (2024b) for a discussion of the motivation behind these functions. In classification problems, we assume  $y \in \{-1, 1\}$ . The convex component includes common loss functions in regression and classification such as squared loss  $s(u) = u^2$  and the negative log-likelihood function in the

Table 1: Concave component with  $\theta > 0$ .

Concave function	$g(z), z \geq 0$
hcave	$\begin{cases} z & \text{if } z \leq \theta^2/2 \\ \theta(2z)^{\frac{1}{2}} - \frac{\theta^2}{2} & \text{if } z > \theta^2/2 \end{cases}$
acave	$\begin{cases} \theta^2(1 - \cos(\frac{(2z)^{\frac{1}{2}}}{\theta})) & \text{if } z \leq \theta^2\pi^2/2 \\ 2\theta^2 & \text{if } z > \theta^2\pi^2/2 \end{cases}$
bcave	$\frac{\theta^2}{6} (1 - (1 - \frac{2z}{\theta^2})^3 I(z \leq \theta^2/2))$
ccave	$\theta^2 (1 - \exp(\frac{-z}{\theta^2}))$
dcave	$\frac{1}{1 - \exp(-\theta)} \log(\frac{1+z}{1+z \exp(-\theta)})$
ecave	$\begin{cases} \frac{2 \exp(\frac{-\delta}{\theta})}{\sqrt{\pi \theta \delta}} z & \text{if } z \leq \delta \\ \operatorname{erf}(\sqrt{\frac{z}{\theta}}) - \operatorname{erf}(\sqrt{\frac{\delta}{\theta}}) + \frac{2 \exp(\frac{-\delta}{\theta})}{\sqrt{\pi \theta \delta}} \delta & \text{if } z > \delta \end{cases}$
gcave	$\begin{cases} \frac{\delta^{\theta-1}}{(1+\delta)^{\theta+1}} z & \text{if } z \leq \delta \\ \frac{1}{\theta} (\frac{z}{1+z})^\theta - \frac{1}{\theta} (\frac{\delta}{1+\delta})^\theta + \frac{\delta^\theta}{(1+\delta)^{\theta+1}} & \text{if } z > \delta \end{cases}$
	where $\delta = \begin{cases} \rightarrow 0+ & \text{if } 0 < \theta < 1 \\ \frac{\theta-1}{2} & \text{if } \theta \geq 1 \end{cases}$
tcave	$\min(\theta, z), \theta \geq 1$ for classification; $\theta > 0$ otherwise

exponential family adopted by generalized linear models. Other examples include negative log-likelihood functions for multinomial logistic regression, Tweedie regression, and the accelerated failure time model for time-to-event data subject to censoring (Barnwal et al., 2022).

Given a set of observations  $(\mathbf{x}_i, y_i), i = 1, \dots, n$ , where  $y_i \in \mathbb{R}$  and  $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top \in \mathbb{R}^p$ , denote  $\Omega$  as the linear span of a set  $H$  of base learners, including regression trees and linear predictor functions. An estimation of  $\mathbf{f} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^\top \in \Omega$  can be obtained by minimizing an empirical loss function

$$\sum_{i=1}^n \ell(y_i, f(\mathbf{x}_i)), \quad (1)$$

where  $\ell$  is a CC-family member,  $\ell = g \circ s = g(s(u))$ . With some abuse of notation,  $s(u)$  is also used to denote  $s(y, f(\mathbf{x}))$ . For instance,  $s(u) = s(y - f(\mathbf{x}))$  in regression, and  $s(u) = s(yf(\mathbf{x}))$  in a margin-based classification. To simplify the notation,  $f$  is often used to replace  $f(\mathbf{x})$ .

The robust function estimation problem (1) can be addressed using Algorithm 1, employing an iteratively reweighted boosting approach. This two-layer algorithm operates as follows: in each iteration at the outer layer, observation weights are updated based on the current loss values from the convex component. With these weights, the inner layer updates function estimation and loss values from the convex component through weighted boosting iterations minimizing weighted convex loss functions. In a new cycle, the outer layer weights are updated using the boosting results. Subsequently, the inner layer runs new boosting iterations with the updated

weights. The process repeats until convergence. This approach represents a generalization of the iteratively reweighted least squares method commonly employed in robust estimation (Maronna et al., 2019; Wang, 2024b).

Step 3 involves  $\varphi$ , the Fenchel conjugate of  $-g$  defined as:

$$\varphi(v) = \sup_{z \in \text{dom } g} \left( zv + g(z) \right).$$

Here,  $\partial(-g(z))$  means the subdifferential of function  $-g$  at point  $z$ , which is a set of slopes that touch the graph of  $-g$  at  $(z, -g(z))$  and bound the graph from below. If the function is differentiable at  $z$ , then  $\partial(-g(z))$  contains only one element  $-g'(z)$ . Let

$$\rho(\mathbf{f}^{(k)}) = \sum_{i=1}^n \ell(y_i, f_i^{(k)}),$$

where  $\mathbf{f}^{(k)}$  is generated by the algorithm. We then have the following convergence results for IRBoost.

---

**Algorithm 1** IRBoost.

---

- 1: **Input:** training samples  $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ , concave component  $g$  with parameter  $\theta$ , convex component  $s$ , starting points  $z_i, i = 1, \dots, n$  and iteration count  $K$ .
  - 2: **for**  $k = 1$  to  $K$  **do**
  - 3:   Compute subgradient  $v_i^{(k)}$  via  $v_i^{(k)} \in \partial(-g(z_i))$  or  $z_i \in \partial\varphi(v_i^{(k)})$ ,  $i = 1, \dots, n$ .
  - 4:   Compute  $\mathbf{f}^{(k)} = \operatorname{argmin}_{\mathbf{f} \in \Omega} \sum_{i=1}^n s(y_i, f_i)(-v_i^{(k)})$ .
  - 5:   Compute  $z_i = s(y_i, f_i^{(k)})$ ,  $i = 1, \dots, n$ .
  - 6: **end for**
  - 7: **Output:**  $v_i^{(K)}$  and  $\mathbf{f}^{(K)}$ .
- 

**Theorem 1.** *Suppose that  $g$  is a concave component in the CC-family, and  $g$  is bounded below. Then the loss function values  $\rho(\mathbf{f}^{(k)})$  generated by Algorithm 1 are nonincreasing and converge.*

This result extends Theorem 6 in Wang (2024b), which focused on linear predictor functions. Our investigation encompasses more expansive function spaces. Conversely, when  $\Omega$  represents a linear space of linear models, Theorem 1 aligns with the results detailed in Wang (2024b). Algorithm 1 is an MM algorithm, and the proof follows the same argument as Theorem 6 in Wang (2024b). Hence, only a sketch of the proof is given below, specifically assuming that  $g$  is differentiable.

For a differentiable concave function  $g$ , the first-order condition is  $\forall u, v \in \text{dom } g$

$$g(u) \leq g(v) + g'(v)(u - v). \quad (2)$$

Substitute  $u$  with  $s(u)$ , and  $v$  with  $s(v)$  in (2), we get

$$g(s(u)) \leq g(s(v)) + g'(s(v))(s(u) - s(v)). \quad (3)$$

Substitute  $s(u) = s(y_i, f_i)$ ,  $s(v) = s(y_i, f_i^{(k)})$  in (3), and sum up for  $i = 1, \dots, n$ , we get

$$\sum_{i=1}^n g(s(y_i, f_i)) \leq \sum_{i=1}^n g\left(s(y_i, f_i^{(k)})\right) + g'\left(s(y_i, f_i^{(k)})\right) \left(s(y_i, f_i) - s(y_i, f_i^{(k)})\right). \quad (4)$$

Let  $Q(\mathbf{f}|\mathbf{f}^{(k)})$  denote the right hand side of (4). We then have

$$\rho(\mathbf{f}) \leq Q(\mathbf{f}|\mathbf{f}^{(k)}), \quad \rho(\mathbf{f}^{(k)}) = Q(\mathbf{f}^{(k)}|\mathbf{f}^{(k)}). \quad (5)$$

This implies that  $Q(\mathbf{f}|\mathbf{f}^{(k)})$  majorizes  $\rho(\mathbf{f})$  at  $\mathbf{f}^{(k)}$ . The algorithm iterates as follows: given an estimate  $\mathbf{f}^{(k)}$  in the  $k$ th iteration,  $Q(\mathbf{f}|\mathbf{f}^{(k)})$  is minimized in the  $k+1$  iteration to obtain an updated minimizer  $\mathbf{f}^{(k+1)}$ . This process is repeated until convergence. Therefore, the algorithm generates a descent sequence of estimates:

$$\rho(\mathbf{f}^{(k+1)}) \leq Q(\mathbf{f}^{(k+1)}|\mathbf{f}^{(k)}) \leq Q(\mathbf{f}^{(k)}|\mathbf{f}^{(k)}) = \rho(\mathbf{f}^{(k)}). \quad (6)$$

Alternatively, the majorization (5) can be constructed from a different surrogate function derived from the Fenchel convex conjugate. The Fenchel-Moreau theorem states that the following result holds (Wang, 2024b)

$$g(s(u)) = \inf_{v \in \text{dom } \varphi} (s(u)(-v) + \varphi(v)).$$

Let

$$R(\mathbf{f}|\mathbf{f}^{(k)}) = \sum_{i=1}^n s(y_i, f_i)(-v_i^{(k)}) + \varphi(v_i^{(k+)}).$$

$R(\mathbf{f}|\mathbf{f}^{(k)})$  is another function that majorizes  $\rho(\mathbf{f})$  at  $\mathbf{f}^{(k)}$ . The algorithm generates a sequence of estimates in a descending order, similar to (6).

In step 3 of IRBoost, weights are computed using two different methods, each associated with a surrogate function, namely  $Q$  and  $R$ . Remarkably, the solutions obtained from these different approaches have been demonstrated to be identical (Wang, 2024b).

## 2.2 Boosting Algorithm for Function Estimation

In this section, we describe methods to compute step 4 in Algorithm 1. For ease of notation, we present methods for unweighted estimation, as weighted estimation does not pose technical difficulties:

$$\operatorname{argmin}_{\mathbf{f} \in \Omega} \sum_{i=1}^n s(y_i, f_i). \quad (7)$$

Boosting techniques addressing the problem (7) are well-documented in the literature; see Chen and Guestrin (2016); Sigrist (2021), and references therein. Briefly, the boosting solution is an additive model given by

$$\hat{f}_i = F_M(\mathbf{x}_i) = \sum_{m=1}^M t_m(\mathbf{x}_i), \quad i = 1, \dots, n, \quad (8)$$

where  $F_M(\mathbf{x}_i)$  is stagewise constructed by sequentially adding an update  $t_m(\mathbf{x}_i)$  to the current estimate  $F_{m-1}(\mathbf{x}_i)$ :

$$F_m(\mathbf{x}_i) = F_{m-1}(\mathbf{x}_i) + t_m(\mathbf{x}_i), \quad m = 1, \dots, M. \quad (9)$$

There are different ways to compute  $\mathbf{t}_m(\mathbf{x}) = (t_m(\mathbf{x}_1), \dots, t_m(\mathbf{x}_n))^T$ : gradient and Newton-type updates are the most popular. When the second derivative of the loss function exists, the Newton-type update is preferred over the gradient update to achieve fast convergence (Sigrist, 2021).

$$\mathbf{t}_m(\mathbf{x}) = \operatorname{argmin}_{\mathbf{f} \in H} \sum_{i=1}^n h_{m,i} \left( -\frac{d_{m,i}}{h_{m,i}} - f(x_i) \right)^2, \quad (10)$$

where the first and second derivatives of the loss function  $s$  for observation  $i$  are given by:

$$d_{m,i} = \frac{\partial}{\partial f} s(y_i, f)|_{f=F_{m-1}(x_i)},$$

$$h_{m,i} = \frac{\partial^2}{\partial f^2} s(y_i, f)|_{f=F_{m-1}(x_i)}.$$

For quadratic loss  $s(y_i, f) = \frac{(y_i - f)^2}{2}$ , we obtain  $h_{m,i} = 1$ . In this case, the Newton-update is reduced to the gradient update. Note that  $f \in \Omega$  in (7) and  $f \in H$  in (10), indicating that optimal base learners can be found in  $H$  and a linear combination of them can be found in  $\Omega$  through the boosting algorithm.

To prevent overfitting, boosting also implements a step-size shrinkage parameter  $0 < \eta \leq 1$  in the update (9):

$$F_m(\mathbf{x}_i) = F_{m-1}(\mathbf{x}_i) + \eta t_m(\mathbf{x}_i), \quad m = 1, \dots, M.$$

### 2.3 Penalized Estimation

Another strategy to avoid overfitting is to add a regularization term to the objective function (1):

$$\sum_{i=1}^n \ell(y_i, f_i) + \sum_{m=1}^M \Lambda(t_m), \quad (11)$$

where  $\Lambda$  penalizes the model complexity.

If  $H$  is the linear space of linear models with a  $p$ -dimensional predictor, i.e.,  $t_m(\mathbf{x}_i) = \mathbf{x}_i^\top \boldsymbol{\beta}_m$ ,  $\boldsymbol{\beta}_m = (\beta_{1m}, \dots, \beta_{pm})^\top$ , let

$$\Lambda(t_m) = \frac{1}{2} \lambda \sum_{j=1}^p \beta_{jm}^2 + \alpha \sum_{j=1}^p |\beta_{jm}|,$$

where  $\lambda \geq 0$ ,  $\alpha \geq 0$  are the  $L_2$  and  $L_1$  regularization parameters, respectively. Note that  $\Lambda(t_m)$  provides shrinkage estimators and can conduct variable selection.

Suppose that  $H$  is the linear space of regression trees. Each regression tree splits the predictor space into disjoint hyper-rectangles with sides parallel to the coordinate axes (Wang, 2018b). Specifically, denote the hyper-rectangles in the  $m$ -th boosting iteration as  $R_{jm}$ ,  $j = 1, \dots, J$ . Let  $t_m(\mathbf{x}_i) = \beta_{jm}$ ,  $\mathbf{x}_i \in R_{jm}$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, J$ . With  $\gamma \geq 0$ , the penalty can be defined as in Chen and Guestrin (2016):

$$\Lambda(t_m) = \gamma J + \frac{1}{2} \lambda \sum_{j=1}^J \beta_{jm}^2 + \alpha \sum_{j=1}^J |\beta_{jm}|.$$

Consequently, the boosting estimation in Section 2.2 requires modifications. The optimization problem (7) is changed to

$$\operatorname{argmin}_{f \in \Omega} \sum_{i=1}^n s(y_i, f_i) + \sum_{m=1}^M \Lambda(t_m).$$

The Newton-type update (10) is modified accordingly:

$$\mathbf{t}_m(\mathbf{x}) = \operatorname{argmin}_{f \in H} \sum_{i=1}^n h_{m,i} \left( -\frac{d_{m,i}}{h_{m,i}} - f(x_i) \right)^2 + \Lambda(t_m).$$

Step 4 in Algorithm 1 can be modified accordingly.

## 2.4 Implementation and Tuning Parameter Selection

The requirement of  $z \geq 0$  on the domain of  $g$ , that is,  $s(u) \geq 0$ , may be relaxed for some concave functions  $g$  in Table 1. To satisfy all  $g$  functions, however, it is simpler to require  $s(u)$  to be non-negative. When  $s(u) < 0$ , such as in the case of a negative log-likelihood value for the gamma distribution, we can ensure a nonnegative loss by subtracting some data-dependent constant. Specifically, if  $s(y, f(\mathbf{x})) < 0$ , one remedy is to subtract a constant  $C$  such that  $s(y, f(\mathbf{x})) - C \geq 0$ . For the exponential family,  $s(y, f(\mathbf{x}_i)) \geq s(y, y)$  holds, as  $s(y, y)$  represents the negative log-likelihood value of a saturated model. Thus, a desired convex loss is  $s(y_i, f(x_i)) - \min_{i=1, \dots, n} s(y_i, y_i) \geq 0$ .

To minimize the penalized convex loss function (11) or its weighted version in step 4 of Algorithm 1, we employ the popular boosting R package, *xgboost*. There are two layers of iterations in the algorithm: the outer layer updates weights, and the inner layer comprises boosting iterations, where early stopping of iterations does not guarantee convergence. On the other hand, the output  $\mathbf{f}^{(K)}$  may overfit the data. In practice, a two-stage process may be considered: In the first stage, Algorithm 1 is applied to obtain optimal weights for the given observations. In the second stage, a data-driven method such as cross-validation can be used to compute a reliable estimate of errors for a weighted boosting model. This process can also be utilized to determine an optimal parameter. For instance, Algorithm 2 can be used to select an optimal robust parameter  $\theta$  and the corresponding IRBoost model. To illustrate the process, a data example will be presented in Section 3.1. Other parameters, such as the boosting iteration  $M$ , penalty number  $\gamma$  for trees, regularization terms  $\lambda$  and  $\alpha$ , can be similarly chosen. Alternatively, since the  $\theta$  parameter is typically considered a hyperparameter, a more computationally convenient approach in the literature is to conduct estimation for different values of  $\theta$  and compare the results. One can begin with a large value of  $\theta$  for less robust estimation and move towards smaller values of  $\theta$  for more robust results.

---

### Algorithm 2 IRBoost tuning parameter selection.

---

- 1: **Input:** training samples  $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ , concave component  $g$  with parameters  $\theta_1, \dots, \theta_T$ , convex component  $s$ , starting points  $z_i, i = 1, \dots, n$  and iteration count  $K$ .
  - 2: **for**  $t = 1$  to  $T$  **do**
  - 3:   Compute IRBoost model  $\mathbf{f}_t$  and robustness weights  $-v_{it}^{(K)}, i = 1, \dots, n$  with  $\theta_t$  using Algorithm 1.
  - 4:   Compute cross-validation errors  $\epsilon_t$  using weighted boosting algorithm with weights  $-v_{it}^{(K)}, i = 1, \dots, n$ .
  - 5: **end for**
  - 6: **Output:**  $\hat{t} = \operatorname{argmin}_{1 \leq t \leq T} \epsilon_t, \hat{\theta} = \theta_{\hat{t}}, \hat{v}_i^{(K)} = v_{i\hat{t}}^{(K)}, \hat{\mathbf{f}} = \mathbf{f}_{\hat{t}}$ .
- 

Nonconvex loss optimization algorithms may lead to local solutions depending on the initial values. In the implementation of *irboost*, a user can specify the initial values  $z_i \geq 0, i = 1, \dots, n$ , with default values of  $z_i$  being the initial weights if provided or the vector of 1s in that order. The user may then compare the results from potentially different solutions and obtain an optimal solution afterward.

The source version of the *irboost* package is freely available from the Comprehensive R Archive Network (<http://CRAN.R-project.org>). The reader can install the package directly from the R prompt via:



```
install.packages("irboost")
```

All analyses presented below are contained in a package vignette. The rendered output of the analyses is available by the R-command:

```
library("irboost")
vignette("static_irbst", package = "irboost")
```

## 3 Applications

### 3.1 Robust Boosting for Regression

In this example, we predict the median value of owner-occupied homes in the suburbs of Boston, using data publicly available from the UCI machine learning data repository. The dataset comprises 506 observations and 13 predictors. Wang (2024b) provides an alternative robust estimation for comparison.

```
urlname <- "https://archive.ics.uci.edu/ml/"
filename <- "machine-learning-databases/housing/housing.data"
dat <- read.table(paste0(urlname, filename), sep = "", header = FALSE)
dat <- as.matrix(dat)
colnames(dat) <- c("CRIM", "ZN", "INDUS", "CHAS", "NOX", "RM",
  "AGE", "DIS", "RAD", "TAX", "PTRATIO", "B", "LSTAT", "MEDV")
p <- dim(dat)[2]
```

We apply IRBoost with the concave component `bcave` and the convex component least squares.

```
library("irboost")
param <- list(objective = "reg:squarederror", max_depth = 2)
fit1 <- irboost(data = dat[, -p], label = dat[, p], cfun = "bcave",
  s = 10, params = param, verbose = 0, nrounds = 50)
plot(fit1$weight_update, ylab = "Weight") # plot robustness weights
id <- sort.list(fit1$weight_update)[1:4] # 4 obs. with smallest weights
text(id, fit1$weight_update[id] - 0.02, id, col = "red") # highlight 4 obs.
```

Figure 1 displays the observation weights used when IRBoost converges, highlighting the four smallest values, which are considered outliers. We plot the observed median housing prices against the predicted values in Figure 2(a). Notably, the four observations with the smallest weights deviate significantly from their predicted values, but this outcome is not surprising. IRBoost returns a weighted boosting estimation. The implementation of `irboost` is equivalent to `xgboost` with weights. Using the weights learned from IRBoost with XGBoost results in predictions that are identical to those produced by IRBoost. This equivalence is illustrated in the following example with Figure 2(b).






Figure 1: Robustness weights of IRBoost for the Boston housing data.




Figure 2: Prediction for the Boston housing data.

```

par(pty = "s")
plot(dat[, p], predict(fit1, newdata = dat[, -p]), xlab = "Observations",
     ylab = "Predictions") # obs. vs predictions
text(dat[id, p], predict(fit1, newdata = dat[id, -p]) - 1, id,
     col = "red") # highlight 4 obs. with smallest weights
abline(0, 1, col = "red") # 45-degree line
library("xgboost")
fit_xg <- xgboost::xgboost(data = dat[, -p], label = dat[, p],
                          weight = fit1$weight_update, params = param, verbose = 0,
                          nrounds = fit1$niter)
par(pty = "s") # plot type square between irboost and xgboost
plot(predict(fit1, newdata = dat[, -p]), predict(fit_xg, newdata = dat[,

```

```
-p]), xlab = "Predictions by irboost", ylab = "Predictions by xgboost")
abline(0, 1, col = "red") # 45-degree line
```

We can compare computing times between `irboost` and `xgboost`. As shown below, both computing tasks completed within one second on an Intel® Core™ i9-10900X CPU @ 3.70GHz × 8 processor with 16GB of RAM. Since the former involves iterative reweighting runs of `xgboost`, it is expected to take more computing time than a single run of `xgboost`.

```
# computing time for irboost
system.time(irboost(data = dat[, -p], label = dat[, p], cfun = "bcave",
  s = 10, params = param, verbose = 0, nrounds = 50))["elapsed"]
```

```
## elapsed
## 0.941
```

```
# computing time for xgboost
system.time(xgboost::xgboost(data = dat[, -p], label = dat[,
  p], weight = fit1$weight_update, params = param, verbose = 0,
  nrounds = fit1$niter))["elapsed"]
```

```
## elapsed
## 0.127
```

Feature importance from the learned model is displayed in Figure 3. The figure reveals that the top two factors for predicting median housing prices are the average number of rooms per dwelling (RM) and the percentage values of the lower status of the population (LSTAT).

```
importance_matrix <- xgboost::xgb.importance(model = fit1) # importance metric
xgboost::xgb.plot.importance(importance_matrix = importance_matrix) # plot
```




Figure 3: Variable importance measures for the Boston housing data.

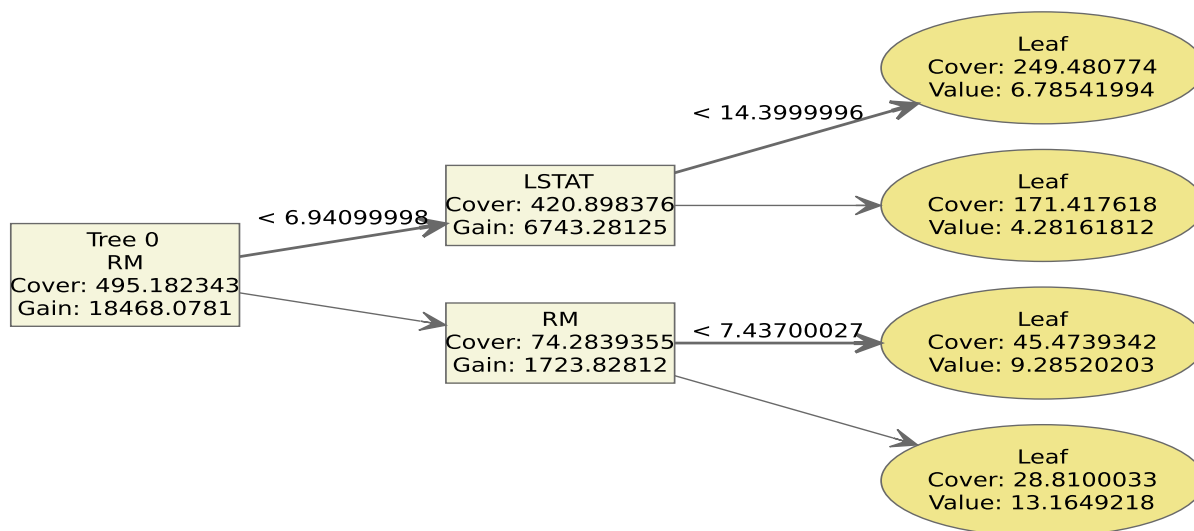


Figure 4: First tree in IRBoost for the Boston housing data.

The first tree used to build the model is depicted in Figure 4.

```
xgboost::xgb.plot.tree(model = fit1, trees = 0)
```

We can optimize tuning parameters using cross-validation with the built-in functions in *xgboost*. First, update the data format with the estimated robustness weights, then determine the optimal IRBoost iteration.

```
dtrain <- xgboost::xgb.DMatrix(data = dat[, -p], label = dat[,
  p]) # create a DMatrix for training data
xgboost::setinfo(dtrain, "weight", fit1$weight_update) # set weight information
param <- list(booster = "gbtree", objective = "reg:squarederror")
set.seed(136) # set the seed for reproducibility
xgbcv <- xgboost::xgb.cv(params = param, data = dtrain, nrounds = 200,
  early_stopping_rounds = 20, nfold = 5, prediction = TRUE)
```

```
xgbcv$best_iteration
```

```
## [1] 37
```

Continuing in the same vein, we can identify the optimal robustness parameter  $\theta$ , and the corresponding IRBoost iteration. For instance, to select a preferable  $\theta$  from the set  $\{5, 10\}$ , the cross-validation results below indicate that  $\theta = 5$  yields a smaller root mean squared error on the test data. Moreover, this procedure identifies the optimal IRBoost iteration as 39.

```
dtrain_cv <- xgboost::xgb.DMatrix(data = dat[, -p], label = dat[,
  p]) # training data
robustness_param <- c(5, 10) # two theta values
```

```

res <- NULL
for (i in 1:length(robustness_param)) {
  fit_init <- irboost(data = dat[, -p], label = dat[, p], cfun = "bcave",
    s = robustness_param[i], params = list(objective = "reg:squarederror",
      max_depth = 2), verbose = 0, nrounds = 50) # fit irboost model
  xgboost::setinfo(dtrain_cv, "weight", fit_init$weight_update) # new weights
  set.seed(136)
  xgbcv <- xgboost::xgb.cv(params = param, data = dtrain_cv,
    early_stopping_rounds = 20, nrounds = 200, nfold = 5,
    prediction = TRUE) # 5-fold CV
  reslog <- xgbcv$evaluation_log[xgbcv$best_iteration] # best values in CV
  tmp <- unlist(c(theta = robustness_param[i], reslog)) # combine theta
  res <- rbind(res, tmp) # combine results from previous theta
  rownames(res) <- NULL
}

```

```
print(res, digits = 2)
```

```

##      theta iter train_rmse_mean train_rmse_std test_rmse_mean
## [1,]     5   39           0.30           0.019           2.3
## [2,]    10   37           0.37           0.036           3.0
##      test_rmse_std
## [1,]           0.31
## [2,]           0.33

```

### 3.2 Robust Logistic Boosting

A binary classification problem, as proposed by Long and Servedio (2010), involves a response variable  $y$  randomly chosen to be  $-1$  or  $+1$  with equal probability. Symbols A, B, and C are randomly generated with probabilities 0.25, 0.25, and 0.5, respectively. The predictor vector  $\mathbf{x}$  with 21 elements is generated as follows: if A is obtained,  $x_j = y$  for  $j = 1, \dots, 21$ . If B is generated,  $x_j = y$  for  $j = 1, \dots, 11$ , and  $x_j = -y$  for  $j = 12, \dots, 21$ . If C is generated,  $x_j = y$ , where  $j$  is randomly chosen from the range of 1 to 11 with a selection of 5 elements, and from the range of 12 to 21 with a selection of 6 elements. For the remaining  $j \in \{1, 2, 3, \dots, 21\}$ ,  $x_j = -y$ . The training data is generated with  $n = 400$  samples, and the test data with  $n = 200$  samples.

We fit a robust logistic boosting model with the concave component `acave`, setting the maximum depth of a tree to 5. With a large parameter  $\theta = 100$ , the robustness weights are very close to 1 as shown below.

```

set.seed(1947)
dat2 <- dataLS(ntr = 400, nte = 200, percon = 0) # percon=0 means clean data
param <- list(objective = "binary:logitraw", max_depth = 5)
fit2 <- irboost(data = dat2$xtr, label = dat2$ytr, cfun = "acave",
  s = 100, params = param, verbose = 0, nrounds = 100)
range(fit2$weight_update) # range of robustness weights

```

```
## [1] 0.9999975 1.0000000
```




Figure 5: Robustness weights of IRBoost with  $\theta = 3$  for the contaminated simulation data.

To add outliers, we simulate data with 10% contamination in the response variables of the training data and then apply IRBoost. Figure 5 displays the robustness weights obtained from the algorithm.

```
set.seed(158)
dat3 <- dataLS(ntr = 400, nte = 200, percon = 0.1) # 10% data contamination
param <- list(objective = "binary:logitraw", max_depth = 5)
fit3 <- irboost(data = dat3$xtr, label = dat3$ytr, cfun = "acave",
  s = 3, params = param, verbose = 0, nrounds = 100)
plot(fit3$weight_update, ylab = "Weight") # plot robustness weights
```

In the third robust logistic boosting, we set the robustness hyperparameter value  $\theta$  to 1 ( $s=1$  in the `irboost` function) for a more robust estimation. Consequently, certain observations exhibit decreased weights, as illustrated in Figure 6.




Figure 6: Robustness weights of IRBoost with  $\theta = 1$  for the contaminated simulation data.




Figure 7: Classification errors and IRBoost iterations for the simulation data.

```
param <- list(objective = "binary:logitraw", max_depth = 5)
fit4 <- irboost(data = dat3$xtr, label = dat3$ytr, cfun = "acave",
  s = 1, params = param, verbose = 0, nrounds = 100)
plot(fit4$weight_update, ylab = "Weight") # plot robustness weights
```

The prediction accuracy can be compared for different models. The prediction error of the test data at each IRBoost iteration is depicted in Figure 7, demonstrating that the most accurate predictions come from the robust model, even in the presence of outliers.

```
err2 <- err3 <- err4 <- rep(NA, 100)
for (i in 1:100) {
  pred2 <- predict(fit2, newdata = dat2$xte, iterationrange = c(1,
    i + 1)) # prediction with the first i trees
  err2[i] <- mean(sign(pred2) != dat2$yte) # error at iteration i
  pred3 <- predict(fit3, newdata = dat3$xte, iterationrange = c(1,
    i + 1))
  err3[i] <- mean(sign(pred3) != dat3$yte)
  pred4 <- predict(fit4, newdata = dat3$xte, iterationrange = c(1,
    i + 1))
  err4[i] <- mean(sign(pred4) != dat3$yte)
}
plot(err2[1:100], ylim = c(0.05, 0.3), type = "l", xlab = "IRBoost iteration",
  ylab = "Classification error")
points(err3[1:100], col = "red", type = "l", lty = "dashed")
points(err4[1:100], col = "blue", type = "l", lty = "dotted")
legend("topright", lty = c("solid", "dashed", "dotted"), col = c("black",
  "red", "blue"), legend = c(expression(paste("clean data with ",
  theta == 100)), expression(paste("cont'd data with ", theta ==
  3)), expression(paste("cont'd data with ", theta == 1))))
```




Figure 8: Robustness weights of IRBoost for the iris data.

### 3.3 Robust Multiclass Boosting

In a 3-class classification using the `iris` dataset, we run IRBoost with concave function `acave` and  $\theta = 1$ . The robustness weights are illustrated in Figure 8, and the model achieves perfect prediction.

```
lb <- as.numeric(iris$Species) - 1 # convert text to numeric values
num_class <- 3
set.seed(11)
param <- list(objective = "multi:softprob", max_depth = 4, eta = 0.5,
              nthread = 2, subsample = 0.5, num_class = num_class)
fit5 <- irboost(data = as.matrix(iris[, -5]), label = lb, cfun = "acave",
               s = 1, params = param, verbose = 0, nrounds = 10)
```

```
plot(fit5$weight_update, ylab = "Weight") # plot robustness weights
```

```
# compute num_class probabilities per case
pred5 <- predict(fit5, newdata = as.matrix(iris[, -5]))
# reshape to a num_class-columns matrix
pred5 <- matrix(pred5, ncol = num_class, byrow = TRUE)
pred5_labels <- max.col(pred5) - 1 # probabilities to labels
sum(pred5_labels != lb) # classification errors
```

```
## [1] 0
```

### 3.4 Robust Poisson Boosting

A survey, collected from 3066 Americans, studied health care utilization (Heritier et al., 2009; Wang, 2024b). The dataset includes information on doctor office visits and 24 risk factors. A






Figure 9: Robustness weights of IRBoost for the doctor office visits data.

robust Poisson boosting model is fitted with the concave component `ccave`, and the estimated robustness weights are illustrated in Figure 9. Doctor office visits ranging from 200 to 750 in two years are highlighted for the 8 patients with the smallest weights.

```
data(docvisits, package = "mpath")
# convert factors and other types of variables into a
# numeric matrix
x <- model.matrix(~age + factor(gender) + factor(race) + factor(hispan) +
  factor(marital) + factor(arthri) + factor(cancer) + factor(hipress) +
  factor(diabet) + factor(lung) + factor(heart) + factor(stroke) +
  factor(psych) + factor(iadla) + factor(adlwa) + edyears +
  feduc + meduc + log(income + 1) + factor(insur) + 0, data = docvisits)
param <- list(objective = "count:poisson", max_depth = 1)
fit6 <- irboost(data = x, label = docvisits$visits, cfun = "ccave",
  s = 20, params = param, verbose = 0, nrounds = 50)
```

```
plot(fit6$weight_update, ylab = "Weight") # plot robustness weights
id <- sort.list(fit6$weight_update)[1:8] # 8 obs. with smallest weights
text(id, fit6$weight_update[id] - 0.02, docvisits$visits[id],
  col = "red") # highlight 8 obs.
```

### 3.5 Robust Survival Boosting with Accelerated Failure Time Model

Cox regression in survival analysis is based on a partial likelihood function. A robust extension in the CC-family is beyond the scope of this article. Alternatively, one may apply robust survival regression with the accelerated failure time model in *irboost*. The following code provides robust survival analysis for patients with advanced lung cancer from the North Central Cancer Treatment Group. Model performance is evaluated with multiple measures for survival data. Figure 10 illustrates the robustness weights using IRBoost.




Figure 10: Robustness weights of IRBoost for the lung cancer data.

```

library("survival")
lung1 <- lung[complete.cases(lung), ] # remove missing data
y_upper_bound <- rep(NA, dim(lung1)[1])
# set right-censoring obs. to y_upper_bound = Inf
for (i in 1:dim(lung1)[1]) {
  if (lung1$status[i] == 2) {
    y_upper_bound[i] <- lung1$time[i]
  } else y_upper_bound[i] <- Inf
}
x <- as.matrix(lung1[, !names(lung1) %in% c("time", "status")]) # predictors
dtrain <- xgboost::xgb.DMatrix(data = x, label_lower_bound = lung1$time,
  label_upper_bound = y_upper_bound) # input data format
param <- list(objective = "survival:aft", eval_metric = "aft-nloglik",
  aft_loss_distribution = "normal", aft_loss_distribution_scale = 1.2,
  max_depth = 3)
library("Hmisc")
fit7 <- irb.train(params = param, data = dtrain, cfun = "hcafe",
  s = 3, nrounds = 50)
# evaluate model prediction accuracy
Hmisc::rcorr.cens(predict(fit7, newdata = dtrain), Surv(time = lung1$time,
  event = lung1$status))

```

```

##          C Index          Dxy          S.D.          n
##  9.805945e-01  9.611889e-01  5.940054e-03  1.670000e+02
##          missing      uncensored Relevant Pairs      Concordant
##  0.000000e+00  1.200000e+02  2.112800e+04  2.071800e+04
##          Uncertain
##  6.572000e+03

```

```

plot(fit7$weight_update, ylab = "Weight") # plot robustness weights

```

## 4 Discussion

In this article, we introduced IRBoost as a unified robust boosting algorithm and demonstrated its versatility in regression, generalized linear models, classification, and time-to-event data analysis. The method served well for outlier detection and yielded more robust predictive models. Leveraging existing weighted boosting software, the approach could conveniently conduct tuning parameter selections and further explore the developed models for variable importance. The R package *irboost* proved to be a valuable tool in machine learning applications.

To make the IRBoost algorithm scalable for large-scale data sets, one direction is to use the stochastic majorization-minimization schemes (Mairal, 2013). These stochastic variants of the MM algorithm can efficiently optimize the objective function by using only a subset of the data at each iteration, rather than the entire dataset. This allows the algorithm to scale to very large problem sizes. Stochastic IRBoost remains a future research topic. The current R package *irboost* does not support user extensibility with other distributions. However, such functionality may be developed in a future release of the package.

## Acknowledgments

The author would like to thank the Editor and an Associate Editor for their constructive comments, which substantially help improve an early draft of the manuscript. The author expresses gratitude for the contributions of the R package *xgboost* to this work.

## Funding

This work was partially supported by the National Institute of Diabetes and Digestive and Kidney Diseases of the National Institutes of Health under Award Number R21DK130006.

## Supplementary Material

The R code necessary to reproduce the analysis presented in the manuscript is provided.

## References

- Barnwal A, Cho H, Hocking T (2022). Survival regression with accelerated failure time model in XGBoost. *Journal of Computational and Graphical Statistics*, 31(4): 1292–1302. <https://doi.org/10.1080/10618600.2022.2067548>
- Bühlmann P, Hothorn T (2007). Boosting algorithms: Regularization, prediction and model fitting (with discussion). *Statistical Science*, 22(4): 477–505.
- Chen T, Guestrin C (2016). Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- Chen T, He T, Benesty M, Khotilovich V, Tang Y, Cho H, et al. (2024). Xgboost: extreme gradient boosting. R package version 1.7.7.1.
- Friedman J (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5): 1189–1232. <https://doi.org/10.1214/aos/1013203451>

- Friedman J, Hastie T, Tibshirani R (2000). Additive logistic regression: A statistical view of boosting (with discussion and a rejoinder by the authors). *The Annals of Statistics*, 28(2): 337–407. <https://doi.org/10.1214/aos/1016218223>
- Heritier S, Cantoni E, Copt S, Victoria-Feser MP (2009). *Robust Methods in Biostatistics*, volume 825. John Wiley & Sons.
- Hothorn T, Bühlmann P, Kneib T, Schmid M, Hofner B, Otto-Sobotka F, et al. (2023). *mboost: Model-Based Boosting*. R package version 2.9-9.
- Li AH, Bradic J (2018). Boosting in the presence of outliers: Adaptive classification with non-convex loss functions. *Journal of the American Statistical Association*, 113(522): 660–674.
- Long PM, Servedio RA (2010). Random classification noise defeats all convex potential boosters. *Machine Learning*, 78(3): 287–304. <https://doi.org/10.1007/s10994-009-5165-z>
- Mairal J (2013). Stochastic majorization-minimization algorithms for large-scale optimization. In: *NIPS 2013 - Advances in Neural Information Processing Systems, 26, Dec 2013, South Lake Tahoe, United States*, 2283–2291.
- Maronna RA, Martin RD, Yohai VJ, Salibián-Barrera M (2019). *Robust Statistics: Theory and Methods (with R)*. John Wiley & Sons, Hoboken, NJ.
- Park SY, Liu Y (2011). Robust penalized logistic regression with truncated loss functions. *Canadian Journal of Statistics*, 39(2): 300–323. <https://doi.org/10.1002/cjs.10105>
- Sigrist F (2021). Gradient and Newton boosting for classification and regression. *Expert Systems with Applications*, 167: 114080. <https://doi.org/10.1016/j.eswa.2020.114080>
- Wang Z (2018a). Quadratic majorization for nonconvex loss with applications to the boosting algorithm. *Journal of Computational and Graphical Statistics*, 27(3): 491–502. <https://doi.org/10.1080/10618600.2018.1424635>
- Wang Z (2018b). Robust boosting with truncated loss functions. *Electronic Journal of Statistics*, 12(1): 599–650. <https://doi.org/10.1214/18-EJS1434>
- Wang Z (2024a). *irboost: Iteratively Reweighted Boosting for Robust Analysis*. R package version 0.1-15.
- Wang Z (2024b). Unified robust estimation. *Australian & New Zealand Journal of Statistics*, 66(1): 77–102. <https://doi.org/10.1111/anzs.12409>
- Wang Z, Hothorn T (2023). *bst: Gradient Boosting*. R package version 0.3-24.
- Wu Y, Liu Y (2007). Robust truncated hinge loss support vector machines. *Journal of the American Statistical Association*, 102(479): 974–983. <https://doi.org/10.1198/016214507000000617>
- Zhao L, Mammadov M, Yearwood J (2010). From convex to nonconvex: A loss function analysis for binary classification. In: *2010 IEEE International Conference on Data Mining Workshops*, 1281–1288. IEEE.