

# UC San Diego

## UC San Diego Previously Published Works

### Title

Boosting in the Presence of Outliers: Adaptive Classification With Nonconvex Loss Functions

### Permalink

<https://escholarship.org/uc/item/8d26g7ts>

### Journal

Journal of the American Statistical Association, 113(522)

### ISSN

0162-1459

### Authors

Li, Alexander Hanbo  
Brdic, Jelena

### Publication Date

2018-04-03

### DOI

10.1080/01621459.2016.1273116

Peer reviewed

# Boosting in the presence of outliers: adaptive classification with non-convex loss functions

Alexander Hanbo Li<sup>†</sup> and Jelena Bradic<sup>†</sup>

*Department of Mathematics  
University of California at San Diego<sup>†</sup>*

## Abstract

This paper examines the role and efficiency of the non-convex loss functions for binary classification problems. In particular, we investigate how to design a simple and effective boosting algorithm that is robust to the outliers in the data. The analysis of the role of a particular non-convex loss for prediction accuracy varies depending on the diminishing tail properties of the gradient of the loss – the ability of the loss to efficiently adapt to the outlying data, the local convex properties of the loss and the proportion of the contaminated data. In order to use these properties efficiently, we propose a new family of non-convex losses named  $\gamma$ -robust losses. Moreover, we present a new boosting framework, *Arch Boost*, designed for augmenting the existing work such that its corresponding classification algorithm is significantly more adaptable to the unknown data contamination. Along with the Arch Boosting framework, the non-convex losses lead to the new class of boosting algorithms, named adaptive, robust, boosting (ARB). Furthermore, we present theoretical examples that demonstrate the robustness properties of the proposed algorithms. In particular, we develop a new breakdown point analysis and a new influence function analysis that demonstrate gains in robustness. Moreover, we present new theoretical results, based only on local curvatures, which may be used to establish statistical and optimization properties of the proposed Arch boosting algorithms with highly non-convex loss functions. Extensive numerical calculations are used to illustrate these theoretical properties and reveal advantages over the existing boosting methods when data exhibits a number of outliers.

## 1 Introduction

Recent advances in technologies for cheaper and faster data acquisition and storage have led to an explosive growth of data complexity in a variety of research areas such as high-throughput genomics, biomedical imaging, high-energy physics, astronomy and economics. As a result, noise accumulation, experimental variation and data inhomogeneity have become substantial. Therefore, developing classification methods that are highly efficient and accurate in such settings, is a problem that is of great practical importance. However, classification in such settings is known to pose many statistical challenges and calls for new methods and theories. For binary classification problems, we assume the presence of separable, noiseless data that belong to two classes and in which an adversary has corrupted a number of observations from both classes independently. There are a number of setups that belong to this general framework. A random flipped label design, in which the labels of the class membership were randomly flipped, is one example that can occur very frequently, as labeling is prone to a number of errors, human or otherwise. Another example is the presence of outliers in the observations, in which a small number of observations from both classes have a variance that is larger than the noise of the rest of the observations. Such situations

may naturally occur with the new era of big and heterogeneous data, in which data are corrupted (arbitrarily or maliciously) and subgroups may behave differently; a subgroup might only be one or a few individuals in small studies that would appear to be outliers within class data.

Considerable effort has therefore been focused on finding methods that adapt to the relative error in the data. Although this has resulted in algorithms, e.g. Grünwald and Dawid (2004), that achieve provable guarantees (Natarajan et al., 2013; Kanamori et al., 2007) when contamination model (Scott et al., 2013) is known or when multiple noisy copies of the data are available (Cesa-Bianchi et al., 2011), good generalization errors in the test set are by no means guaranteed. This problem is compounded when the contamination model is unknown, where outliers need to be detected automatically. Despite progress on outlier-removing algorithms, significant practical challenges (due to exceedingly restrictive conditions imposed therein) remain. In this paper, we concentrate on the ensemble algorithms. Among these, AdaBoost (Freund and Schapire, 1997) has proven to be simple and effective in solving classification problems of many different kinds. The aesthetics and simplicity of AdaBoost and other forward greedy algorithms, such as LogitBoost (Friedman, et al., 2000), also facilitated a tacit defense from overfitting, especially when combined with early termination of the algorithm (Zhang and Yu, 2005). Friedman, et al. (2000) developed a powerful statistical perspective, which views AdaBoost as a gradient-based incremental search for a good additive model using the exponential loss. The gradient boosting (Friedman, 2001) and AnyBoost (Mason et al., 1999) have used this approach to generalize the boosting idea to wider families of problems and loss functions. This criterion was motivated by the fact that the exponential loss is a convex surrogate of the hinge or 0 – 1 loss. Nevertheless, in the presence of label noise and/or outliers, the performance of all of them deteriorates rapidly (Dietterich, 2000). Although algorithms like LogitBoost, MadaBoost (Domingo and Watanabe, 2000), Log-lossBoost (Collins, et al., 2002) are able to better tolerate noise than AdaBoost, they are still not insensitive to outliers. Hence, they are efficient when the data is observed with little or no noise. However, Long and Servedio (2010) pointed out that any boosting algorithm with convex loss functions is highly susceptible to a random label noise model. They constructed a simple example, from hereon denoted Long/Servedio problem, that cannot be “learned” by the boosting algorithms above.

Center to our analysis is the work by Freund (2009). He proposed a robust boosting algorithm based on the Boost by majority (BBM) (Freund, 1995) and BrownBoost (Freund, 2001) algorithm, which uses a non-convex loss function. Instead of maximizing the margin, the algorithms achieve robustness by allowing a preassigned  $\theta$  error of margin maximization (Servedio, 2003). Moreover, in each step, the algorithms update and solve a differential equation and update the preassigned remaining time  $c$  or the target error  $\epsilon$ . As the loss function changes with each iteration of the algorithm, they do not agree with the general boosting interpretation through additive models. Furthermore, with at least two preassigned parameters, each of them is difficult to implement and is highly inconsistent with respect to minor changes in the settings. Surprisingly, statistical and convergence properties of these two algorithms are still unknown. This leads to a natural question: how do we develop a simple but effective boosting algorithm that has a non-convex loss function, that preserves boosting interpretation and that is robust to the noise in the data? In this paper, we address this question and propose a fully automatic estimator, with no tuning parameters to be chosen, that has provable guarantees.

To design the new framework we will explore and amend the drawbacks of the AdaBoost algorithm in the contaminated data setting. We successfully identify that the AdaBoost’s sensitivity to outliers comes from the unbounded weight assignment of the misclassified observations. As outliers are more likely to be misclassified, they are very likely to be assigned large weights and will be repeatedly refitted in the following iterations. This refitting will deteriorate seriously the generalization performance of the algorithm, as the algorithm “learns” incorrect data distribution. To achieve robustness, the algorithm should be able to abandon observations that are on the extreme,

incorrect side of the boundary. Here, we theoretically and computationally investigate the applicability of non-convex loss functions for this purpose. We illustrate that the best weight updating rule is to assign a weight of  $-\phi'(y_i F(x_i))$  to each data point  $(x_i, y_i)$ , in which  $\phi$  is the appropriate non-convex loss function. We use a tilting argument, with non-convex losses. It is shown that, if we use a non-convex loss, sufficiently tilted, i.e. such that  $-\phi'(v)$  is small for all  $v \ll 0$ , then the outliers are eliminated successively. Hence, constant tilting or “trimming” is not sufficient for outlier removal. In tilting the loss function, we are effectively preserving as much fidelity to the data as possible, while redistributing emphasis to different observations. We propose a new *Arch Boosting* framework that implements the above tilting method and adjusts for optimality by a new search of the optimal weak hypothesis. Moreover, we show that the new framework avoids overfitting much similar to the AdaBoost. We propose a sufficient set of conditions needed for a loss function to allow for good properties of the ArchBoost. We show that not every non-convex function satisfies such conditions; an example is the sigmoid loss. However, we propose a family of loss functions that balances both the benefits of non-convexity and the empirical risk interpretation of boosting.

Properties of the boosting algorithms based on convex loss functions have been extensively studied (e.g. [Koltchinskii and Panchenko \(2002\)](#); [Zhang and Yu \(2005\)](#)). Comparatively little is known for the non-convex losses, as the existing techniques do not apply. We show that local convexity properties are sufficient for statistical consistency. Furthermore, even though the proposed loss function is shown to enjoy the aforementioned local convexity, it is largely unknown whether numerical algorithms can identify this local minimizer. Moreover, as our algorithm is not defined as a gradient descent algorithm, we require a new approach for the proof of numerical convergence. We develop a new sufficient optimality condition based on the “hardness condition” in the technical proofs. By hardness property, we mean the orthogonality of the new reweighted classifier and the class membership vector. Furthermore, we address the robustness and efficiency of the proposed method, with respect to the outliers. Although it is straightforward to provide such analysis for parametric linear models, computations for classification with the nonparametric boundaries are far more challenging. We provide a novel analysis, for which we propose a new finite sample breakdown point theory ([Hampel, 1968](#)) and show that the influence function ([Hampel, 1974](#)) is bounded for appropriate class of classification problems. To the best of our knowledge, this is the first result regarding the robustness properties of the boosting algorithms, with respect to the presence of outliers. Our analysis allows for both convex and non-convex loss function. We finalize the analysis with a proof of statistical consistency of the proposed method that includes many non-convex losses; we do so by exploring local curvatures of the loss.

In essence, this paper investigates the effects of non-convex losses on a variety of boosting algorithms in the presence of unknown contamination of the data. In particular, we focus on how to design a new boosting framework in order to improve the prediction accuracy of classification methods for the data with outliers. The rest of the paper is organized as follows. We present a new *Arch boosting* framework in [Section 2](#); it is designed for augmenting the boosting framework such that its corresponding classification algorithm is significantly more adaptable to outliers. [Section 3](#) outlines a new family of loss functions that explores non-convexity and present sufficient conditions for non-convex loss to be robust. We present theoretical analysis in [Section 4](#): the numerical and statistical convergence of the proposed algorithms are discussed in [4.1](#) and [4.3](#), respectively. Moreover, we show theoretical robust properties in the [Section 4.2](#) with the breakdown point discussed in [Section 4.2.2](#) and influence function in [Section 4.2.3](#). [Section 5](#) contains numerical experiments on a number of examples of the loss functions belonging to the introduced family:  $\gamma$ -robust loss, least squares, logistic, exponential and truncated exponential loss. We demonstrate both how to use these methods in practice and compare them to the alternative of applying the non-augmented AdaBoost algorithm to the noisy data. The subsection [5.1](#) varies the  $\gamma$  parameter and considers examples of a contaminated Gaussian distribution. All examples clearly illustrate that the methods

outlined in Section 3 are far more successful than the existing boosting methods. Section 5.2 deals with the more complex situation of Long/Servedio data for which we show that our Arch Boost method outperforms the robust boosting method of Freund (2009). We also discuss outlier detection examples in Section 5.3, and apply our methods to three real datasets in Section 5.4.

## 2 Arch Boost

We consider a binary classification problem, with  $\mathcal{X}$  denoting the domain of the  $d$ -dimensional variable  $X$ , and  $\mathcal{Y}$  denoting the class label set that equals  $Y \in \{-1, 1\}$ . We aim to estimate a function  $F(X) : \mathbb{R}^p \rightarrow \mathbb{R}$  and assume that the training data  $\{(X_i, Y_i), i = 1, \dots, n\}$  are i.i.d. copies of  $(X, Y) \in \mathcal{X} \times \mathcal{Y}$  with unknown distribution. The data consists of samples from the contaminated distribution that is composed of the true (uncontaminated) data and a fixed and unknown number of outliers in each of the classes.

From now on, we let  $\phi$  denote any differentiable loss function. Note that  $\phi$  does not need to be convex necessarily. For such  $\phi$ , we define the  $\phi$ -risk  $R_\phi$  and the empirical  $\phi$ -risk  $\widehat{R}_{\phi,n}$  as

$$R_\phi(F) = \mathbb{E}_{(X,Y) \sim \mathbb{P}}[\phi(YF(X))], \quad \widehat{R}_{\phi,n}(F) = \frac{1}{n} \sum_{i=1}^n \phi(Y_i F(X_i)), \quad (1)$$

where  $n$  is the sample size,  $\mathcal{S}_n = \{(X_i, Y_i)\}_{i=1}^n$  are i.i.d observations, and  $\mathbb{P}$  is the true probability distribution on  $\mathcal{X} \times \mathcal{Y}$ . For simplicity of notation, we write  $R_\phi(F) = \mathbb{E}[\phi(YF(X))]$ . Note that the observed samples  $(X_i, Y_i)$  can come from a contaminated distribution, i.e., they are i.i.d. samples from  $\delta\mathbb{P} + (1 - \delta)\Delta$  for small, but positive contamination  $\delta > 0$ .

We view boosting as a method that iteratively builds an additive model,

$$F_T(x) = \sum_{t=1}^T \alpha_t h_{j_t}(x)$$

where  $h_{j_t}$  belongs to a large (but we will assume finite VC-dimension) space of weak hypotheses, denoted with  $\mathcal{H}$ . The use of our framework in the presence of countably infinite features, also known as the task of feature induction, can easily be established. Next, we introduce the new framework of the boosting in the presence of the noise, which we call an *Arch Boosting* framework.

We design the Arch Boosting framework as a stage-wise, iterative, minimization of the  $\phi$ -risk (1). However, when  $\phi$  is non-convex this cannot be done by simply applying the well known Friedman's Gradient Boosting (Friedman, 2001). The explicit updates are usually unavailable and standard numerical methods like Newton-Raphson, which are suitable only for convex functions, are used. Instead, we constrain the stage-wise minimization of the  $\phi$ -risk to keep one additional important property of the boosting algorithms, namely the *hardness condition* of Freund and Schapire (1997). It is shown that hardness condition can easily distinguish between the outliers and the center of the data, when non-convex loss is used. Moreover, it allows for a fine tuning of the appropriate optimal hypothesis assignments such that the minimization of the  $\phi$ -risk is approximately kept. Therefore, hardness condition allows us to simultaneously escape the non-convexity in the minimization and to use non-convexity to separate the outliers from the inliers. This property, from iteration  $t$  to  $t + 1$ , is defined as

$$\mathbb{E}_{w_{t+1}}[Y h_t(X) | X = x] = 0, \quad (2)$$

where  $\mathbb{E}_w$  is defined as a weighted conditional expectation

$$\mathbb{E}_w[g(X, Y) | X = x] = \frac{\mathbb{E}[w(X, Y)g(X, Y) | X = x]}{\mathbb{E}[w(X, Y) | X = x]}. \quad (3)$$

Table 1: The list of commonly used loss functions and its corresponding  $F^*$

Classification Method	Population parameters	
	Loss function $\phi(v)$	Optimal Minimizer $F^*(x)$
Logistic	$\log(1 + e^{-v})$	$(\log \mathbb{P}(y = 1 x) - \log \mathbb{P}(y = -1 x))$
Exponential	$e^{-v}$	$\frac{1}{2} (\log \mathbb{P}(y = 1 x) - \log \mathbb{P}(y = -1 x))$
Least Squares	$(v - 1)^2$	$\mathbb{P}(y = 1 x) - \mathbb{P}(y = -1 x)$
Modified Least Squares	$[(1 - v)_+]^2$	$\mathbb{P}(y = 1 x) - \mathbb{P}(y = -1 x)$

The equation (2) can be explained as progressing from step  $t$  to  $t + 1$ ; the weights are updated from  $w_t$  to  $w_{t+1}$ , such that  $h_t(X)$  is orthogonal to  $Y$  with respect to the inner product defined on the reweighed data  $w_{t+1}$ . In a certain sense, the weights  $w_{t+1}$  are chosen as the most difficult for the weak hypothesis  $h_t$ . For the weight vector  $w_{t+1}$  the hypothesis  $h_t$  is not better than a random guess.

Recall that our goal is to find the optimal  $F^* \in \mathcal{F}$ , which minimizes the  $\phi$ -risk  $R_\phi(F)$  for a suitable class of measurable functions  $\mathcal{F}$ . In the binary case, each instance  $X_i$  is associated with a label  $Y_i \in \{-1, 1\}$  and the goal of learning is then to find a classifier  $F_T$ , such that the sign of  $F_T(X_i)$  is equal to  $Y_i$ . In order to minimize  $R_\phi(F) = \mathbb{E}[\phi(YF(X))]$ , we minimize it at every point  $x \in \mathcal{X}$  – that is, given any  $x \in \mathcal{X}$ , considering  $F(x)$  as a parameter and denoting  $\Phi(F(x)) = \mathbb{E}[\phi(YF(X))|X = x] = \mathbb{E}_Y[\phi(YF(x))]$ , the problem is to find

$$F^*(x) = \operatorname{argmin}_{F \in \mathcal{F}} \Phi(F(x)). \quad (4)$$

Table 1 contains a list of commonly used loss functions  $\phi$  and the corresponding optimal classifiers  $F^*$  when  $\mathcal{F} = \mathcal{M}$ , which is the family of all measurable functions.

Provided that  $\mathcal{F} = \mathcal{M}$  and for any  $x \in \mathcal{X}$ ,  $\Phi(F(x))$  has only one critical point at  $F^*(x)$  that is the global minimum, we can find  $F^*(x)$  by the first order optimality condition

$$\frac{d}{dF(x)} \Phi|_{F=F^*}(F(x)) = 0. \quad (5)$$

Expanding on the above first order optimality conditions, we obtain

$$\mathbb{E}[Y \phi'(YF^*(X))|X = x] = 0,$$

where  $\phi'$  is defined as the first order derivative  $\frac{d}{dv} \phi(v)$ . In classification problems, the input  $v$  of loss function  $\phi$  is  $v = YF(X)$  – that is, the margin of a classifier  $F$  applied to a data point  $(X, Y)$ . Rewriting the expectation in terms of the class probabilities, we obtain the following representation of the first order optimality conditions

$$\frac{\phi'(-F^*(x))}{\phi'(F^*(x))} = \frac{\mathbb{P}(Y = 1|X = x)}{\mathbb{P}(Y = -1|X = x)}. \quad (6)$$

From now on, without confusion, we write  $\mathbb{P}(Y = 1|x)$  instead of  $\mathbb{P}(Y = 1|X = x)$ . We aim to mimic equation above in each of the iteration steps of the proposed framework. In more details, at iteration  $t$ , with the current estimate  $F_{t-1}(x) = h_1(x) + \dots + h_{t-1}(x)$  at hand, we wish to find a new weak hypothesis  $h_t \in \mathcal{H}$ , such that  $F_t(x) = F_{t-1}(x) + h_t(x)$  with  $h_t(x)$  that solving the following equation

$$\frac{\phi'(-F_{t-1}(x) - h_t(x))}{\phi'(F_{t-1}(x) + h_t(x))} = \frac{\mathbb{P}(Y = 1|x)}{\mathbb{P}(Y = -1|x)}. \quad (7)$$

If we can always accurately find such a weak hypothesis (i.e.  $\mathcal{H}$  is rich enough and we know the true distribution), then this process will terminate in just one step. However, complications arise from two aspects. First, we only have a restricted family  $\mathcal{H}$  of weak hypotheses. Therefore, at each iteration, an approximated weak hypothesis, which is the closest approximation in  $\mathcal{H}$ , will be found. Secondly, the equation (7) alone cannot be efficiently utilized since  $\mathbb{P}(Y = 1|x)$ , which is the ultimate goal of classification problems, is unknown to us. However, we propose to solve approximated equation, where we replace  $\mathbb{P}(Y = 1|x)$  with the weighted conditional probability  $\mathbb{P}_{w_t}(Y = 1|x)$  at each iteration  $t$  of the proposed algorithm. Hence,  $\mathbb{P}_{w_t}(Y = 1|x) = \mathbb{E}_{w_t}[\mathbb{1}_{[Y=1]}|X = x]$  with the weighted conditional expectation defined in (3).

In more details, at each iteration  $t$ , for a given  $\mathbb{P}_{w_t}(Y = 1|x)$ , we find  $h_t(x)$  such that it solves the estimating equation

$$\frac{\phi'(-F_{t-1}(x) - h_t(x))}{\phi'(F_{t-1}(x) + h_t(x))} = \mathbb{C}_{t-1}(x) \frac{\mathbb{P}_{w_t}(Y = 1|x)}{\mathbb{P}_{w_t}(Y = -1|x)}, \quad (8)$$

for a suitable function  $\mathbb{C}_{t-1}(x)$  to be specified later. In order to do this, we will first discuss how to define the weights  $w_t$  at each iteration  $t$  and postpone the method of finding  $h_t$  for later.

Since  $h_t$  is only an approximation to the optimal increment at step  $t$ , instead of adding itself to  $F_{t-1}$  we multiply it by a constant  $\alpha_t$  and search for the best constant  $\alpha_t$ . The best  $\alpha_t$  should be such that the updating classifier  $\alpha_t h_t(x)$  approaches the optimal one, defined in equation (6). Hence, we define the optimal  $\alpha_t$  as the solution to the following optimization problem

$$\alpha_t = \operatorname{argmin}_{\alpha \in \mathbb{R}} \mathbb{E}[\phi(Y(F_{t-1}(X) + \alpha h_t(X)))]. \quad (9)$$

Similarly to the AdaBoost, at each iteration  $t$ , the data will be reweighed according to the weights  $w_t$ . We explore relation (9) further to find the optimal weights  $w_t$ . For that purpose, we observe that for a given  $F_{t-1}$  (from a previous iteration) and  $h_t$  (solving (8)) the optimal  $\alpha_t$  then satisfies

$$\mathbb{E}\left[-\phi'(YF_{t-1}(X) + Y\alpha_t h_t(X)) \cdot \alpha_t Y h_t(X)\right] = 0. \quad (10)$$

Then, we recall the hardness condition of the AdaBoost algorithm: the weights  $w_t$  should be updated such that the weighted misclassification error of the most recent weak hypothesis is about 50%. According to (2), this can be achieved by defining the weights  $w_{t+1}(X, Y)$  such that

$$\mathbb{E}[w_{t+1}(X, Y) \cdot Y \alpha_t h_t(X)] = 0. \quad (11)$$

Now, in the Arch Boost framework the most recent weak hypothesis is  $\alpha_t h_t$ . Hence, by contrasting the last two equations, (10) and (11), we define the weights to be

$$w_{t+1}(X, Y) = -\phi'(YF_{t-1}(X) + Y\alpha_t h_t(X)) = -\phi'(YF_t(X))$$

such that both the hardness condition and the optimality of the updating hypothesis are satisfied.

Having defined the weight updating rule, we go back to equation (7) to define the optimal weak hypothesis  $h_t(x)$ . We do so by finding a relationship between  $\mathbb{P}_{w_t}(Y = -1|x)$  and  $\mathbb{P}(Y = -1|x)$  and defining the function  $\mathbb{C}_{t-1}$  of (8). We multiply (7) with  $\frac{\phi'(F_{t-1}(x))}{\phi'(-F_{t-1}(x))}$  on both sides to obtain

$$\frac{\phi'(F_{t-1}(x))\phi'(-F_{t-1}(x) - h_t(x))}{\phi'(-F_{t-1}(x))\phi'(F_{t-1}(x) + h_t(x))} = \frac{\phi'(F_{t-1}(x))\mathbb{P}(Y = 1|x)}{\phi'(-F_{t-1}(x))\mathbb{P}(Y = -1|x)}. \quad (12)$$

Then, it is easy to observe that

$$\begin{aligned}
\phi'(F_{t-1}(x))\mathbb{P}(Y = 1|x) &= \mathbb{E}[\mathbb{1}_{[Y=1]}\phi'(YF_{t-1}(X))|X = x] \\
&= -\mathbb{E}[\mathbb{1}_{[Y=1]}w_t(X, Y)|X = x] = -\mathbb{E}_{w_t}[\mathbb{1}_{[Y=1]}|X = x] \\
&= -\mathbb{P}_{w_t}(Y = 1|x).
\end{aligned}$$

Hence, we have

$$\frac{\phi'(F_{t-1}(x))\phi'(-F_{t-1}(x) - h_t(x))}{\phi'(-F_{t-1}(x))\phi'(F_{t-1}(x) + h_t(x))} = \frac{\mathbb{P}_{w_t}(Y = 1|x)}{\mathbb{P}_{w_t}(Y = -1|x)}, \quad (13)$$

that is, equation (8) is true for

$$\mathbb{C}_{t-1}(x) = \phi'(-F_{t-1}(x))/\phi'(F_{t-1}(x)).$$

Observe that the right hand side of (13) can be estimated by a weak learner at each iteration  $t$ . For reference, we list the weak hypothesis  $h(x)$  for several commonly used loss functions in the Table 2.

Table 2: The list of commonly used loss functions and their weak hypotheses  $h$

Classification Method	Population parameters	
	Loss function $\phi(v)$	Optimal weak hypotheses $h(x)$
Logistic	$\log(1 + e^{-v})$	$\log \mathbb{P}_w(Y = 1 x) - \log \mathbb{P}_w(Y = -1 x)$
Exponential	$e^{-v}$	$\frac{1}{2}(\log \mathbb{P}_w(Y = 1 x) - \log \mathbb{P}_w(Y = -1 x))$
Least Squares	$(v - 1)^2$	$C(1 - F(x))(1 + F(x))/(CF(x) + 1)$
Modified Least Squares	$[(1 - v)_+]^2$	$C(1 - F(x))(1 + F(x))/(CF(x) + 1)$

\*  $C = \mathbb{P}_w(Y = 1|x) - \mathbb{P}_w(Y = -1|x)$

Note that the weak hypotheses of the least squares loss and modified least squares loss depend on the current estimate  $F(x)$  and the weighted conditional probability  $\mathbb{P}_{w_t}(Y = 1|x)$ , which is different from that of Gradient boosting (Friedman, 2001). Lastly, we summarize the above procedure in the following Algorithm 1, which we call *Arch Boost*. The assumption that  $\Phi(F(x))$  has only one critical point does not require the loss function  $\phi$  to be convex and hence the Arch Boost algorithm can be applied to many non-convex loss functions. We illustrate this in Section 3 by proposing a family of robust boosting algorithms based on the Arch Boosting framework.

In the step (b) of the Algorithm 1, any classifier can be used; one example is a decision tree for which case  $\mathbb{P}_{w_t}(Y = 1|x)$  is the proportion of the training samples with label 1 in the terminal node, where  $x$  ends up. For instance, in terminal region  $R$ , let  $\mathcal{I}_+$  stands for the index set of data points with positive label, then  $\mathbb{P}_{w_t}(Y = 1|x) = \sum_{i \in \mathcal{I}_+} w_t(x_i, y_i) / \sum_{j \in R} w_t(x_j, y_j)$ .

Note that in the step (c), whenever  $\mathbb{P}_{w_t}(Y = 1|x)\mathbb{P}_{w_t}(Y = -1|x) = 0$ , then  $h_t(x)$  can be  $\infty$  or  $-\infty$ . In order to deal with this problem, one method is to only update the points with  $-\infty < F_t(x) < \infty$ . If  $F_t(x)$  becomes infinity after some step  $t$ , then we keep it to be infinity and do not update further. Finally, we just set  $sign(\infty) = 1$  and  $sign(-\infty) = -1$ . Another method is simply applying a map  $v \mapsto av + \frac{1-a}{2}$  to the class probability estimations where  $a < 1$  is a constant very close to 1 (e.g.  $a = 0.9999$ ).

### 3 Robust boosting algorithms

In this section, we propose a new class of boosting algorithms especially designed to be resilient to the presence of extraneous noise in the data. Whether this noise comes in the forms of mislabeled



---

**Algorithm 1** Arch Boost ( $\phi$ )

---

**procedure** ARCH BOOST( $\phi$ )  
Given:  $(x_1, y_1), \dots, (x_n, y_n)$   
3: Initialize the vector of weights  $w_0$ , e.g.  $w_0(x_i, y_i) = 1/n$   
**for**  $t = 1, \dots, T$  **do**  
    (a) Normalize the weight by assigning  $w_t \leftarrow w_t / \sum_i w_t(x_i, y_i)$   
6:     (b) Fit the classifier to obtain a class probability estimate  $\mathbb{P}_{w_t}(Y = 1|x) \in [0, 1]$   
        using current weights  $w_t$  on the training data.  
    (c) Set  $h_t(x)$  to be the solution of (13).  
9:     (d) Find  $\alpha_t$  by solving (9).  
    (e) Set  $F_t(x) = F_{t-1}(x) + \alpha_t h_t(x)$   
    (f) Update the weights  $w_t \leftarrow -\phi'(yF_t(x))$   
12: **end for**  
    Output the classifier:  
$$\text{sign}(F_T(x))$$
  
**end procedure**

---

data points, additional variance within class data, or malicious data, the proposed algorithm adapts to it with great success. We begin the section by proposing a new loss function, which we have named two-robust loss function. Furthermore, we characterize regularity conditions that a loss function needs to satisfy, in order to be appropriate for Arch Boosting framework. Finally, we provide a family of loss functions that is not convex and that satisfies the newly defined regularity conditions and the corresponding family of robust boosting algorithms called ARB- $\gamma$ .

### 3.1 A non-convex loss

An approach of non-convex functions has been recognized as successful in the existing literature. Both Freund (2001) and Freund (2009) have utilized it to propose boosting methods that are resilient to the presence of outliers in the data. However, both loss functions require a number of tuning parameters. The behavior of the algorithm is hindered by the optimal choice of such parameters. Moreover, such optimal choices are data dependent and not isolated as universal by either approach.

Driven by the need to propose an alternative loss function that can be easily used in many situations, we propose the following non-convex loss function:

$$\phi(v) = 4(1 + e^v)^{-2}. \tag{14}$$

The proposed loss function (14) has optimal  $F^*$ , the weak hypothesis, and the weight updating rule as presented in Table 3. We call the function (14) *two-robust loss* and it belongs to a family of loss functions that will be discussed in Section 3.3. By adopting the Arch Boosting framework of Algorithm 1 we illustrate the nice downweighting property of the new weight updating rule with this new loss function. In Figure 1, we plot the two-robust loss (14), together with exponential and logistic regression losses and the corresponding weight updating rules.

We observe that the two-robust loss is non-convex and bounded from above when  $v \rightarrow -\infty$ . Therefore, the outliers, even if large in size, will only have bounded influence on the classification. Moreover, by investigating the weight updating rule, we observe that the more misclassified the data point is, the smaller the weight updating function will be. The algorithm, in fact, abandons the data points that are repeatedly misclassified and are far from the Bayes boundary. This phenomenon disappears when one uses exponential loss or logistic loss and, in fact, any commonly used convex

Table 3: Properties of the two-robust loss

Arch 2-boosting	Optimal parameters			
	Loss function $\phi(v)$	Optimal minimizer $F^*(x)$	Weak hypothesis $h(x)$	Weight vector $w(x, y)$
Arch	$4(1 + e^v)^{-2}$	$\log \frac{\mathbb{P}(Y=1 x)}{\mathbb{P}(Y=-1 x)}$	$\log \frac{\mathbb{P}_{w_t}(Y=1 x)}{\mathbb{P}_{w_t}(Y=-1 x)}$	$\frac{e^{yF_t(x)}}{(1+e^{yF_t(x)})^3}$

loss function. To the best of our knowledge, no existing loss function satisfies this nice property without requiring a-priori fixed tuning parameters. By plugging the two-robust loss into Arch Boost, we obtain new boosting algorithm that we named *Adaptive Robust Boost-2*, denoted with ARB-2 from hereon.

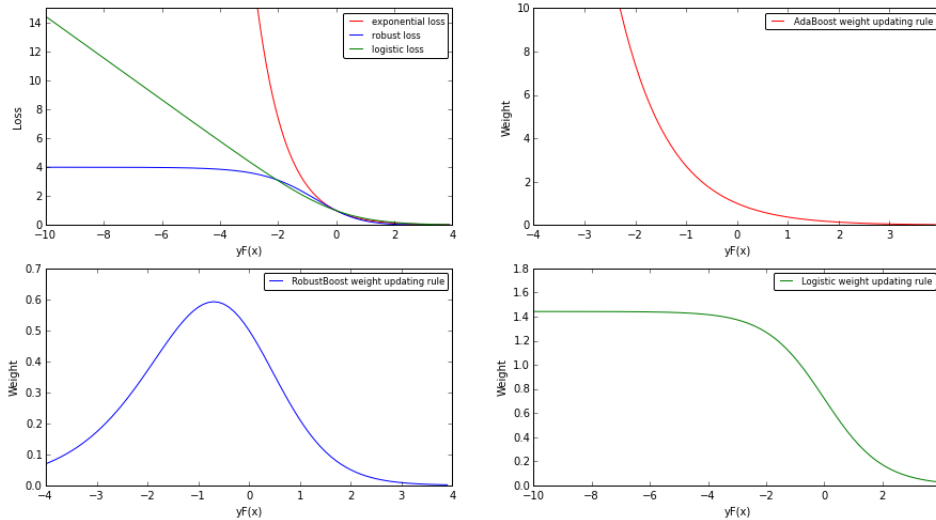


Figure 1: Comparison of the two-robust loss, the exponential loss and the logistic loss functions (left-top) together with the corresponding weight updating functions.

Next, we consider a simple two dimensional binary classification example consisting of two mislabeled points near the classification boundary. The example is artificially created to illustrate the point of an adaptive classifier, when there are outliers in the data. We evenly put 121 points in the region  $[0, 3] \times [0, 3]$ , that is,  $\mathcal{X} = \{\mathbf{x}_{i,j} : i = 0, \dots, 10, j = 0, \dots, 10\} \subset \mathbb{R}^2$ , where  $\mathbf{x}_{i,j} = (0.3i, 0.3j)$ . For each  $\mathbf{x}_{i,j}$ , the corresponding  $y_{i,j} = 1$  if  $i + j \leq 60$  and  $y_{i,j} = -1$  otherwise. Then we flip the labels of  $(\mathbf{x}_{2,4}, y_{2,4})$  and  $(\mathbf{x}_{3,4}, y_{3,4})$ . The step sizes,  $\alpha_t$ , for both algorithms are set to be constant 0.5, and the number of iterations is set as 400. We show the difference between the Real AdaBoost and the ARB-2 in Figure 2. It can be seen that the decision boundary of AdaBoost are influenced by the two blue “abnormal” data points. But for ARB-2, the decision boundary stays the same as if there were no such abnormal points, hence it adapts to the outliers.

### 3.2 Loss functions for Arch Boost

In the previous section, we have seen an example of a robust boosting algorithm. Nevertheless, there are plenty of other non-convex functions which leads to the question: can we use any of them as a loss function for the Arch Boost? The answer is no, if we do not impose auxiliary conditions. In this section, we discuss what kinds of conditions need to be imposed so that a non-convex function becomes a suitable loss function for the Arch Boost framework 1.




Figure 2: The classification boundary of the Real AdaBoost (left) and of the ARB-2 (right).

Recall that in a binary classification problem, we want to minimize  $\Phi(F(x))$ , to find

$$\min_{v \in \mathbb{R}} \left[ \mathbb{P}(Y = 1|x)\phi(v) + \mathbb{P}(Y = -1|x)\phi(-v) \right]. \quad (15)$$

Since  $\mathbb{P}(Y = 1|x) + \mathbb{P}(Y = -1|x) = 1$ , the equation above becomes a convex combination of  $\phi(v)$  and  $\phi(-v)$ . Therefore, one necessary condition on loss functions is that (15) has a unique optimal solution in  $\mathbb{R}$ . Observe that if  $\mathcal{F}$  is a class of all measurable functions, then  $F(x)$  can take any real value for every  $x \in \mathcal{X}$ . This condition is not equivalent to the convexity of the loss function but rather to the local convexity around the true parameter of interest,  $F^*(x)$ . In the next section we present a family of non-convex loss functions that possess this property. Hence, we present the set of regularity conditions in the Definition below.

**Definition 1.** A function  $\phi$  is an **Arch boosting loss function** if it is differentiable and

- (i)  $\phi(v) \geq 0$  for all  $v \in \mathbb{R}$  and  $\inf_{v \in \mathbb{R}} \phi(v) = 0$ ;
- (ii) for any  $0 < \alpha < 1$ ,  $\alpha\phi(v) + (1 - \alpha)\phi(-v)$  has only one critical point  $v^*$  which is the global minimum;
- (iii) for any  $0 \leq \alpha \leq 1$  and  $\alpha \neq \frac{1}{2}$ ,  $\inf\{\alpha\phi(v) + (1 - \alpha)\phi(-v) : v(2\alpha - 1) \leq 0\} > \inf\{\alpha\phi(v) + (1 - \alpha)\phi(-v) : v \in \mathbb{R}\}$ .

Conditions (i) and (iii) together imply that  $\phi$  is an upper bound of the 0-1 loss up to a constant scaling. Condition (iii) is a “classification calibration” (Bartlett et al., 2006) and is considered the weakest possible condition imposed on  $\phi$  for which a measurable function  $F(x)$  which minimizes the  $\phi$ -risk  $R_\phi(F)$  will also have the risk close to the minimal one; in other words, close to that of the optimal  $F^*(c)$  that creates a Bayes boundary. If the function  $\phi$  is convex, then condition (iii) is satisfied as long as  $\phi$  is differentiable and  $\phi'(0) < 0$ . However, when considering non-convex losses, the set of regularity conditions doesn’t exist in the current literature. The above framework includes non-convex loss functions  $\phi$  and differs from the existing literature on convex losses in that it includes an additional condition (ii) (see Section 4).

**Lemma 1.** A positive function  $\phi$  that is continuously differentiable, convex and such that  $\phi'$  is not equal to a constant satisfies Condition (ii).

By Lemma 1, we know that the logistic, the exponential, the least square and the modified least square loss are all Arch boosting loss functions. Differentiability of the loss function is a technical

condition and is not crucial for the proposed framework. The hinge loss  $\phi_{hinge}(v) = (1 - v)_+$  is not differentiable but can be shown to satisfy Conditions (i)-(iii). However, if we plug in the hinge loss to the equation (6), we cannot easily find the optimal solution  $F^*(x)$ . However, not every differentiable and non-convex loss function satisfies regularity conditions above.

**Remark 1.** *The sigmoid loss  $\phi_{sig}(v) = (1 + e^v)^{-1}$  is differentiable and satisfies Condition (iii), but it does not satisfy Condition (ii) and hence is not an Arch Boosting loss function.*

Finally, we provide a more general characterization for the loss functions, including not necessarily convex loss functions, which satisfy Condition (ii).

**Lemma 2.** *Let  $\phi$  be a positive, continuously differentiable loss function such that  $\phi'(v) \neq 0$  for all  $v \in \mathbb{R}$ . Let  $g : (0, \infty) \rightarrow (0, 1)$ , be defined as  $g(v) := \frac{\phi'(v)}{\phi'(-v)}$ . Then  $\phi$  satisfies Condition (ii) as long as function  $g$  is a decreasing, bijection.*

### 3.3 A family of non-convex functions and ARB- $\gamma$ algorithms

Recall that the optimal classifier  $F^*(x)$  satisfies equation (6). We observe that the right hand side of this equation does not depend on the loss function  $\phi$  and can take values in the positive real line,  $\mathbb{R}_+$ . Hence, we can parameterize it with any real-valued function whose range is  $\mathbb{R}_+$ , as follows

$$\frac{\phi'(-v)}{\phi'(v)} = g(v) \tag{16}$$

for any surjective, decreasing function  $g : \mathbb{R} \rightarrow \mathbb{R}_+$ . The classical motivation for reparametrization (McCullagh and Nelder, 1989) – often called link functions – is that often one uses a parametric representation that has a natural scale matching the desired one. We choose to use the function  $g(v) = e^{(\gamma-1)v}$  for any constant  $\gamma > 1$ , which is a surjection. This parametrization is not unique but it admits a solution to the following differential equation

$$\frac{\phi'(-v)}{\phi'(v)} = e^{(\gamma-1)v}. \tag{17}$$

Solving it for  $\phi$  (exact steps are presented in the Appendix A), we obtain a family of non-convex functions

$$\phi_{a,\gamma}(v) = \frac{2^\gamma}{(1 + e^{av})^\gamma}, \quad a > 0. \tag{18}$$

Observe that parameters  $a$  and  $\gamma$  are not tuning parameters, but rather an index set of a family of non-convex losses much like Huber and Tukey’s biweight losses are. Note that when  $\gamma > 1$ , the right hand side of (17) can take all the real values and is monotonically increasing and consequently the corresponding loss  $\phi$  will always give a unique solution to (6). We name each element in the family (18) as  $\gamma$ -robust loss function. Later, we will show that the positive parameter  $a$  is irrelevant in our algorithms, so we can fix it to be 1. Note that for  $a = 1$  and  $\gamma = 2$ , we obtain the non-convex loss function (14). Moreover, each loss function  $\phi_{a,\gamma}$  is a bounded function with upper bound equal to  $2^\gamma$ . Therefore, the effects of the outlier will be necessarily bounded. Moreover, the weight updating rule will also downweight the largely misclassified data points. We plot  $\phi_{1,\gamma}$  with different  $\gamma$  values and the corresponding normalized weight updating rules in the Figure 3. From Figure 3, we can see that for different  $\gamma$ , the peak point of each of the weight updating rules shifts to the left when  $\gamma$  increases. When  $\gamma = 1$ , the weight updating curve is symmetric about  $yF(x) = 0$  and in fact, this is equivalent to the Sigmoid loss function  $\phi(v) = 1 - \tanh(\lambda v)$  when  $\lambda = \frac{1}{2}$  (Mason




Figure 3: Comparison of  $\phi_{1,\gamma}$  with different  $\gamma$  values and the corresponding normalized weight updating rules.

et al., 1999). Moreover, for  $a = 2$  and  $\gamma = 2$  the loss  $\phi_{a,\gamma}$  is equivalent to the Savage loss function  $\phi(v) = (1 + e^{2v})^{-2}$  of Mesnadi-Shirazi and Vasconcelos (2009), in which they used the probability elicitation technique from Savage (1971) to design new loss functions.

**Lemma 3.** For all  $a > 0, \gamma > 1$ ,  $\phi_{a,\gamma}$  is an Arch Boosting loss function.

Lemma 3 allows us to plug  $\phi_{a,\gamma}$  with  $a > 0, \gamma > 1$  into the Arch Boost framework and obtain a family of robust boosting algorithms which we name *Adaptive Robust Boost- $\gamma$  (ARB- $\gamma$ )*. The details are presented in the Algorithm 2 with computations relegated to the Appendix A. Note that

---

**Algorithm 2** ARB- $\gamma$  ( $\gamma \geq 1$ )

---

**procedure** ARB- $\gamma$

Given:  $(x_1, y_1), \dots, (x_n, y_n)$

3: Initialize the weight vector  $w_0$ , e.g.  $w_0(x_i, y_i) = 1/n$

**for**  $t = 1, \dots, T$  **do**

(a) Normalize the weight vector  $w_t \leftarrow w_t / \sum_i w_t(x_i, y_i)$

6: (b) Compute the classifier to obtain a class probability estimate  $\mathbb{P}_{w_t}(Y = 1|x) \in [0, 1]$ , using weights  $w_t$  on the training data.

(c) Set  $h_t(x) \leftarrow \log \frac{\mathbb{P}_{w_t}(Y=1|x)}{\mathbb{P}_{w_t}(Y=-1|x)} \in \bar{\mathbb{R}}$ .

9: (d) Find  $\alpha_t$  by solving (9).

(e) Set  $F_t(x) \leftarrow F_{t-1}(x) + \alpha_t h_t(x)$

(f) Set  $w_{t+1} \leftarrow e^{yF_t(x)}(1 + e^{yF_t(x)})^{-\gamma-1}$

12: **end for**

Output the classifier:

$$\text{sign}(F_T(x)) \tag{19}$$

**end procedure**

---

the form of the weak hypothesis in Algorithm 2 is  $h_t(x) = \log \frac{\mathbb{P}_{w_t}(Y=1|x)}{\mathbb{P}_{w_t}(Y=-1|x)}$  instead of the solution  $\frac{1}{\gamma-1} \log \frac{\mathbb{P}_{w_t}(Y=1|x)}{\mathbb{P}_{w_t}(Y=-1|x)}$  to equation (8), because we absorb the constant  $\frac{1}{\gamma-1}$  into the step size  $\alpha_t$  at each iteration  $t$ .

## 4 Theoretical Considerations

Despite the substantial body of existing work on Gradient Boosting classifiers and AdaBoost in particular (e.g. Bartlett et al. (2006); Bartlett and Traskin (2007); Freund (1995); Friedman, et al. (2000); Schapire (1999); Schapire et al. (1998); Zhang and Yu (2005); Breiman (2004); Koltchinskii and Panchenko (2002)), research on robust boosting classifiers has mostly been limited to methodological proposals with little supporting theory (e.g., Gentile and Littlestone (1999); Kearns and Li (1993); Littlestone (1991)). However, whereas the loss functions studied in those papers are finite-sample versions of globally convex functions, many important robust classifiers, such as those arising in Freund (2009) and the proposed ARB- $\gamma$ , only possess convex curvature over local regions even at the population level. In this paper, we present new theoretical results, based only on local curvatures, which may be used to establish statistical and optimization properties of the proposed Arch boosting algorithms, with highly non-convex loss functions.

### 4.1 Numerical convergence of Arch Boost algorithm

We show that the empirical risk will always decrease and that the Arch Boost can be viewed as the step-wise iterative minimization of the empirical risk  $\hat{R}_{\phi,n}$ . Similar results can be found in Koltchinskii and Panchenko (2002) and Zhang and Yu (2005). The main difference is that the authors use the gradient descent rule in the first or an approximate minimization in the second paper, while here we use the hardness condition to select the optimal weak hypothesis  $h$ . We use  $\{w(X_i, Y_i)\}_{i=1}^n$  to denote the weights on the data such that  $\sum_{i=1}^n w(X_i, Y_i) = 1$ . Recall that for any classifier  $h$  and any data point  $(X, Y)$ , the term  $Yh(X)$  always stands for the margin. For any weak hypothesis  $h_t$ , we denote the *expected margin* as  $\mu(h, w) = \mathbb{E}_w[Yh(X)]$ , and the *empirical margin* as  $\hat{\mu}(h, w) = \sum_{i=1}^n w(X_i, Y_i)Y_i h(X_i)$ . To introduce the notation used in the theorem, for any family of weak classifiers  $\mathcal{H}$ , we denote

$$\mathcal{F}^T = \left\{ F : F = \sum_{t=1}^T \alpha_t h_t, \alpha_t \in \mathbb{R}, h_t \in \mathcal{H} \right\}, \quad (20)$$

a set of  $T$ -combinations ( $T \in \mathbb{N}$ ) of functions in  $\mathcal{H}$ . Let  $\{\bar{f}_t\}$  be a sequence of functions with empirical risk converging to  $R_{\phi,n}^*$ , defined as  $R_{\phi,n}^* = \inf_{F \in \cup_{T=1}^{\infty} \mathcal{F}^T} \hat{R}_{\phi,n}(F)$ . Then,  $\bar{f}_t \in \mathcal{H}_f \subset \cup_{T=1}^{\infty} \mathcal{F}^T$  can be represented as  $\sum_{h \in H_f} \alpha^h h$ . In this respect, we define its  $l_1$  norm as  $|H_f|^{-1} \sum_{h \in H_f} |\alpha^h|$ .

**Theorem 1.** *Assume  $\phi$  is an Arch boosting loss function. Furthermore, assume that the weak learner is able to provide disjoint regions on the domain  $\mathcal{X}$  at each iteration  $t$  (e.g. decision tree). We apply the Arch Boost ( $\phi$ ) algorithm to a sample  $\mathcal{S}_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$  for  $T$  iterations.*

- (i) *If at each iteration  $t$  the weak hypotheses  $h_t \in \mathcal{H}$  satisfies  $\hat{\mu}(h_t, w_t) > 0$ , then  $\hat{R}_{\phi,n}(F_T)$  will converge in  $\mathbb{R}$  as  $T \rightarrow \infty$ .*
- (ii) *At each iteration  $t$ , let  $\{R_t^j\}_{j=1}^J$  be the set of disjoint regions on  $\mathcal{X}$  returned by the weak hypothesis  $h_t$ , and  $p_t^j := \mathbb{P}_{w_t}(Y = 1|x \in R_t^j)$  be the class probability estimation in that region. If there exists a strictly increasing function  $\theta : [0, 1] \rightarrow \mathbb{R}$  with  $\theta(\frac{1}{2}) = 0$ , and a positive constant  $K > 0$  such that  $h_t$  satisfies the representation*

$$h_t(x) = K\theta(p_t^j)$$

*for all  $x \in R_t^j$ , then  $\hat{\mu}(h_t, w_t) > 0$ .*

- (iii) *Using any Arch boosting loss function  $\phi$  for the Arch Boost Algorithm 1, at any iteration  $t$ , there exists such a function  $\theta$  that satisfies (ii).*

(iv) Let  $\{\bar{f}_t\}$  be a sequence of functions with empirical risk converging to  $R_{\phi,n}^*$  and such that

$$\|\bar{f}_t - F_t\|_1 = o(\log t), \quad \|\bar{f}_t - F_t\|_2^2 \leq \frac{\|\bar{f}_t - F_t\|_1^2}{t^{c_t}}$$

where  $c_t \in (0, 1)$  and  $c_t \rightarrow 0$  as  $t \rightarrow \infty$ . Furthermore, assume  $\phi$  is Lipschitz differentiable,  $\theta$  is bounded and  $\hat{\mu}(h_t, w_t) \rightarrow 0$  as  $t \rightarrow \infty$ . Then, if a sequence of step sizes  $\alpha_t$  is such that

$$\sum_{t=1}^{\infty} \alpha_t = \infty, \quad \sum_{t=1}^{\infty} \alpha_t^2 < \infty, \quad \sum_{t=1}^{\infty} \frac{\alpha_{t+1} \xi_t \log t}{t^{c_t}} < \infty,$$

for a sequence of positive numbers  $\xi_t = o(1)$ , we have  $\hat{R}_{\phi,n}(F_T) \rightarrow R_{\phi,n}^*$  as  $T \rightarrow \infty$ .

Theorem 1 suggests that for any ARB- $\gamma$  algorithm, the weak hypothesis  $h_t$  at each iteration  $t$  has a corresponding function  $\theta$  that satisfies (ii). Therefore,  $h_t$  satisfies  $\hat{\mu}(h_t, w_t) > 0$  and with it, by (i), the ARB- $\gamma$  algorithm converges when the number of iterations increases. The conditions in part (iv) are somewhat different compared to the equivalent one obtained for the gradient boosting with convex losses (Zhang and Yu, 2005). The reference sequence  $\bar{f}_t$  needs to be in a local neighborhood of  $F_T$  in the sense that  $\bar{f}_t - F_t$  cannot blow up too rapidly. Moreover, the size of  $\mathcal{H} \supset \mathcal{H}_t \ni \bar{f}_t$  cannot be smaller than  $t^{c_t}$  implying that the size of  $\mathcal{H}$  needs to converge to  $\infty$  faster than a polynomial of  $T$ . Additionally, the choice of  $\alpha_t$  depends on  $c_t$  and  $\xi_t$ . The classical conditions that are guarding against infinitely small step sizes are now supplemented with an additional constraint  $\sum_{t=1}^{\infty} \frac{\alpha_{t+1} \xi_t \log t}{t^{c_t}} < \infty$ . For example, if  $\xi_t = O(\frac{1}{t})$ , then we can choose  $\alpha_t = O(\frac{1}{t^c})$  where  $c \in (\frac{1}{2}, 1)$  and  $c_t$  can converge to 0 at any speed. However, if  $\xi_t = O(\frac{1}{\log t})$ , then we need  $c_t \rightarrow 0$  slowly (e.g.  $O(\frac{1}{\log \log t})$ ) and  $\alpha_t$  can be chosen as  $O(\frac{1}{t})$ . The additional constraint in the step size choice acts as a penalty on allowing non-convex loss functions. However, unlike existing results Theorem 1 does not require any additional algorithmic tuning parameters (see Theorem 3.1 of Zhang and Yu (2005) and choices of  $\varepsilon_t, \Lambda_t$ ). Results in Bartlett and Traskin (2007) (e.g., Theorem 6) provide similar bounds under an assumption of an unbounded step size of the boosting algorithm and assume a positive lower bound on the hessian of the empirical risk (a condition strictly violated for non-convex losses).

We provide examples of the function  $\theta$  in the Table 4. Note that this distribution  $\mathbb{P}$  may potentially be a contaminated distribution – the proof is not affected by the contamination. Furthermore,

Table 4:  $\theta$  functions for several loss functions

Classification Method	Population parameters	
	Loss function $\phi(v)$	$\theta$ functions
Logistic	$\log(1 + e^{-v})$	$\log(t) - \log(1 - t)$
Exponential	$e^{-v}$	$\log(t) - \log(1 - t)$
Least Squares	$(v - 1)^2$	$\frac{(2t-1)(1-F^2)}{(2t-1)F+1}, F \in [-1, 1]$
Modified Least Squares	$[(1 - v)_+]^2$	$\frac{(2t-1)(1-F^2)}{(2t-1)F+1}, F \in [-1, 1]$
$\gamma$ -robust ( $\gamma > 1$ )	$2^\gamma / (1 + e^v)^\gamma$	$\log(t) - \log(1 - t)$

we remark that the result of Theorem 1 (i) to (iii) requires very weak conditions. Namely, the approximate minimization step (9) can be inexact (by contrast, see Theorem 6 of Bartlett and Traskin (2007)). The weak hypothesis  $h_t$  at each iteration is obtained by preserving the “hardness” property of the AdaBoost, and this method is applicable to non-convex functions. For a certain loss function  $\phi$ ,  $h_t$  may not apriori point to the gradient descend direction of the empirical  $\phi$ -risk. Hence, we

cannot use numerical methods like Newton’s method as needed in the gradient descent, but provide a novel way to find a weak hypothesis that is also suitable for non-convex loss functions. Then, we show that the direction of such weak hypothesis will indeed be a descending direction for the non-convex loss, that is, we guarantee that  $\widehat{R}_{\phi,n}(F + \alpha h) \leq \widehat{R}_{\phi,n}(F)$  with appropriate step size  $\alpha$ , where  $F$  is the current estimate.

## 4.2 Robustness

In Section 2, we gave an informal explanation why non-convex losses lead to a more robust algorithm. The weight updating rule or the first derivative of the loss functions plays an important role in their robustness. Unlike convex functions, many non-convex functions can have a diminishing first derivative  $\phi'(v)$  when  $v$  tends to both infinity and negative infinity. In this section, we quantify the robustness and justify the robustness of Arch Boosting Algorithms 2 through the point of view of the finite sample Breakdown point, as well as that of the Influence Functions, the population measure of robustness.

### 4.2.1 An invex function view of robustness

In this section, we will use invex function properties to show why our non-convex functions leads to more robust algorithms. We first recall some definitions (Ben-Israel and Mond, 1986).

**Definition 2** (Invex function). *Assume  $X \subset \mathbb{R}^n$  is an open set. The differentiable function  $f : X \rightarrow \mathbb{R}$  is invex if there exists a vector function  $\eta : X \times X \rightarrow \mathbb{R}^n$  such that*

$$f(x) - f(y) \geq \eta(x, y)^T \nabla f(y), \quad \forall x, y \in X. \quad (21)$$

It is well known that if  $g : \mathbb{R}^n \rightarrow \mathbb{R}$  is differentiable and convex and  $A : \mathbb{R}^r \rightarrow \mathbb{R}^n$  ( $r \geq n$ ) is differentiable with  $\nabla A$  of rank  $n$ , then  $f = g \circ A$  is invex (Mishra et al., 2008). Craven and Glover (1985) proved that a function  $f$  is invex if and only if every stationary point is a global minimizer. So the second condition (ii) in Definition 1 is equivalent to say that  $\alpha\phi(v) + (1 - \alpha)\phi(-v)$  is an invex function with exactly one critical point for all  $\alpha \in (0, 1)$ . We can then show that with  $\phi$  being our two-robust loss function, for any sample  $\mathcal{S}$ , the empirical risk  $\widehat{R}_{\phi,n}$  is an invex function on the set  $\mathbb{F}^n := \{(F(x_1), \dots, F(x_n)) : F \in \mathcal{F}\} \subset \mathbb{R}^n$ . Now the problem is whether we can decompose this empirical risk as a composition of a convex function and a differentiable function, that is, whether we can write  $\widehat{R}_{\phi,n} = g \circ A$  where  $g : \mathbb{R}^n \rightarrow \mathbb{R}$  is a differentiable convex function and  $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is a differentiable vector function with  $\nabla A$  of rank  $n$ . We let  $g$  be the empirical risk  $\widehat{R}_{\text{exp},n}$  with exponential loss function. Then we want to find a function  $A(\vec{x}) = (A_1(\vec{x}), \dots, A_n(\vec{x}))$  such that

$$\widehat{R}_{\phi,n}(\mathbf{F}) = \frac{1}{n} \sum_{i=1}^n \frac{1}{1 + e^{y_i F(x_i)}} = \frac{1}{n} \sum_{i=1}^n e^{-y_i A_i(\mathbf{F})} = \widehat{R}_{\text{exp},n}(A(\mathbf{F})), \quad (22)$$

where  $\mathbf{F} := (F(x_1), \dots, F(x_n))$  and we ignore the constant 4 in the two-robust loss. Comparing each term in (22), we get  $\frac{1}{1 + e^{y_i F(x_i)}} = e^{-y_i A_i(\mathbf{F})}$ , that is,

$$y_i \log(1 + e^{y_i F(x_i)}) = A_i(\mathbf{F}). \quad (23)$$

The equation (22) also means that minimizing empirical risk  $\widehat{R}_{\phi,n}$  on the set  $\mathbb{F}^n$  is equivalent to minimizing the empirical risk  $\widehat{R}_{\text{exp},n}$  on the transformed set  $A(\mathbb{F}^n)$ , where  $A$  is defined in (23). On each data point, we have  $F_{\text{exp}}(x_i) = A_i(\mathbf{F}) = y_i \log(1 + e^{y_i F(x_i)})$ . We plot the function  $l(v) := \log(1 + e^v)$  together with the identity function in Figure 4. If we compare two empirical risk



minimization (ERM) problems (1)  $\min_{\mathbf{F} \in \mathbb{F}^n} \widehat{R}_{\text{exp},n}(\mathbf{F})$  and (2)  $\min_{\mathbf{F} \in \mathbb{F}^n} \widehat{R}_{\text{exp},n}(A(\mathbf{F}))$ , then  $A$  can be viewed as an "influence trimming" function. For any current estimate  $\mathbf{F}$ , we observe that whenever  $y_i F(x_i) \gg 0$ , then  $F_{\text{exp}}(x_i) \approx F(x_i)$ ; otherwise if  $y_i F(x_i) \ll 0$ , then  $F_{\text{exp}}(x_i) \approx 0$ , that is, instead of saying  $F$  made a severe mistake at  $(x_i, y_i)$ , after the  $A$ -transformation, we just say  $F_{\text{exp}}$  is uncertain about this point.




Figure 4: Comparison of function  $l$  and the identity function.

#### 4.2.2 Breakdown point

Empirical robustness properties defined as breakdown point in [Donoho and Huber \(1983\)](#) has proved most successful in the context of location, scale and regression problems (e.g. [Rousseeuw \(1984\)](#); [Stromberg and Ruppert \(1992\)](#); [Tyler \(1994\)](#), etc.). This success has sparked many attempts to extend the concept to other situations (e.g. [Ruckstuhl and Welsh \(2001\)](#); [Genton and Lucas \(2003\)](#); [Davies and Gather \(2005\)](#), etc.). However, very little work has been done in the classification context.

The breakdown point, as defined in [Hampel \(1968\)](#), is roughly the smallest amount of contamination that may cause an estimator to take on arbitrarily large aberrant values. The breakdown points of  $1/n$  for the mean and  $1/2$  for the median do reflect their finite-sample behavior. However, an alternative view is desired in the classification context as the magnitude of an estimator may not relate to necessarily bad classification – that is, the size of the weak hypothesis is only marginally related to the classification boundary. [Henning \(2004\)](#) exploits cluster stability for the breakdown point analysis of estimators in the mixture models. However, in the context of boosting it is not clear if the same concept can differentiate between AdaBoost and LogisticBoost algorithms; the first is known not to be robust and the second is known to have better finite sample analysis. We argue that the case when the gradient of the classification loss takes on alternating directions is in a certain sense similar to diverging estimator values in regression setting. In this sense, the proposed breakdown study resembles Hampel’s original definition. Hence, we look for the largest number of perturbed data points that keep the gradient of the risk minimization in the correct direction. Moreover, we relate this concept to the weight updating rule of Arch Boost algorithms.

Let  $\mathcal{S}_n = \{(X_1, Y_1), \dots, (X_m, Y_m)\} \cup \{(X_{m+1}, Y_{m+1}), \dots, (X_n, Y_n)\}$  be a set of observed, contaminated samples among which  $\mathcal{O}_n = \{(X_{m+1}, Y_{m+1}), \dots, (X_n, Y_n)\}$  is a set of outliers. On the original dataset  $\mathcal{S}$ , let  $h_t$  be the weak hypothesis at iteration  $t$ , and denote  $\mathbf{h}_t = (h_t(X_1), \dots, h_t(X_n))$ . We also denote  $-\mathbf{g}_t = (-g_t(X_1), \dots, -g_t(X_n)) = (Y_1 w_t(X_1, Y_1), \dots, Y_n w_t(X_n, Y_n))$  as the negative gradient of the empirical risk  $\widehat{R}_{\phi,n}$  on the sample  $\mathcal{S}$ . Similarly,  $-\mathbf{g}_o = (-g_t(X_1), \dots, -g_t(X_m), 0, \dots, 0) = (Y_1 w_t(X_1, Y_1), \dots, Y_m w_t(X_m, Y_m), 0, \dots, 0)$  is obtained by embedding the negative gradient of

the empirical risk  $\widehat{R}_{\phi,m}$  on the sample  $\mathcal{S} \setminus \mathcal{O}$  without outliers, into  $\mathbb{R}^n$ . We consider the inner product of the weak hypothesis  $\mathbf{h}_t$  and the modified negative gradient  $-\mathbf{g}_o$ . We have the following theorem.

**Theorem 2.** *Using the same notations as above, let  $\eta_j = \min^{-1}(p_j, 1 - p_j)|p_j - \frac{1}{2}|$ , for every region  $R^j$ , where  $p_j \in (0, 1)$  and  $p_j \neq \frac{1}{2}$ . Then, if any Arch Boost algorithm, at iteration  $t$ , conditional on the realizations  $\{(X_i, Y_i) = (x_i, y_i)\}_{i=1}^n$ , satisfies for all  $R_j$*

$$\sum_{i: x_i \in \mathcal{O}_n \cap R^j} w_t(x_i, y_i) \leq \eta_j \sum_{i: x_i \in R^j \setminus \mathcal{O}_n} w_t(x_i, y_i),$$

then the gradient descent direction is preserved, that is,  $-\langle \mathbf{g}_o, \mathbf{h}_t \rangle \geq 0$ .

Theorem 2 suggests that any Arch Boosting algorithm that satisfies conditions above, preserves the direction of the descent of the non-contaminated empirical  $\phi$ -risk, hence disregarding the outliers. In a sense it is an oracle property as it minimizes the risk on the “clean” data. A few remarks are imminent. Conditions of the above theorem are very mild. When  $p_j = \frac{1}{2}$  – that is, the total weight on positive labels is the same as that of the negative ones in region  $R^j$  – the elements of  $\mathbf{h}$  corresponding to the points in  $R^j$  are 0 and consequently have no influence on the sign of  $-\langle \mathbf{g}_o, \mathbf{h} \rangle$ . Moreover, the case of  $p_j = 0$  or 1 is not of the main interest as in this case when all the data in that region have the same labels and hence it is reasonable to say there are no outliers. For all other cases, when  $p_j \in (0, 1) \setminus \{\frac{1}{2}\}$ , Theorem 2 establishes that whenever the summation of weights on the outliers is no larger than a constant  $\eta_j$  times the summation of weights on the data points that are not outliers, then  $\mathbf{h}$  will also be the direction along which the empirical risk of the non-contaminated data will decrease. Figure 3 shows the weight updating rule of the ArchBoost algorithm for various  $\gamma$  and clearly illustrates that the condition above is more likely to be satisfied than the AdaBoost or LogitBoost algorithm. For example, in the case of the ArchBoost, if  $y_i = -1$  and  $\mathbb{P}(Y = -1|X = x_i) = 0.001$ , then for Real AdaBoost,  $\frac{w(x_i, y_i)}{w_b} \simeq 32$ , and for ARB-2,  $\frac{w(x_i, y_i)}{w_b} \simeq 0.008$  where  $w_b$  is the weight for the outlying data point  $(x_b, y_b)$  on the boundary, that is, the data such that  $F^*(x_b) = 0$ . Hence, it can be seen that AdaBoost puts 4000 times larger weight on this data compared to the ARB-2.

### 4.2.3 Influence function

The richest quantitative robustness measure is provided by the influence function  $u \rightarrow IF(u; T, G)$  of  $T$  at  $G$  proposed by Hampel (1974) and Hampel et al. (1986). It is defined as the first Gâteaux derivative of  $T$  at  $G$ , where the point  $u$  plays the role of the coordinate in the infinite-dimensional space of probability distributions. In more details, the influence function of any statistical functional  $T$  at a distribution  $\mathbb{P}$  is the special Gâteaux derivative (if it exists)

$$IF(z; T, \mathbb{P}) = \lim_{\epsilon \rightarrow 0^+} \frac{T((1 - \epsilon)\mathbb{P} + \epsilon\Delta_z) - T(\mathbb{P})}{\epsilon},$$

where  $\Delta_z$  is the Dirac distribution at the point  $z$  such that  $\Delta_z(\{z\}) = 1$ . It describes the effect of an infinitesimal contamination at the point  $u$  on the estimate, standardized by the mass of the contamination. Additionally, it gives the effect that an outlying observation may have on an estimator. If it is bounded, then the effect of an outlier is infinitesimal. It is straightforward to obtain the influence functions associated with the parametric estimators in linear regression models. However, the nonparametric regression models are far more complicated.

To simplify the analysis, we consider a subclass of binary classification models, in which the true boundary  $F^*$  is assumed to belong to a class of functions  $H$ . Here,  $H$  is defined as a Reproducing Kernel Hilbert Space (denoted with RKHS from now on) with a bounded kernel  $k$  and the induced

norm  $\|\cdot\|_H$ . Observe that ArchBoost, much like AdaBoost and LogitBoost, belongs to a broad class of empirical risk minimization methods (see Theorem 1) and converges only if it is properly regularized (stopped after a certain number of steps; see Theorem 4). Hence, to study its robustness properties we consider

$$f_{\mathbb{P},\lambda} = \operatorname{argmin}_{f \in H} \{ \mathbb{E}_{\mathbb{P}} [\phi(Y, f(X))] + \lambda \|f\|_H^2 \},$$

which can be viewed as the population quantity of interest. Here, the loss  $\phi$  is a function of tuple  $(Y, f(X))$  for convenience. The solution  $f_{\mathbb{P},\lambda}$  is viewed as a map that for every fixed value of the regularization parameter,  $\lambda$ , assigns an element of the RKHS  $H$  to every distribution  $\mathbb{P}$  on a given set  $Z \subset \mathcal{X} \times \mathcal{Y}$ . Observe that  $\mathbb{P}$  is not contaminated as the analysis herein is completed at the population level. In the risk minimization context,  $T(\mathbb{P}) = f_{\mathbb{P},\lambda}$ , and for each  $\mathbb{P}$ , the influence function  $IF(z; T, \mathbb{P}) \in H$ . We also denote the map  $\Psi : \mathcal{X} \rightarrow H$ , defined as  $\Psi(x) = k(x, \cdot)$ . The influence function of  $f_{\mathbb{P},\lambda}$  then takes the form

$$IF(z; T, \mathbb{P}) = -S^{-1} \circ J, \tag{24}$$

where  $\circ$  is defined to mean  $S^{-1}$  acting on  $J$  and operators  $S : H \rightarrow H$  and  $J \in H$  are defined as

$$\begin{aligned} S &= \mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, f_{\mathbb{P},\lambda}(X)) \langle \Psi(X), \cdot \rangle \Psi(X) \right] + 2\lambda id_H, \\ J &= \phi'(z_y, f_{\mathbb{P},\lambda}(z_x)) \Psi(z_x) - \mathbb{E}_{\mathbb{P}} [\phi'(Y, f_{\mathbb{P},\lambda}(X)) \Psi(X)], \end{aligned}$$

where  $id_H : H \rightarrow H$  is the identity mapping and  $z = (z_x, z_y) \in \mathcal{X} \times \mathcal{Y}$  is the contamination point. In the above display, the derivative is defined as  $\phi'(u, v) := \frac{\partial}{\partial v} \phi(u, v)$ . For the proof of this result, we refer to the proof of Theorem 4 in Christmann and Steinwart (2004) and of Theorem 15 in Christmann and Steinwart (2007). In the above the convexity of the loss function with respect to the second argument is required for  $S$  to be invertible for all  $\lambda > 0$ . For a non-convex loss function  $\phi$ ,  $\phi''$  is not guaranteed to be nonnegative. However, we observe that it is sufficient to have the non-negativity of the expectation rather than of the second derivative itself. Moreover the non-negativity is only required in a local neighborhood of the true parameter of interest,  $F^*$ . In the following lemma, we show the properties of  $\mathbb{E}_{\mathbb{P}} [\phi''(Y, \cdot)]$ .

**Lemma 4.** *For a binary classification problem, given any distribution  $\mathbb{P}$ , whenever  $\phi$  is a twice continuous differentiable Arch boosting loss function and  $F^*$  is obtained by (6), then  $\mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, F^*(X)) q^2(X) \right] \geq 0$  for any measurable function  $q : \mathcal{X} \rightarrow \mathbb{R}$ .*

*Furthermore, if  $\mathbb{P}$  and  $\mathcal{X}$  are such that  $\mathbb{P}(Y = 1|X = x) \in [\delta, 1 - \delta]$  for some  $0 < \delta < \frac{1}{2}$ , and if  $p\phi''(1, v_p^*) + (1 - p)\phi''(-1, v_p^*) > 0$  at the global minimum  $v_p^*$  for all  $p \in [\delta, 1 - \delta]$ , then there exists  $r > 0$  such that  $\mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, G(X)) q^2(X) \right] \geq 0$  for all measurable function  $G$  with  $\|G - F^*\|_{\infty} < r$ .*

A few comments are necessary. Conditions of the above lemma are satisfied for the case of the  $\gamma$ -robust loss function. For example, for the case of the two-robust loss function, one can show that for any  $x$ ,  $\mathbb{E}_Y [\phi''(Y, F^*(X)) q^2(X) | X = x] = 2p_x^2(1 - p_x)^2 q^2(x) \geq 0$  where  $p_x = \mathbb{P}(Y = 1|X = x)$ . Thus,  $\mathbb{E}_{\mathbb{P}} \phi''(Y, F^*(X)) q^2(X) \geq 0$ . Furthermore, if  $p_x \in [\delta, 1 - \delta]$  for some  $\delta \in (0, \frac{1}{2})$ , then  $p_x \phi''(1, F^*(x)) + (1 - p_x) \phi''(-1, F^*(x)) = 2p_x^2(1 - p_x)^2 \geq 2\delta^2(1 - \delta)^2 > 0$  for all  $p_x \in [\delta, 1 - \delta]$ . Observe that the condition of  $p_x \in [\delta, 1 - \delta]$  for some  $\delta \in (0, \frac{1}{2})$  restricts our setting to the ‘‘low-noise’’ setting where the true probability of the class membership is bounded away from 0 or 1. An example of a model where such condition holds is  $\mathbb{P}(Y = 1|X = x) = (1 + e^{-\gamma f(x)})^{-1}$  with  $\mathcal{X}$  being a compact space.

Recall that  $F^*$  is the solution to the risk minimization problem defined over all measurable functions. If we restrict the search to the RKHS,  $H$ , with bounded kernel, then  $F^* \neq f_{\mathbb{P},\lambda}$  in general. However, as long as  $H$  is rich enough such that  $f_{\mathbb{P},\lambda}$  is close to  $F^*$  in the sense that  $\|f_{\mathbb{P},\lambda} - F^*\|_\infty$  is small, we have the following Theorem. This Theorem should be treated as a clue of the robustness of non-convex loss functions as it shows that the influence function is bounded in  $H$  and decreases when the contamination point  $z$  is more of an outlier. Furthermore, this theorem again emphasizes that the robustness mainly comes from the fact that  $|\phi'| \rightarrow 0$ , when the magnitude of the “margin” is large.

**Theorem 3.** *For a binary classification problem, let  $\phi : \mathbb{R} \rightarrow [0, \infty)$  be a twice continuously differentiable Arch boosting loss function and let  $H$  be a RKHS with bounded kernel  $k$ . Assume  $\mathbb{P}$  is a distribution on  $\mathcal{X} \times \mathcal{Y}$  such that  $\mathbb{P}(Y = 1|X = x) \in [\delta, 1 - \delta]$  for all  $x \in \mathcal{X}$  and for some  $0 < \delta < \frac{1}{2}$  and  $p\phi''(1, v_p^*) + (1 - p)\phi''(-1, v_p^*) > 0$  at the global minimum  $v_p^*$  for all  $p \in [\delta, 1 - \delta]$ . Then there exists  $r > 0$  such that for all  $\|f_{\mathbb{P},\lambda} - F^*\|_\infty < r$*

$$\|IF(z; f_{\mathbb{P},\lambda}, \mathbb{P})\|_H \leq \sqrt{\frac{C_\phi}{\lambda}} + \frac{M_k |\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))|}{2\lambda}, \quad (25)$$

where  $M_k$  is the upper bound of the kernel  $k$  and  $C_\phi = \phi(0, 0)$ .

For any non-convex Arch boosting loss function, due to Assumption 2, we have  $|\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))| \rightarrow 0$  when  $z_y f(z_x) \rightarrow -\infty$  or  $z_y f(z_x) \rightarrow \infty$ . Result of Theorem 3 implies that for all smooth enough classification boundaries  $F^*$ , the Arch Boosting algorithm has a bounded influence function whenever  $\lambda > 0$ . If we plot  $\|IF(z; f_{\mathbb{P},\lambda}, \mathbb{P})\|_H$  versus  $z_y f_{\mathbb{P},\lambda}(z_x)$ , then it will decrease towards a constant far from the origin, just like the loss function of a redescending M-estimator (Huber, 1981). Moreover, we observe that  $\|IF(z; f_{\mathbb{P},\lambda}, \mathbb{P})\|_H$  is unbounded for the exponential loss (AdaBoost) and bounded but not diminishing for the logistic loss (Logit Boost).

### 4.3 Statistical Consistency of Arch Boost

Per Theorem 1, we observe that the Arch Boost Algorithm 1, can be formulated as an empirical risk minimization procedure, for which the consistency has been established in the case of convex loss functions (Bartlett et al., 2006; Bartlett and Traskin, 2007; Jiang, 2004; Zhang and Yu, 2005). However, we extend this framework to include non-convex losses by exploring local curvatures. Given any training sample  $\mathcal{S}_n = \{(X_i, Y_i)\}_{i=1}^n$ , we compute a classifier  $f_n$  on  $\mathcal{S}_n$ , and denote the misclassification error to be  $L(f_n) = \mathbb{P}(f_n(X) \neq Y|\mathcal{S}_n)$ . Moreover, the Bayes risk is  $L^* = \inf_{f \in \mathcal{M}} L(f) = \mathbb{E}_X[\min(\eta(X), 1 - \eta(X))]$ , where  $\eta(X) = \mathbb{P}(Y = 1|X)$ . Let  $\mathcal{M}$  be the family of measurable functions. The three key steps for proving consistency (?) include: (I) utilizing the property of the loss that whenever the  $\phi$ -risk  $R_\phi(f_n)$  converges to the minimal risk

$$R_\phi^* = \inf_{f \in \mathcal{M}} R_\phi(f),$$

the misclassification error converges to the Bayes risk; (II) there exists a deterministic sequence  $\widetilde{f}_n$  for which the  $\phi$ -risk approaches the minimal risk  $R_\phi^*$  as  $n \rightarrow \infty$ ; (III)  $\phi$ -risk of the estimated  $f_n$  can be approximated by the  $\phi$ -risk of the deterministic, reference sequence  $\widetilde{f}_n$ . Observe that  $f_n$  is sample dependent, that is,  $f_n$  is a function of  $\mathcal{S}_n$ , and hence  $L(f_n)$  is again a random variable. Hence, in step (II), we choose a deterministic reference sequence such that  $F_n$  is only a function of the sample size  $n$ . The first condition (I) is automatically satisfied if we use Arch boosting loss function because of the classification calibration condition. Note that the above inequalities are true for any sample  $\mathcal{S}_n$  and sample size  $n$ . From now on, we will have several assumptions. The first one is about the class  $\mathcal{H}$  and the distribution  $\mathbb{P}$ .

**Assumption 1.** Let the class of the weak hypothesis  $\mathcal{H}$  and the probability distributions  $\mathbb{P}$  be such that  $\lim_{T \rightarrow \infty} \inf_{f \in \mathcal{F}^T} R_\phi(f) = R_\phi^*$  and  $d_{VC}\{\mathcal{H}\} < \infty$ .

For certain rich enough class  $\mathcal{H}$ , Assumption 1 is true for any distribution  $\mathbb{P}$  (Bartlett and Traskin, 2007). One example is the class  $\mathcal{T}$  of binary trees with the number of terminal nodes larger or equal to  $d + 1$ , where  $d$  is the dimension of  $\mathcal{X}$  (Breiman, 2004). The next assumption is about the loss function  $\phi$ .

**Assumption 2.** Let the loss function  $\phi$  be a bounded, decreasing, Lipschitz function that belongs to the class of Arch Boosting loss functions.

If  $\phi$  satisfies Assumption 2, then we know both  $\lim_{v \rightarrow \infty} \phi(v)$  and  $\lim_{v \rightarrow -\infty} \phi(v)$  exist in  $\mathbb{R}$ . A class of loss functions that satisfy Assumption 2 incorporates a class of loss functions for which the first derivative converges to zero away from the origin; an example is a redescending Hampel's three part loss function. This lessens the effect of gross outliers and in turn leads to many good robust properties of the resulting estimator. All the  $\gamma$ -robust loss functions ( $\gamma \geq 1$ ) satisfy this condition. Commonly used loss functions like least squares, exponential loss and logistic loss do not satisfy this condition because they are not bounded.

Next we state two results important for proving the consistency of the Arch Boosting estimator. Much like the consistency of AdaBoost, the proof hinges upon the optimal choice of stopping times. However, the statement is not dependent on an additional truncation level of the functions  $f$  (like Lemma 4 in Bartlett and Traskin (2007)) because of the boundedness of our loss function  $\phi$ . The nice decaying property of the derivative of the proposed Arch Boosting loss enables us to avoid additional parameters. The proof is given in the Appendix A.

**Theorem 4.** Assume  $\mathcal{H}$  and distribution  $\mathbb{P}$  satisfy Assumption 1. Let  $\phi$  be an Arch boosting loss function satisfying Assumption 2. Then for any sample  $\mathcal{S}_n$  and for a non-negative sequence of stopping times  $T_n = n^{1-\varepsilon}$  with  $\varepsilon \in (0, 1)$ , we have for  $f_n^* = \operatorname{argmin}_{f \in \mathcal{F}^{T_n}} \widehat{R}_{\phi,n}(f)$  as  $n \rightarrow \infty$ ,

$$(a) \sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| \rightarrow 0 \text{ a.s.} \quad \text{and} \quad (b) R_\phi(f_n^*) \rightarrow R_\phi^* \text{ a.s.}$$

Theorem 4 illustrates the uniform deviation between the  $\phi$ -risk and the empirical  $\phi$ -risk. Note that we want  $T_n \rightarrow \infty$  as  $n \rightarrow \infty$  but not too fast (slower than  $\mathcal{O}(n)$ ) in order to make sure that the uniform deviation converges to zero when  $n \rightarrow \infty$ . Moreover, from part (b) of Theorem 4, we know there exists a sequence of samples  $\{S_n^*\}_{n=1}^\infty$  such that  $R_\phi(\tilde{f}_n) \rightarrow R_\phi^*$  as  $n \rightarrow \infty$  where  $\tilde{f}_n$  is the optimal classifier obtained by minimizing the empirical risk on  $S_n^*$ . In another word, from  $\omega \in \Omega$  such that  $R_\phi(f_n^*(\omega)) \rightarrow R_\phi^*$  as  $n \rightarrow \infty$ , we pick a  $\omega_0$  and set  $\tilde{f}_n = f_n^*(\omega_0)$ . We let  $\{\tilde{f}_n\}$  be our reference sequence, and note that  $\tilde{f}_n$  only depends on sample size  $n$ , that is,  $\{\tilde{f}_n\}$  is a deterministic sequence. This is necessary to decouple the dependence of the estimator and the sample. Next we state the intermediary lemma that connects the reference sequence,  $\tilde{f}_n$ , to the Arch Boost estimator,  $F_{T_n}$ .

**Lemma 5.** For the above reference sequence  $\{\tilde{f}_n\}_{n=1}^\infty$  and a non-negative sequences  $T_n = n^{1-\varepsilon}$ ,  $\varepsilon \in (0, 1)$ , we have as  $n \rightarrow \infty$ , (a)  $\left(\widehat{R}_{\phi,n}(\tilde{f}_n) - R_\phi(\tilde{f}_n)\right)_+ \rightarrow 0$  a.s. and (b)  $\left(\widehat{R}_{\phi,n}(F_{T_n}) - \widehat{R}_{\phi,n}(\tilde{f}_n)\right)_+ \rightarrow 0$  a.s.

Lemma 5 is proved easily with the help of Theorems 1 and 4. Together with the result of Theorem 4 we can finally state the consistency result.

**Corollary 1.** Assume the class  $\mathcal{H}$  and distribution  $\mathbb{P}$  satisfy Assumption 1 and that the Arch boosting loss function  $\phi$  satisfies Assumption 2. Then, as  $n \rightarrow \infty$ , the sequence of classifiers  $F_{T_n}$  returned by the Arch Boost algorithm stopped at the step  $T_n$ , chosen in Theorem 4, satisfies  $L(\operatorname{sign}(F_{T_n})) \rightarrow L^*$  a.s.

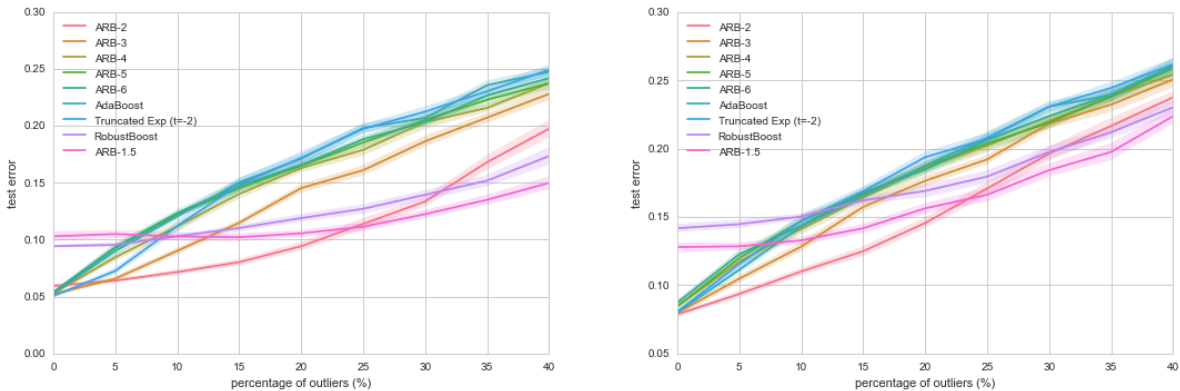
## 5 Numerical Experiments

In this section, we will test the Arch Boost ( $\phi$ ) algorithm under different loss functions  $\phi$  on binary classification problems.

### 5.1 Simulated Examples

We will generate datasets using 'make-hastie-10-2' and 'make-gaussian-quantiles' (Pedregosa et al., 2011). For 'make-hastie-10-2', the data have 10 features that are standard independent Gaussian and the target  $y_i = 1$  if  $\sum_{i=1}^{10} x_i^2 > \chi_{10,1/2}^2$  (Friedman, et al., 2001). For 'make-gaussian-quantiles', in our case, the dataset is constructed by taking a 20-dimensional normal distribution  $\mathcal{N}(\mathbf{0}, 2I_{20})$  and defining two classes as  $y = 1$  if  $\sum_{i=1}^{20} x_i^2 > 4\chi_{20,1/2}^2$ . In both datasets, the classes are separated by a concentric multi-dimensional spheres with origin at  $\mathbf{0}$  such that roughly equal numbers of samples are in each class.

We test the robustness of ARB- $\gamma$  algorithms by adding noise to different percentages of the training samples. In the following two examples, we add independent t-distribution (df = 4) noise to the features of a % of selected training samples. The results are summarized in Figure 5. For each dataset, we generate 14000 samples using the corresponding methods. Among the 14000 data, we use 2000 for training, 2000 for cross validation, and the rest 10000 for testing. The number of weak classifiers is 1000, and by cross validation, we set the step sizes,  $\alpha_t$ , to be 0.78 for ARB-1.5, 0.45 for ARB-2, 0.28 for ARB-3, 0.20 for ARB-4, 0.14 for ARB-5, 0.10 for ARB-6, and 0.80 for Real AdaBoost. For RobustBoost, we tune the target parameter for each percentage of errors using bisection search. In each figure, we plot the average test errors and the corresponding 95% confidence intervals.



(a) make-hastie-2 dataset.

(b) Gaussian quantiles dataset.

Figure 5: Comparison of average test errors of ARB- $\gamma$ , AdaBoost, Arch Boost with truncated exponential loss (truncated at  $t = -2$ ) and RobustBoost.

From Figure 5, we have several observations. First, the test errors of the ARB- $\gamma$  algorithms are all less than that of the Real AdaBoost. Moreover, they are smaller than that of the truncated exponential loss, indicating that the introduced non-convex family of losses substantially outperforms traditional truncation losses. Second, when the percentage of outliers is less than a certain level (around 23% in Figure 5a and Figure 5b), the performances of ARB-2 is the best. When the noise level is higher, ARB-1.5 behaves the best. Moreover, RobustBoost has higher test error when the noise level is low. We run the RobustBoost algorithm for 1000 iterations and obtain an error that is larger than zero. However, if we run RobustBoost for long enough when error rate is 0,

its performance should converge to that of Real AdaBoost. RobustBoost and ARB-1.5 have very similar performance. ARB-2 is worse than RobustBoost when the noise level is very high. However, note that we tune a number of tuning parameters of the RobustBoost at each noise level. If we were to also "tune"  $\gamma$  for ARB- $\gamma$  algorithms, for example, choose ARB-2 when noise level is less than 25% and ARB-1.5 otherwise in Figure 5a, then we could observe that ARB- $\gamma$  is uniformly better than the RobustBoost.

To illustrate the importance of the loss function choice and the Arch Boosting method, we implement a Gradient Descend Boosting algorithm with a "trimmed" version of the exponential loss, i.e., the truncated exponential loss function. We observe that the improvement over AdaBoost is extremely minor and it disappears when the dimensionality of the problem grows. For 'make-hastie-10-2' dataset truncated exponential loss is better than  $\gamma$ -robust losses for  $\gamma \geq 4$  and  $\epsilon < 15\%$  after which point it is very much indistinguishable from the AdaBoost. Situation is even better for 'make-gaussian-quantiles' dataset as truncated exponential is almost identical as the AdaBoost as soon as  $\epsilon > 7\%$ . This suggests that the Arch Boosting framework is essential for robust and generalizable performance.

## 5.2 Long/Servedio problem

Long and Servedio (2010) constructed a challenging classification setting described as follows. The input  $X \in \mathbb{R}^{21}$  with binary features  $X_i \in \{-1, +1\}$  and label  $y_i \in \{-1, +1\}$ . First, the label  $y$  is chosen to be  $-1$  or  $+1$  with equal probability. Then for any given  $y$ , the features  $X_i$  are generated according to the following mixture distribution:

- **Large margin:** With probability  $\frac{1}{4}$ , set  $X_i = y$  for all  $1 \leq i \leq 21$ .
- **Pullers:** With probability  $\frac{1}{4}$ , set  $X_i = y$  for  $1 \leq i \leq 11$  and  $X_i = -y$  for  $12 \leq i \leq 21$ .
- **Penalizers:** With probability  $\frac{1}{2}$ , randomly choose 5 coordinates from the first 11 features and 6 from the last 10 to be equal to  $y$ . The remaining features are set to  $-y$ .

The data from this distribution can be perfectly classified by  $sign(\sum_i X_i)$ . We generate 800 samples from this distribution and flip each label with probability  $\epsilon \in [0, 0.5)$ . Then we train the classifier on the noisy data and test the performance on the original clean data. We first generate 20 datasets according to the distribution, and on each of them, we randomly flip  $\epsilon = 10\%$  of the labels. The result is in Table 5, where we record the average test error and also report the sample deviations in the brackets. We can see that the ARB-2 outperforms Real AdaBoost and LogitBoost (Friedman, et al., 2000), and is even better than RobustBoost (target parameter  $\theta = 0.15$ ) (Freund, 2001).

data type	Real AdaBoost	LogitBoost	RobustBoost ( $\theta = 0.15$ )	ARB-2
noise( $\epsilon = 0.1$ )	28.24%(1.53%)	26.61%(1.51%)	11.04%(0.67%)	<b>9.82%</b> (0.43%)
clean	25.07%(1.92%)	22.59%(1.74%)	0.21%(0.35%)	<b>0.02%</b> (0.04%)

Table 5: Long/Servedio problem

We also compare the performance of different ARB- $\gamma$  and plot the average test errors and 95% confidence intervals in Figure 6. We can see from Figure 6 that ARB-1.5 behaves the best on this dataset among all these algorithms. When  $\gamma$  increases, the performance of ARB- $\gamma$  is approaches that of the Real AdaBoost. The breakdown point will get higher when  $\gamma \rightarrow 1^+$ , implying that the smaller  $\gamma$  lead to the better robustness properties. When  $\gamma = 2$ , then breakdown point is about

15%, and when  $\gamma = 1.5$  and 1.3, the breakdown point is about 20%. But for ARB-1.3, the test error when labels are not flipped is not zero.




Figure 6: Comparison of ARB- $\gamma$  on Long/Servedio problem with different  $\epsilon$  (the probability of flipping the labels).

### 5.3 Outlier detection

We have shown in previous sections that ARB- $\gamma$  algorithms are more robust to the noise. Therefore, because of the robustness, ARB- $\gamma$  should be able to detect the outliers. Intuitively, if a point is an outlier, then it should be misclassified by most of the weak hypotheses of ARB-2. In this experiment, we generate 2000 data points using 'make-hastie-10-2' and randomly shuffle them. Then we add a noise drawn from a t-distribution ( $df = 4$ ) to each of the 10 features of the first  $\epsilon$  percentage data points. After running the algorithms for 800 iterations, we record the times that each data point is misclassified, and count the number of points that are misclassified more than 600 times (denoted as  $T$ ), and count how many of them (denoted as  $T_o$ ) actually belong to the noisy set that we add noise to. Finally we calculate the ratio  $\frac{T_o}{T}$ . This ratio describes the chance that a data point is an outlier. By cross-validation, we set the step size  $\alpha = 0.5$  for the ARB-2 and  $\alpha = 0.8$  for the Real AdaBoost. The results are shown in Table 6. The x-axis stands for the index of the training points ranging from 1 to 2000, and the y-axis stands for the times a point is misclassified, ranging from 0 to 800. We can see that when the percentage of outliers is less than 15%, for the ARB-2, more than 99% of the points that have been misclassified for 600 times or above, are indeed the outliers, but for the Real AdaBoost, this number is only around 31%. Informally, for ARB-2, when  $\epsilon \leq 15\%$ , we have more than 99% "confidence" to conclude that a data point, that is misclassified for more than 600 times, is an outlier.



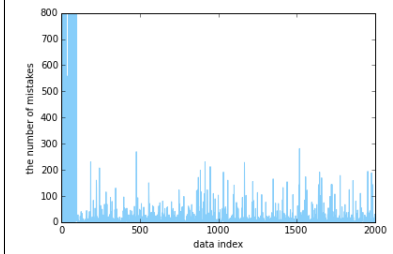
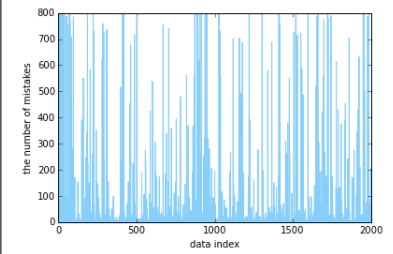
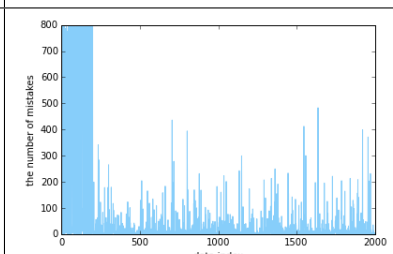
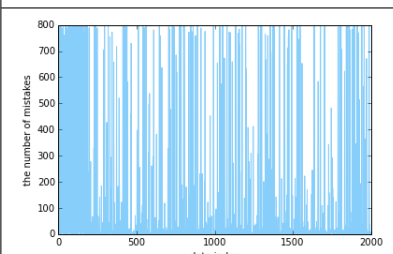
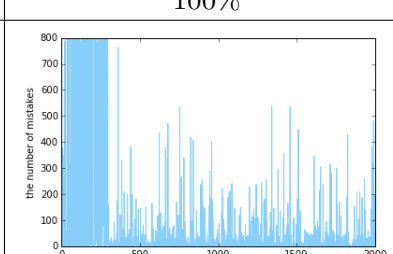
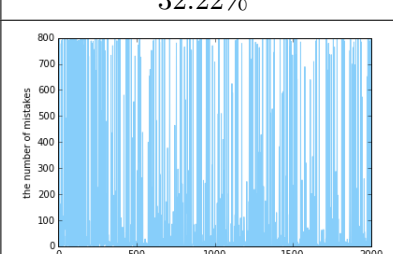
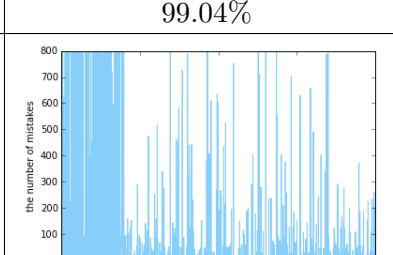
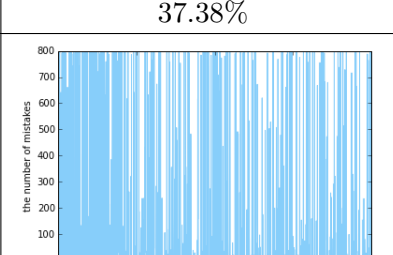
$\epsilon$	ARB-2	Real AdaBoost
5%		
$T_o/T$	100%	30.49%
10%		
$T_o/T$	100%	32.22%
15%		
$T_o/T$	99.04%	37.38%
20%		
$T_o/T$	85.48%	44.40%

Table 6: Outliers detection. The x-axis stands for the index of the training points ranging from 1 to 2000, and the y-axis stands for the times a point is misclassified, ranging from 0 to 800.

## 5.4 Real data application

### 5.4.1 Wisconsin (diagnostic) breast cancer data set

We test our ARB- $\gamma$  algorithms on the Wisconsin (diagnostic) breast cancer data set of [Street et al. \(1993\)](#), which is available on the machine learning repository website at the University of California, Irvine: [https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic)). The data set was created by taking measurements from a digitized image of a fine needle aspirate of a breast mass for each of 569 individuals, with 357 benign and 212 malignant instances. Ten real-valued features are computed for each cell nucleus: radius, texture, perimeter, area, smoothness, compactness, concavity, concave points, symmetry, fractal dimension.

Table 7: Comparison of the average test errors and sample deviation of four algorithms on the Wisconsin breast cancer dataset.

Percentage of flipped labels	Methods			
	ARB-2	ARB-1.5	Robust Boost	Ada Boost
0%	<b>3.35%</b> (1.29%)	3.49%(1.23%)	4.52%(1.61%)	4.06%(1.48%)
5%	4.71%(1.76%)	<b>4.43%</b> (1.63%)	4.78%(1.66%)	5.38%(1.92%)
10%	5.71%(1.71%)	<b>5.03%</b> (1.60%)	5.35%(1.77%)	6.24%(1.97%)
15%	6.92%(2.06%)	<b>5.84%</b> (2.04%)	6.47%(2.14%)	7.01%(2.28%)

We randomly split the data into an equally balanced training set with 150 benign samples and 150 malignant samples, and the rest samples were used for testing. The maximum iterations is set to be 200, and a five-fold cross-validation is implemented on the training set to select the step size and stopping time ( $\leq 200$ ) for each algorithm. This procedure was repeated for 100 times and an average of the test error is reported in Table 7. Boxplots of the test error are presented in Figure 7. We observe that ARB-2 behaves the best on the original data set, and ARB-1.5 outperforms others




Figure 7: Comparison of ARB-2, ARB-1.5, RobustBoost and Real AdaBoost on the UCI Wisconsin breast cancer dataset. In each subfigure, from left to right are the box plots for the test errors of ARB-2, ARB-1.5, RobustBoost and Real AdaBoost.

when there is noise. Compared to [Stefanski et al. \(2014\)](#) who obtain the best test error rate of about 4%, all of our methods uniformly achieve smaller test error rate, on the clean and comparable test error rates on the perturbed datasets.

### 5.4.2 Sensorless drive diagnosis data set

We compare ARB-2, ARB-1.5, RobustBoost and Real AdaBoost on the dataset sensorless drive diagnosis, which is also available on the UCI machine learning repository: <https://archive.ics.uci.edu/ml/datasets/Dataset+for+Sensorless+Drive+Diagnosis>.

Table 8: Comparison of the average test errors and sample deviation of four algorithms on the Sensorless drive diagnosis dataset.

Percentage of flipped labels	Methods			
	ARB-2	ARB-1.5	Robust Boost	Ada Boost
0%	<b>5.84%</b> (0.43%)	6.52%(0.43%)	7.37%(0.36%)	7.04%(0.44%)
5%	9.68%(0.58%)	<b>8.79%</b> (0.57%)	8.96%(0.51%)	11.52%(0.78%)
10%	12.35%(1.03%)	10.88%(0.94%)	<b>10.81%</b> (0.74%)	14.63%(0.98%)
15%	14.16%(0.85%)	12.50%(0.76%)	<b>11.95%</b> (0.99%)	17.88%(0.81%)

This dataset contains 58509 instances and each of them has 49 features all extracted from the electric current drive signals. A range of typical defects in drive train applications are considered with 11 different classes present. We combine the data points with label  $\leq 6$  into one class and the rest into the other class. Then at each time, we randomly choose 14000 points and use 2000 for training, 2000 for cross validation and 10000 for testing. We use the cross validation set to choose the stopping time ( $\leq 1000$  iterations) and step sizes. According to the noise levels, a certain proportion of the labels of the training data points will be randomly flipped. We summarized the test errors using box plots in Figure 8 and calculated the mean and sample deviation in Table 8.




Figure 8: Comparison of ARB-2, ARB-1.5, RobustBoost and Real AdaBoost on the UCI sensorless drive diagnosis dataset. In each subfigure, from left to right are the box plots for the test errors of ARB-2, ARB-1.5, RobustBoost and AdaBoost.

We observed that ARB-2 performs the best on the original data without adding any extra noise. RobustBoost again behaves worse than others on the original dataset. One reason is that it needs more time to terminate when target error is near 0, and the other reason is that it cannot distinguish outliers and hard inliers (Kobetski et.al, 2013). When we flipped 10% of the labels, ARB-1.5 outperformed others and when we flipped 10% or 15% of the labels, RobustBoost behaved the best. But in all of the three cases with noise, the test errors of ARB-1.5 and RobustBoost are very close. However, ARB-1.5 does not need to fine tune any target parameters at each different noise levels.

### 5.4.3 MAQC-II Project: human breast cancer (BR) data set

We next test our Algorithms on a dataset that is part of the 'MicroArray quality control II' project. It is available from the gene expression omnibus database with accession number GSE20194: <http://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE20194>.

Table 9: Comparison of the average test errors and sample deviation of four algorithms on the GSE20194 gene dataset.

Percentage of flipped labels	Methods			
	ARB-2	ARB-1.5	RobustBoost	AdaBoost
0%	9.66%(1.90%)	<b>9.45%</b> (1.93%)	10.24%(2.10%)	10.51%(2.20%)
5%	11.35%(2.50%)	<b>11.16%</b> (2.20%)	11.82%(2.95%)	12.40%(2.89%)
10%	12.92%(3.43%)	<b>12.17%</b> (3.50%)	12.22%(3.04%)	14.76%(3.54%)
15%	13.46%(4.91%)	13.33%(5.03%)	<b>13.29%</b> (3.13%)	15.85%(6.85%)

The dataset contains 278 newly diagnosed breast cancer patients, aged from 26 to 79 years with population spanning all three major races and their mixtures. Patients received 6 months of preoperative chemotherapy followed by surgical resection of the cancer. Estrogen-receptor status helps guide treatment for breast cancer patients because breast cancer contains many estrogen receptors. Of 278 patients, 164 had positive estrogen-receptor status and 114 have negative estrogen-receptor status. Each sample is described by 22283 biomarker probe-sets.

To alleviate computational burden we choose 3000 probe-sets with the smallest p-values in the two-sample t-test and standardize each feature. Such simplification is often considered in high dimensional data (e.g. Zhang et. al (2014)). We randomly choose 50 samples with positive estrogen receptor status and 50 samples with negative estrogen receptor status for a training set and use the rest for the testing set. We randomly flip labels of the samples in training set according to the preassigned noise level and repeat the analysis 100 times. Then a five-fold cross-validation is implemented on the training set to select the stopping time ( $\leq 100$ ) and step sizes. We summarize the results in Table 9 and Figure 9. This dataset was previously analyzed in Deshwar and Morris (2014) and Zhang et. al (2014) where the best obtained test error was 15% and 9%, respectively. However, our methods achieve error comparable to those even when the labels were perturbed at random. This suggests that our method is extremely stable even in high-dimensional models.

## 5.5 Discussion

We showed that ArchBoost  $-\gamma$  is a robust alternative to the popular Gradient Boost  $-\gamma$ -type algorithms. The algorithmic part, presented by Theorem1, works for quite a very general class of loss functions that satisfy the Arch-Boost loss properties, presented in Definition 1. The differentiability condition is imposed artificially and we believe can be avoided by considering appropriate sub-differential analysis. However, the Condition (ii) is crucial for the analysis and we believe that it cannot be




Figure 9: Comparison of ARB-2, ARB-1.5, RobustBoost and Real AdaBoost on the GSE20194 gene dataset, from left to right are the box plots for the test errors of ARB-2, ARB-1.5, RobustBoost and AdaBoost.

relaxed. Moreover, the robustness properties depend crucially on this condition too. Additionally, the robustness part of the analysis, summarized in Theorems 2 and 3, works for quite for an arbitrary Lipschitz loss function. Hence, it presents novel proof of why is LogitBoost more robust than the AdaBoost, a folklore observation made by many experts in the field. For example, Theorem 2 is more likely to hold for LogitBoost than the AdaBoost and similarly more likely to hold for ArchBoost than the LogitBoost.

Note that Arch Boost framework can be easily explored to define an estimate of the conditional probability  $\mathbb{P}(Y = 1|x)$ . A special case is to plug in the exponential loss function, in which case we will get Real AdaBoost algorithm, and it has conditional probability estimation  $\widehat{P}_{AdaBoost}(Y = 1|X = x) = 1/(1 + e^{-2F_T(x)})$ . In contrast, many of the existing boosting methods, based on the Gradient Boosting ideas, cannot be directly applied for this purpose. In order to propose an estimator of  $\mathbb{P}(Y = 1|x)$ , we find and explore a recursive relationship between  $F_t$  and  $F_{t-1}$ . We observe that by rewriting equation (13), we have

$$\frac{\phi'(-F_t(x))}{\phi'(F_t(x))} = \frac{\phi'(-F_{t-1}(x))}{\phi'(F_{t-1}(x))} \left( \frac{\mathbb{P}_{w_t}(Y = 1|x)}{\mathbb{P}_{w_t}(Y = -1|x)} \right).$$

Suppose  $F_0 \equiv 0$ . By solving these equations recursively, after  $T$  iterations we obtain

$$\frac{\phi'(-F_T(x))}{\phi'(F_T(x))} = \prod_{t=1}^T \frac{\mathbb{P}_{w_t}(Y = 1|x)}{\mathbb{P}_{w_t}(Y = -1|x)}.$$

Now define  $\widetilde{\mathbb{P}}(Y = 1|x)$  as

$$\widetilde{\mathbb{P}}(Y = 1|x) = \frac{\prod_{t=1}^T \mathbb{P}_{w_t}(Y = 1|x)}{\prod_{t=1}^T \mathbb{P}_{w_t}(Y = 1|x) + \prod_{t=1}^T \mathbb{P}_{w_t}(Y = -1|x)} \in [0, 1]. \quad (26)$$

Then, we have the following relation

$$\frac{\phi'(-F_T(x))}{\phi'(F_T(x))} = \frac{\tilde{\mathbb{P}}(Y = 1|x)}{\tilde{\mathbb{P}}(Y = -1|x)}. \quad (27)$$

By comparing equation (6) and (27), if  $F_T(x)$  is close to  $F^*(x)$ , then  $\tilde{\mathbb{P}}(Y = 1|x)$  will be a good approximation of the conditional probability  $\mathbb{P}(Y = 1|x)$ . Observe that due to a nature of weak classifiers,  $\tilde{\mathbb{P}}(Y = 1|x)$  is guaranteed to be bounded in  $[0, 1]$  for any differentiable loss function  $\phi$ . Moreover, non-convex loss functions  $\phi$  seems to be better candidates for the class membership probability estimation as well as for the classifier estimation.

The statistical consistency proof is centered around “tilted” loss functions that are non-convex in particular. We believe that non-convex losses have great and unexplored potential for robust high dimensional statistics. The framework of “tilted” loss functions is very general and can very well be explored for robust variable selection and estimation, through an appropriate penalization scheme. Moreover, it is very well known that the impact of outliers is multiplied in case of inferential problems, such are confidence intervals and testing. By screening out many large outliers, “tilted” losses may significantly improve upon asymptotic efficiency of existing procedures.

## Appendix A Proofs

### A.1 Derivation of ARB- $\gamma$ algorithms

Note that  $\phi'_{a,\gamma}(v) = \frac{-a\gamma 2^\gamma e^{av}}{(1+e^{av})^{\gamma+1}}$ , and  $\frac{\phi'_{a,\gamma}(-v)}{\phi'_{a,\gamma}(v)} = e^{a(\gamma-1)v}$ , and when  $\gamma > 1$ ,

(i) for any data  $(x, y)$ , the optimal  $F^*(x)$  satisfies

$$\frac{\phi'_{a,\gamma}(-F^*(x))}{\phi'_{a,\gamma}(F^*(x))} = e^{a(\gamma-1)F^*(x)} = \frac{\mathbb{P}(Y = 1|x)}{\mathbb{P}(Y = -1|x)},$$

that is,

$$F^*(x) = \frac{1}{a(\gamma-1)} \log \frac{\mathbb{P}(Y = 1|x)}{\mathbb{P}(Y = -1|x)}, \quad \gamma > 1.$$

(ii) After iteration  $t$ , we will update the weights to be

$$w_{t+1}(x, y) = -\phi'_{a,\gamma}(yF(x) + yh_t(x)) = \frac{a\gamma 2^\gamma e^{a(yF(x)+yh_t(x))}}{(1 + e^{a(yF(x)+yh_t(x))})^{\gamma+1}}.$$

Since the constant  $a\gamma 2^\gamma$  will not influence the normalized weights, we can just update the weights to be

$$w_{t+1}(x, y) = \frac{e^{a(yF(x)+yh_t(x))}}{(1 + e^{a(yF(x)+yh_t(x))})^{\gamma+1}}. \quad (28)$$

(iii) At iteration  $t$ , we will update the hypothesis  $F_{t-1}(x)$  to be  $F_{t-1}(x) + h_t(x)$  such that

$$\frac{\phi'_{a,\gamma}(F_{t-1}(x))\phi'_{a,\gamma}(-F_{t-1}(x) - h_t(x))}{\phi'_{a,\gamma}(-F_{t-1}(x))\phi'_{a,\gamma}(F_{t-1}(x) + h_t(x))} = e^{a(\gamma-1)h_t(x)} = \frac{\mathbb{P}_w(Y = 1|x)}{\mathbb{P}_w(Y = -1|x)},$$

that is,

$$h_t(x) = \frac{1}{a(\gamma-1)} \log \frac{\mathbb{P}_w(Y = 1|x)}{\mathbb{P}_w(Y = -1|x)}, \quad \gamma > 1.$$

The algorithms for different  $\gamma$  (set  $a = 1$ ) is given in Algorithm 2. Since  $\frac{1}{\gamma-1}$  is just a constant, we can simply absorb it into  $\alpha_t$  and leave  $h_t = \log \frac{\mathbb{P}(Y=1|x)}{\mathbb{P}(Y=-1|x)}$ . So intuitively, the larger  $\gamma$  is, the smaller the step size  $\alpha_t$  will be. If we use constant step size, then a rule of thumb is to set  $\alpha_t = \frac{\alpha}{\gamma-1}$  for ARB- $\gamma$  where  $\alpha$  is a tuning parameter for the step size of ARB-2.

## A.2 Proof of Lemma 1

For any  $\alpha \in (0, 1)$ , let  $D(v) = \frac{d}{dv}[\alpha\phi(v) + (1-\alpha)\phi(-v)] = \alpha\phi'(v) - (1-\alpha)\phi'(-v)$ . Since  $\phi$  is convex,  $\phi'(v)$  is monotone increasing, and  $\phi'(-v)$  is monotone decreasing. So  $D(\cdot)$  is monotone increasing.

Then note that there exists  $v^*$  such that  $\phi'(v^*) \neq \phi'(-v^*)$ . Otherwise,  $\phi'$  will be both increasing and decreasing, that is, a constant. Contradiction to our assumption. Therefore,  $D(v^*)D(-v^*) < 0$ , and combined with the monotonicity and continuity of  $D(\cdot)$ , we know  $D(v) = 0$  has one and only one solution, and it is a global minimum of  $\phi$ .

## A.3 Proof of Remark 2

Note that  $\phi_{a,1}(v) = 1 - \tanh(\frac{1}{2}v)$  is a sigmoid function. Let  $C_\eta(v) = \eta\phi(v) + (1-\eta)\phi(-v) = 1 + (1-2\eta)\tanh(\frac{1}{2}v)$ , then  $\inf_{v \in \mathbb{R}} C_\eta(v) = 2 \min(\eta, 1-\eta) < 1 = \inf_{v: v(2\eta-1) \leq 0} C_\eta(v)$  for all  $\eta \neq \frac{1}{2}$ . So condition (iii) is satisfied.

But since  $C_\eta(v) = 1 + (1-2\eta)\tanh(v)$ , we know if  $\eta > \frac{1}{2}$ , then  $C_\eta(v)$  is monotone decreasing. And when  $\eta < \frac{1}{2}$ ,  $C_\eta(v)$  is monotone increasing. When  $\eta = \frac{1}{2}$ ,  $C_{\frac{1}{2}}(v) \equiv 1$ . Hence, for every  $\eta \in [0, 1]$ , there is no unique global minimum in  $\mathbb{R}$ .

## A.4 Proof of Lemma 2

Define  $f(v) = a\phi(v) + (1-a)\phi(-v)$ , then  $f'(v) = a\phi'(v) - (1-a)\phi'(-v)$  and  $f''(v) = a\phi''(v) + (1-a)\phi''(-v)$ . For any  $a \in (\frac{1}{2}, 1)$ ,  $f'(v) = 0$  if and only if  $g(v) = \frac{1-a}{a} \in (0, 1)$ . Since  $g : (0, \infty) \rightarrow (0, 1)$  is a bijection, there exists one and only one  $v^* \in \mathbb{R}^*$  such that  $v^* = g^{-1}(\frac{1-a}{a})$ . Note that if  $v > v^*$ , then  $g(v) < g(v^*)$ , that is  $f'(v) > f'(v^*)$ ; if  $v < v^*$ , then similarly we have  $f'(v) < f'(v^*)$ . So  $v^*$  is a minimum for  $\phi$ .

For  $a \in (0, \frac{1}{2})$ , since  $g(-v) = \frac{1}{g(v)}$ , we only need to solve  $g(v) = \frac{a}{1-a}$  for  $v^*$ , then  $-v^*$  will be solution for  $f'(v) = 0$  for this  $a$ . And the minimum claim follows similarly as above.

## A.5 Solution of (17)

Here, we show the derivation of one possible family of solutions. We do so by employing integrating factor method and adapting it to the nonlinear ordinary differential equation (17). Since we have  $\frac{\phi'(-v)}{\phi'(v)} = e^{(\gamma-1)v}$ , we made a reasonable guess that  $\phi(v) = f(e^v)$ . After plugging this into (17) and let  $x = e^v$ , we have

$$\frac{f'(1/x)}{f'(x)} = x^{\gamma+1}. \quad (29)$$

Furthermore, for a non-convex function  $\phi$  satisfying Assumption 2, we know  $\phi'(v) \rightarrow 0$  as  $v \rightarrow \pm\infty$ , thereafter  $xf'(x) \rightarrow 0$  as  $x \rightarrow \infty$ , that is,  $f'(x) = o(\frac{1}{x})$ . By rewriting equation (29), we have

$$\left(\frac{1}{x}\right)^{\frac{\gamma+1}{2}} f'\left(\frac{1}{x}\right) = x^{\frac{\gamma+1}{2}} f'(x). \quad (30)$$

Let  $G(x) := x^{\frac{\gamma+1}{2}} f'(x)$ , from (30), we have  $G(\frac{1}{x}) = G(x)$  and  $\lim_{x \rightarrow \infty} G(x) = \lim_{x \rightarrow \infty} G(\frac{1}{x}) = \lim_{x \rightarrow \infty} (\frac{1}{x})^{\frac{\gamma+1}{2}} f'(\frac{1}{x}) = 0$  provided  $f'(0) < \infty$  and  $f'(x)$  is continuous at 0. One such choice is

$$G(x) = \frac{x}{(1+x)^2}.$$

Then by integrating  $f'(x)$  and substitution of parameter, we get one solution to (17) is

$$\phi(v) = \frac{c}{(1+e^v)^\gamma},$$

for any positive constant  $c$ . The numerator  $2^\gamma$  in (18) is just introduced to make the function  $\phi$  an upper bound of the 0 – 1 loss and  $\phi(0) = 1$ .

## A.6 Proof of Lemma 3

(i) and (ii) are easy to verify.

For (iii), given any  $\alpha \in (0, 1)$ , let  $D(v) = \frac{\alpha 2^\gamma}{(1+e^{av})^\gamma} + \frac{(1-\alpha)2^\gamma}{(1+e^{-av})^\gamma}$ . Let  $D'(v) = 0$ , we have the only solution  $v^* = \frac{1}{a(\gamma-1)} \log \frac{\alpha}{1-\alpha}$ . Note that  $D'(v) = (-\gamma a) \left( \frac{\alpha 2^\gamma}{(1+e^{av})^\gamma} \frac{1}{1+e^{-av}} + \frac{(1-\alpha)2^\gamma}{(1+e^{-av})^\gamma} \frac{-1}{1+e^{av}} \right) = (-\gamma a) \left( \alpha \phi_{a,\gamma}(v) \frac{1}{1+e^{-av}} + (1-\alpha) \phi_{a,\gamma}(-v) \frac{-1}{1+e^{av}} \right)$ . Since  $\phi_{a,\gamma}$  is decreasing, when  $v > v^*$ , we have  $\left( \alpha \phi_{a,\gamma}(v) \frac{1}{1+e^{-av}} + (1-\alpha) \phi_{a,\gamma}(-v) \frac{-1}{1+e^{av}} \right) < 0$ , and hence  $D'(v) > 0$ . Similarly, when  $v < v^*$ ,  $D'(v) < 0$ . Therefore,  $v^*$  is indeed a global minimum point.

For (iv), let  $C_\eta(\alpha) = 2^\gamma \left( \eta \frac{1}{(1+e^\alpha)^\gamma} + (1-\eta) \frac{1}{(1+e^{-\alpha})^\gamma} \right)$ . Then by setting  $C'_\eta(\alpha) = 0$ , we get the minimum point  $\alpha^* = \frac{1}{\gamma-1} \log \frac{\eta}{1-\eta}$ , and  $C_\eta(\alpha^*) = \eta \left( \frac{2}{(1+\frac{2}{1-\eta})^\gamma} \right)^\gamma + (1-\eta) \left( \frac{2}{(1+\frac{2}{\eta})^\gamma} \right)^\gamma$  which can be shown to attain the global maximum when  $\gamma = \frac{1}{2}$ , and  $C_{1/2}(\alpha^*) = 1$ . We also have  $\inf_{\alpha: \alpha(2\eta-1) \leq 0} C_\eta(\alpha) = C_\eta(0) = 1 > C_\eta(\alpha^*)$  when  $\eta \neq \frac{1}{2}$ . By Bartlett et al. (2006), we have (iv) holds.

## A.7 Proof of Theorem 1

The proof is completed by showing that at each iteration  $t$ , as long as the empirical margin  $\widehat{\mu}(w_t, h_t)$  is positive, the empirical risk decreases by adding the weak hypotheses  $h_t$  to the current estimate. Then, we show that the weak hypothesis returned by our Arch Boost algorithm, always has a positive empirical margin before convergence.

(i) On the sample  $\mathcal{S}_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ , at each iteration  $t$ , denote

$$\mathbf{F}_{t-1} = (F_{t-1}(x_1), \dots, F_{t-1}(x_n)).$$

Recall that the empirical risk  $\widehat{R}_{\phi,n}(F) = \frac{1}{n} \sum_{i=1}^n \phi(Y_i F(X_i))$  and it can be viewed as a multivariate function of  $\mathbf{F} = (F(X_1), \dots, F(X_n))$ . Denote the partial derivative w.r.t.  $F(X_i)$  at iteration  $t$  as

$$g_t(X_i) = \left[ \frac{\partial \widehat{R}_{\phi,n}(\mathbf{F})}{\partial F(X_i)} \right]_{F(X_i)=F_{t-1}(X_i)} = \frac{1}{n} Y_i \phi'(Y_i F_{t-1}(X_i)).$$

Then the gradient of  $\widehat{R}_{\phi,n}$  at  $\mathbf{F}_{t-1}$  is

$$\nabla \widehat{R}_{\phi,n}(\mathbf{F}_{t-1}) = \frac{1}{n} \begin{pmatrix} g_t(X_1) \\ \vdots \\ g_t(X_n) \end{pmatrix}.$$



At iteration  $t$ , the weight on  $(X_i, Y_i)$  is updated to be

$$w_t(X_i, Y_i) = -\phi'(Y_i F_{t-1}(X_i)), \quad i = 1, \dots, n.$$

Suppose we choose a weak hypothesis  $h_t$  with positive empirical margin w.r.t. weights  $w_t$ , that is,  $\widehat{\mu}(h_t, w_t) > 0$ , and denote  $\mathbf{h}_t = (h_t(X_1), \dots, h_t(X_n))$ . Note that

$$\begin{aligned} \langle -\nabla \widehat{R}_{\phi, n}(\mathbf{F}_{t-1}), \mathbf{h}_t \rangle &= \frac{1}{n} \sum_{i=1}^n Y_i h_t(X_i) (-\phi'(Y_i F_{t-1}(X_i))) \\ &= \frac{1}{n} \sum_{i=1}^n Y_i h_t(X_i) w_t(X_i, Y_i) \\ &= \widehat{\mu}(h_t, w_t) > 0, \end{aligned}$$

where  $\langle \cdot, \cdot \rangle$  is the standard inner product in  $\mathbb{R}^n$ . Therefore, we know

$$\langle -\nabla \widehat{R}_{\phi, n}(\mathbf{F}_{t-1}), \mathbf{h}_t \rangle > 0 \iff \widehat{\mu}(w_t, h_t) > 0.$$

But if  $\langle -\nabla \widehat{R}_{\phi, n}(\mathbf{F}_{t-1}), \mathbf{h}_t \rangle > 0$ , then  $\begin{pmatrix} h_t(X_1) \\ \vdots \\ h_t(X_n) \end{pmatrix}$  is a descending direction of  $\widehat{R}_{\phi, n}(\mathbf{F})$  at  $\mathbf{F}_{t-1}$ ,

therefore

$$\widehat{R}_{\phi, n} \left[ \begin{pmatrix} F_t(X_1) \\ \vdots \\ F_t(X_n) \end{pmatrix} \right] = \widehat{R}_{\phi, n} \left[ \begin{pmatrix} F_{t-1}(X_1) \\ \vdots \\ F_{t-1}(X_n) \end{pmatrix} + \alpha_t \begin{pmatrix} h_t(X_1) \\ \vdots \\ h_t(X_n) \end{pmatrix} \right] < \widehat{R}_{\phi, n} \left[ \begin{pmatrix} F_{t-1}(X_1) \\ \vdots \\ F_{t-1}(X_n) \end{pmatrix} \right]$$

with an appropriate step size  $\alpha_t$  which can be found by line search

$$\alpha_t = \underset{\alpha}{\operatorname{argmin}} \widehat{R}_{\phi, n} \left[ \begin{pmatrix} F_{t-1}(X_1) \\ \vdots \\ F_{t-1}(X_n) \end{pmatrix} + \alpha \begin{pmatrix} h_t(X_1) \\ \vdots \\ h_t(X_n) \end{pmatrix} \right].$$

In summary, we have

$$\widehat{R}_{\phi, n}(F_t) < \widehat{R}_{\phi, n}(F_{t-1}) \tag{31}$$

if at step  $t$ , we choose a base learner  $h_t$  such that  $\widehat{\mu}(w_t, h_t) > 0$  and choose a suitable step size  $\alpha_t$  either by line search or set to be appropriately small. Therefore,  $\widehat{R}_{\phi, n}$  will converge in  $\mathbb{R}$ .

(ii) In any region  $R_t^j$ , we know  $h_t \equiv \gamma_t^j$ . Then  $\begin{pmatrix} -g_t(X_1) \\ \vdots \\ -g_t(X_N) \end{pmatrix} \begin{pmatrix} h_t(X_1) \\ \vdots \\ h_t(X_N) \end{pmatrix}$  is equal to

$$\begin{aligned}
& \sum_{j=1}^{J_t} \sum_{i \in R_t^j} Y_i w_t(X_i, Y_i) \gamma_t^j \\
&= \sum_{j=1}^{J_t} \gamma_t^j \left( \mathbb{P}_{w_t}(Y = 1 | X \in R_t^j) - \mathbb{P}_{w_t}(Y = -1 | X \in R_t^j) \right) \sum_{i \in R_t^j} w_t(X_i, Y_i) \\
&= \sum_{j=1}^{J_t} \theta(\mathbb{P}_{w_t}(Y = 1 | X \in R_t^j)) \left( 2\mathbb{P}_{w_t}(Y = 1 | X \in R_t^j) - 1 \right) \sum_{i \in R_t^j} w_t(X_i, Y_i) \\
&\stackrel{(i)}{\geq} 0.
\end{aligned}$$

The last inequality (i) is because  $\theta(\mathbb{P}_{w_t}(Y = 1 | X \in R_t^j))$  is strictly increasing and has the only root at  $\frac{1}{2}$ , and hence always has the same sign as  $2\mathbb{P}_{w_t}(Y = 1 | X \in R_t^j) - 1$ , and “=” holds if and only if  $\mathbb{P}_{w_t}(Y = 1 | X \in R_t^j) = \frac{1}{2}$  for all  $j = 1, \dots, J_t$ .

(iii) From (13), we have

$$\frac{\phi'(-F_t(x))}{\phi'(F_t(x))} = \frac{\mathbb{P}_{w_t}(Y = 1|x)}{\mathbb{P}_{w_t}(Y = -1|x)} \frac{\phi'(-F_{t-1}(x))}{\phi'(F_{t-1}(x))}.$$

If  $\mathbb{P}_{w_t}(Y = 1|x) > \mathbb{P}_{w_t}(Y = -1|x)$ , then  $\frac{\phi'(-F_t(x))}{\phi'(F_t(x))} > \frac{\phi'(-F_{t-1}(x))}{\phi'(F_{t-1}(x))}$ . By Lemma 2,  $F_t(x) > F_{t-1}(x)$ , that is,  $h_t(x) > 0$ .

(iv) Here, we develop ideas much similar to the proof of Lemma 4.1 and Lemma 4.2 in Zhang and Yu (2005). There are two differences here in comparison to Zhang and Yu (2005). First, the loss is non-convex function and second, the optimal hypothesis is chosen differently. For  $f_1, f_2 \in \cup_{T=1}^{\infty} \mathcal{F}^T$ , let  $H_f \subset \mathcal{H}$  be the set that contains all weak hypotheses in  $f_1$  and  $f_2$ . For example,  $f_1 = \sum_{h \in H_f} \alpha_1^h h$  and  $f_2 = \sum_{h \in H_f} \alpha_2^h h$ . Then denote

$$\|f_1 - f_2\|_2^2 := \sum_{h \in H_f} (\alpha_1^h - \alpha_2^h)^2 \leq \frac{1}{|H_f|} \left( \sum_{h \in H_f} |\alpha_1^h - \alpha_2^h| \right)^2 =: \frac{1}{|H_f|} \|f_1 - f_2\|_1^2.$$

Now let  $\bar{f}_t$  be any reference function in  $\cup_{T=1}^{\infty} \mathcal{F}^T$  satisfying

$$\widehat{R}_{\phi, n}(\bar{f}_t) < \inf_{f \in \cup_{T=1}^{\infty} \mathcal{F}^T} \widehat{R}_{\phi, n}(f) + \frac{1}{t}$$

and  $F_t$  be the classifier returned by Arch Boost at step  $t$ . Moreover, denote

$$\bar{f}_t = \sum_{h \in H_t} \omega_t^h h, \quad F_t = \sum_{h \in H_t} \alpha_t^h h.$$

For notation simplicity, we denote  $R = \widehat{R}_{\phi, n}$  since we have fixed a loss function  $\phi$  and sample size  $n$ . Let  $s^h = \text{sign}(\omega_t^h - \alpha_t^h)$ . By Taylor expansion, we have

$$R(F_t + \alpha_{t+1} s^h h) \leq R(F_t) + \alpha_{t+1} s^h \langle \nabla R(F_t), h \rangle + \frac{\alpha_{t+1}^2}{2} \sup_{\xi \in [0, 1]} R''_{F_t, h}(\xi \alpha_{t+1} s^h),$$

where  $R_{F_t, h}(\alpha) := R(F_t + \alpha h)$ . Since  $\theta$  is bounded, from part (ii) we know there exists  $M > 0$  s.t.  $\sup_{\xi \in [0, 1]} R''_{F_t, h}(\xi \alpha_{t+1} s^h) < M$  if we choose  $h$  by (13). Therefore,

$$R(F_t + \alpha_{t+1} s^h h) \leq R(F_t) + \alpha_{t+1} s^h \langle \nabla R(F_t), h \rangle + \frac{\alpha_{t+1}^2}{2} M.$$

By Algorithm 1 we know that  $R(F_{t+1}) = R(F_t + \alpha_{t+1} h_{t+1})$ . Moreover, by (13),  $h_{t+1}$  is chosen as the  $\arg \min_{h \in \mathcal{H}_t} \mathbb{E}_w [R(F_t + \alpha_{t+1} h)]$ . Hence, for any  $h \in \mathcal{H}_t$ ,  $\mathbb{E}_w [R(F_t + \alpha_{t+1} h_{t+1})] \leq \mathbb{E}_w [R(F_t + \alpha_{t+1} h)]$ . Moreover, for any bounded random variable  $Z$ ,  $|\mathbb{E}_w [Z] - \mathbb{E}[Z]| \leq K$  for a positive constant  $K$ . Combining the above, we have

$$R(F_{t+1}) \leq R(F_t + \alpha_{t+1} s^h h) + 2\epsilon_t + 2K,$$

for

$$\epsilon_t = \sup_{h \in \mathcal{H}_t} \left| R(F_t + \alpha_{t+1} s^h h) - \mathbb{E} \left[ R(F_t + \alpha_{t+1} s^h h) \right] \right|.$$

By the arguments very much similar to Lemmas 6 and 7, it easy to obtain  $\epsilon_t = o_P(1)$ .

Since  $\|\bar{f}_t - F_t\|_1 = o(\log t)$ , and  $\|\bar{f}_t - F_T\|_2^2 \leq \frac{\|\bar{f}_t - F_t\|_1^2}{t^{c_t}}$  where  $c_t \in (0, 1)$  and  $c_t \rightarrow 0$  as  $t \rightarrow \infty$ , we have  $\frac{\|\bar{f}_t - F_t\|_1^2}{t^{c_t}} = o\left(\frac{\log t}{t^{c_t}} \|\bar{f}_t - F_t\|_1\right)$ . Hence,

$$\begin{aligned} & \|\bar{f}_t - F_t\|_2^2 (R(F_{t+1}) - 2\epsilon_t - 2K) \\ &= o \left[ \frac{\log t}{t^{c_t}} \sum_{h \in H_t} |\alpha_t^h - \omega_t^h| R(F_t + \alpha_{t+1} s^h h) \right] \\ &= o \left[ \frac{\log t}{t^{c_t}} \sum_{h \in H_t} |\alpha_t^h - \omega_t^h| \left( R(F_t) + \alpha_{t+1} s^h \langle \nabla R(F_t), h \rangle + \frac{\alpha_{t+1}^2}{2} M \right) \right] \\ &= o \left[ \frac{\log t}{t^{c_t}} \|\bar{f}_t - F_t\|_1 R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t}} \langle \nabla R(F_t), \bar{f}_t - F_t \rangle + \frac{M \alpha_{t+1}^2 \log t}{2 t^{c_t}} \|\bar{f}_t - F_t\|_1 \right] \end{aligned} \quad (32)$$

Now we look at the situation when  $\hat{\mu}(h_k, w_k) = 0$ . From part (ii), we know this happens if and only if  $\mathbb{P}_{w_k}(Y = 1 | X \in R_k^j) = \frac{1}{2}$  in every region  $j$ . In another word,  $\nabla R(F_k) \perp \mathcal{H}$ . Now since  $\hat{\mu}(h_t, w_t) \rightarrow 0$ ,  $\nabla R(F_t)$  is more and more perpendicular to  $\mathcal{H}$  and hence perpendicular to  $\cup_{T=1}^{\infty} \mathcal{F}^T$ , and  $\langle \nabla R(F_t) - \nabla R(\bar{f}_t), \bar{f}_t - F_t \rangle \rightarrow 0$  since  $\bar{f}_t - F_t \in \cup_{T=1}^{\infty} \mathcal{F}^T$ .

Since  $\phi$  is Lipschitz differentiable, we know there exists  $L > 0$  ( $L$  is the Lipschitz constant s.t.  $\|\nabla R(f_1) - \nabla R(f_2)\|_2 \leq L \|f_1 - f_2\|_2$  for all  $f_1$  and  $f_2$ ) s.t.

$$R(F_t) - R(\bar{f}_t) \leq \langle \nabla R(\bar{f}_t), F_t - \bar{f}_t \rangle + \frac{L}{2} \|\bar{f}_t - F_t\|_2^2.$$

Then  $\langle \nabla R(\bar{f}_t), \bar{f}_t - F_t \rangle \leq R(\bar{f}_t) - R(F_t) + \frac{L}{2} \|\bar{f}_t - F_t\|_2^2$ . When  $t$  is large enough, we know there exists sequence  $\tilde{\epsilon}_t \rightarrow 0$  s.t.

$$\langle \nabla R(F_t), \bar{f}_t - F_t \rangle \leq R(\bar{f}_t) - R(F_t) + \frac{L}{2} \|\bar{f}_t - F_t\|_2^2 + \tilde{\epsilon}_t.$$

Then by (32),

$$\begin{aligned}
& \|\bar{f}_t - F_t\|_2^2 (R(F_{t+1}) - 2\epsilon_t - 2K) \\
&= o \left[ \frac{\log t}{t^{c_t}} \|\bar{f}_t - F_t\|_1 R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t}} \langle \nabla R(F_t), \bar{f}_t - F_t \rangle + \frac{\alpha_{t+1}^2 \log t}{2t^{c_t}} \|\bar{f}_t - F_t\|_1 M \right] \\
&= o \left[ \frac{\log t}{t^{c_t}} \|\bar{f}_t - F_t\|_1 R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t}} \left( R(\bar{f}_t) - R(F_t) + \frac{L}{2} \|\bar{f}_t - F_t\|_2^2 + \tilde{\epsilon}_t \right) + \frac{\alpha_{t+1}^2 \log t}{2t^{c_t}} \|\bar{f}_t - F_t\|_1 M \right] \\
&= o \left[ \frac{\log t}{t^{c_t}} \|\bar{f}_t - F_t\|_1 R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t}} (R(\bar{f}_t) - R(F_t)) + \eta_t \right], \tag{33}
\end{aligned}$$

where  $\eta_t := \frac{\alpha_{t+1} \log t}{t^{c_t}} \left( \frac{L}{2} \|\bar{f}_t - F_t\|_2^2 + \tilde{\epsilon}_t \right) + \frac{\alpha_{t+1}^2 \log t}{2t^{c_t}} \|\bar{f}_t - F_t\|_1 M$ .

Then by dividing  $\|\bar{f}_t - F_t\|_2^2$  on both sides of (33), we get

$$\begin{aligned}
R(F_{t+1}) &= o \left[ \frac{\log t}{t^{c_t}} \frac{\|\bar{f}_t - F_t\|_1}{\|\bar{f}_t - F_t\|_2^2} R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t} \|\bar{f}_t - F_t\|_2^2} (R(\bar{f}_t) - R(F_t)) + \bar{\eta}_t + 2\epsilon_t + 2K \right] \\
&= o \left[ \frac{\log t}{t^{c_t/2} \|\bar{f}_t - F_t\|_2} R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t} \|\bar{f}_t - F_t\|_2^2} (R(\bar{f}_t) - R(F_t)) + \bar{\eta}_t + 2\epsilon_t + 2K \right],
\end{aligned}$$

where  $\bar{\eta}_t := \frac{\alpha_{t+1} \log t}{t^{c_t}} \left( \frac{L}{2} + \frac{\tilde{\epsilon}_t}{\|\bar{f}_t - F_t\|_2^2} \right) + \frac{\alpha_{t+1}^2 \log t}{2t^{c_t/2} \|\bar{f}_t - F_t\|_2} M$ .

Therefore,

$$\begin{aligned}
R(F_{t+1}) - R(\bar{f}_t) &= o \left[ \frac{\log t}{t^{c_t} \|\bar{f}_t - F_t\|_2} R(F_t) + \frac{\alpha_{t+1} \log t}{t^{c_t} \|\bar{f}_t - F_t\|_2^2} (R(\bar{f}_t) - R(F_t)) + \bar{\eta}_t + 2\epsilon_t + 2K \right] \\
&\leq \frac{\xi_t \log t}{t^{c_t} \|\bar{f}_t - F_t\|_2} R(F_t) + \frac{\alpha_{t+1} \xi_t \log t}{t^{c_t} \|\bar{f}_t - F_t\|_2^2} (R(\bar{f}_t) - R(F_t)) + \xi_t \bar{\eta}_t + 2\xi_t \epsilon_t + 2K \xi_t,
\end{aligned}$$

for some sequence  $\xi_t \rightarrow 0$  as  $t \rightarrow \infty$ . Now if we assume  $c_t \rightarrow 0$  slowly enough, then by choosing  $\alpha_t$  s.t.  $\sum_{t=1}^{\infty} \alpha_t = \infty$ ,  $\sum_{t=1}^{\infty} \alpha_t^2 < \infty$  and  $\sum_{t=1}^{\infty} \frac{\alpha_{t+1} \xi_t \log t}{t^{c_t}} < \infty$ , and by Lemma 4.2 in Zhang and Yu (2005), we have  $R(F_{t+1}) - R(\bar{f}_t) \rightarrow 0$  as  $t \rightarrow \infty$ , and hence  $R(F_t) \rightarrow R_{\phi, n}^*$  as  $t \rightarrow \infty$ .

## A.8 Proof of Theorem 2

The proof of Theorem consists of a careful decomposition of the inner product between the gradient vector and the weak hypothesis obtained on the complete data. The decomposition is done in such a manner that one of the factors is the inner product between the gradient computed on the noise-free data and the weak hypothesis. Then, the proof is completed by showing that the signs of the two inner products above match.

On the original dataset  $\mathcal{S}$ , suppose after some iteration we obtain a weak hypothesis  $h$  such that  $h(x) = \theta(\mathbb{P}_w(Y = 1|x))$  where  $\theta$  is defined in Theorem 1. In the rest of the proof we will exploit the decomposition proved in Theorem 1,

$$\begin{aligned}
-\langle \mathbf{g}_o, \mathbf{h} \rangle &= \sum_{j=1}^J \theta(\mathbb{P}_w(Y = 1|X \in R^j)) (2\mathbb{P}_w(Y = 1|X \in R^j) - 1) \sum_{i \in R^j} w(x_i, y_i) \\
&= \sum_{j=1}^J \theta(p_j) (2p_j - 1) \sum_{i \in R^j} w(x_i, y_i)
\end{aligned}$$

where  $R^j, j = 1, \dots, J$  are the regions corresponding to terminal nodes, and as long as  $p_j = \mathbb{P}_w(Y = 1|X \in R^j)$  is not  $\frac{1}{2}$  in some region,  $-\langle \mathbf{g}, \mathbf{h} \rangle > 0$ , that is,  $\mathbf{h}$  points to the descending direction of the empirical risk.

For  $\mathbf{g}_o$ , we get  $-\langle \mathbf{g}_o, \mathbf{h} \rangle = \sum_{j=1}^J \theta(p_j) (2\tilde{p}_j - 1) \sum_{i \in R^j \setminus \mathcal{O}} w(x_i, y_i)$  with a different conditional probability estimation  $\tilde{p}_j$ . Recall that

$$2p_j - 1 = \mathbb{P}_w(Y = 1|X \in R^j) - \mathbb{P}_w(Y = -1|X \in R^j) = \frac{\sum_{i \in R^j} y_i w(x_i, y_i)}{\sum_{i \in R^j} w(x_i, y_i)},$$

$$2\tilde{p}_j - 1 = \frac{\sum_{i \in R^j \setminus \mathcal{O}} y_i w(x_i, y_i)}{\sum_{i \in R^j \setminus \mathcal{O}} w(x_i, y_i)}.$$

Then we have that  $(2\tilde{p}_j - 1) \sum_{i \in R^j \setminus \mathcal{O}} w(x_i, y_i) = \sum_{i \in R^j \setminus \mathcal{O}} y_i w(x_i, y_i) = (2p_j - 1) \sum_{i \in R^j} w(x_i, y_i) - \sum_{i \in \mathcal{O} \cap R^j} y_i w(x_i, y_i)$ . Therefore,

$$\begin{aligned} -\langle \mathbf{g}_o, \mathbf{h} \rangle &= -\langle \mathbf{g}, \mathbf{h} \rangle - \sum_{j=1}^J \theta(p_j) \sum_{i \in \mathcal{O} \cap R^j} y_i w(x_i, y_i) \\ &= \sum_{j=1}^J \theta(p_j) \left[ (2p_j - 1) \sum_{i \in R^j} w(x_i, y_i) - \sum_{i \in \mathcal{O} \cap R^j} y_i w(x_i, y_i) \right]. \end{aligned}$$

From previous equation, a sufficient condition for  $-\langle \mathbf{g}_o, \mathbf{h} \rangle \geq 0$  is that each summand reminds non-negative. Since  $\theta(p_j) \geq 0$  if and only if  $p_j \geq \frac{1}{2}$ , the sufficient condition becomes equivalent to

$$(p_j - \frac{1}{2}) \left[ (2p_j - 1) \sum_{i \in R^j} w(x_i, y_i) - \sum_{i \in \mathcal{O} \cap R^j} y_i w(x_i, y_i) \right] \geq 0.$$

Furthermore, this inequality can be reformulated as

$$2(p_j - \frac{1}{2})^2 \sum_{i \in R^j \setminus \mathcal{O}} w(x_i, y_i) \geq (p_j - \frac{1}{2}) \sum_{i \in \mathcal{O} \cap R^j} (1 + y_i - 2p_j) w(x_i, y_i).$$

When  $p_j \neq \frac{1}{2}$ , a sufficient condition for the inequality above is

$$\begin{aligned} 2|p_j - \frac{1}{2}| \sum_{i \in R^j \setminus \mathcal{O}} w(x_i, y_i) &\geq \sum_{i \in \mathcal{O} \cap R^j} \max_{y_i = \pm 1} \left[ \text{sign}(p_j - \frac{1}{2})(1 + y_i - 2p_j) \right] w(x_i, y_i) \\ &= \sum_{i \in \mathcal{O} \cap R^j} 2 \min(p_j, 1 - p_j) w(x_i, y_i). \end{aligned}$$

## A.9 Proof of Lemma 4

Since  $\phi$  is an Arch boosting loss function, we know for any  $p \in (0, 1)$ ,  $p\phi(v) + (1-p)\phi(-v)$  has only one critical point  $v^*$  that is the global minimum. Hence,  $p_x \phi''(F^*(x)) + (1-p_x) \phi''(-F^*(x)) \geq 0$  where  $p_x = \mathbb{P}(Y = 1|X = x)$ . Note that here we treat  $Yf(X)$  as the input of  $\phi$ , and by the chain rule, we know  $p_x \phi''(1, F^*(X)) + (1-p_x) \phi''(-1, F^*(X)) \geq 0$ , where the derivative now is w.r.t. the second argument. Then  $\mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, F^*(X)) q^2(X) | X = x \right] = q^2(x) [\mathbb{P}(Y = 1|X = x) \phi''(1, F^*(x)) + \mathbb{P}(Y = -1|X = x) \phi''(-1, F^*(x))] \geq 0$  since  $\mathbb{P}(Y = 1|X = x) \phi''(1, F^*(x)) + \mathbb{P}(Y = -1|X = x) \phi''(-1, F^*(x)) \geq 0$  for all  $\mathbb{P}(Y = 1|X = x) \in (0, 1)$  and is also nonnegative at the end points  $\{0, 1\}$  because of the continuity of  $\phi''$ .

Now if  $\mathbb{P}$  and  $\mathcal{X}$  satisfy that  $\mathbb{P}(Y = 1|X = x) = p_x \in [\delta, 1 - \delta]$  for all  $x \in \mathcal{X}$  for some  $\delta \in (0, \frac{1}{2})$ , and for all  $p \in [\delta, 1 - \delta]$ ,  $p\phi''(v^*) + (1-p)\phi''(-v^*) > 0$ , that is, it is locally convex near  $v^*$ , then by

the continuity of  $\phi''$ , for each  $p \in [\delta, 1 - \delta]$ , there exists  $r_p > 0$  such that  $p\phi''(v) + (1-p)\phi''(-v) \geq 0$  for all  $v \in \mathbb{R}$  such that  $|v - v^*| < r_p$ . But since  $[\delta, 1 - \delta]$  is compact, there exists  $r > 0$  such that  $p\phi''(v) + (1-p)\phi''(-v) \geq 0$  for all  $|v - v^*| < r$  and for all  $p \in [\delta, 1 - \delta]$ .

Now for any  $x \in \mathcal{X}$ , we know the corresponding  $p_x \in [\delta, 1 - \delta]$ , and for this  $p_x$ , we have  $p_x\phi''(G(x)) + (1-p_x)\phi''(-G(x)) \geq 0$  for all measurable function  $G$  with  $|G(x) - F^*(x)| < r$ . Therefore, if we take measurable function  $G$  s.t.  $\|G - F^*\|_\infty < r$ , then  $p_x\phi''(G(x)) + (1-p_x)\phi''(-G(x)) \geq 0$  for all  $x \in \mathcal{X}$ , that is,  $\mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, G(X))q^2(X) | X = x \right] \geq 0$  for all  $x \in \mathcal{X}$ . Taking expectation w.r.t.  $X$ , we get  $\mathbb{E}_{\mathbb{P}} \left[ \phi''(Y, G(X))q^2(X) \right] \geq 0$ .

### A.10 Proof of Theorem 3

By Lemma 4, there exists  $r > 0$  such that if  $\|f_{\mathbb{P},\lambda} - F^*\| < r$ , then  $\mathbb{E}_{\mathbb{P}}\phi''(Y, f_{\mathbb{P},\lambda}(X))q^2(X) \geq 0$  for any measurable function  $q$ . Therefore one can show that the influence function still exists in form of (24) by the Theorem 4 in Christmann and Steinwart (2004) since in that Theorem, the only place one needs convexity is to show  $\mathbb{E}_{\mathbb{P}}\phi''(Y, f_{\mathbb{P},\lambda}(X))q^2(X) \geq 0$ . Rewriting (24), and let  $IF(z; T, \mathbb{P}) = g_z \in H$ , we have

$$2\lambda g_z + \mathbb{E}_{\mathbb{P}}\phi''(Y, f_{\mathbb{P},\lambda}(X))g_z(X)\Psi(X) = \mathbb{E}_{\mathbb{P}}\phi'(Y, f_{\mathbb{P},\lambda}(X))\Psi(X) - \phi'(z_y, f_{\mathbb{P},\lambda}(z_x))\Psi(z_x).$$

By taking inner product  $\langle \cdot, \cdot \rangle_H$  with  $g_z$  itself, we have

$$2\lambda \|g_z\|_H^2 + \mathbb{E}_{\mathbb{P}}\phi''(Y, f_{\mathbb{P},\lambda}(X))g_z^2(X) = \mathbb{E}_{\mathbb{P}}\phi'(Y, f_{\mathbb{P},\lambda}(X))g_z(X) - \phi'(z_y, f_{\mathbb{P},\lambda}(z_x))g_z(z_x). \quad (34)$$

Assume  $\lambda\|f\|_H^2 + R_\phi(f)$  has only one global minimum  $f_{\mathbb{P},\lambda}$  in  $H$ , then at the minimum, we know the Frechet derivative at  $f_{\mathbb{P},\lambda}$  is a zero mapping, that is,

$$2\lambda \langle f_{\mathbb{P},\lambda}, \cdot \rangle_H + \mathbb{E}_{\mathbb{P}}\phi'(Y, f_{\mathbb{P},\lambda}(X))\Psi(X) = 0,$$

where  $0 : H \rightarrow \mathbb{R}$  is the zero functional. And hence

$$2\lambda \langle f_{\mathbb{P},\lambda}, g_z \rangle_H + \mathbb{E}_{\mathbb{P}}\phi'(Y, f_{\mathbb{P},\lambda}(X))g_z(X) = 0. \quad (35)$$

We also note that since  $f_{\mathbb{P},\lambda}$  is the global minimum, then  $\lambda\|f_{\mathbb{P},\lambda}\|_H^2 + R_\phi(f_{\mathbb{P},\lambda}) \leq \lambda\|0_H\|_H^2 + R_\phi(0_H) = C_\phi$  where  $C_\phi = R_\phi(0_H) = \phi(0, 0)$  is a constant, that is,

$$\lambda\|f_{\mathbb{P},\lambda}\|_H^2 \leq \lambda\|f_{\mathbb{P},\lambda}\|_H^2 + \mathbb{E}_{\mathbb{P}}\phi(Y, f_{\mathbb{P},\lambda}(X)) \leq C_\phi. \quad (36)$$

Finally, we have

$$\begin{aligned} 2\lambda\|g_z\|_H^2 &\leq 2\lambda\|g_z\|_H^2 + \mathbb{E}_{\mathbb{P}}\phi''(Y, f_{\mathbb{P},\lambda}(X))g_z^2(X) \\ &\stackrel{(i)}{=} \mathbb{E}_{\mathbb{P}}\phi'(Y, f_{\mathbb{P},\lambda}(X))g_z(X) - \phi'(z_y, f_{\mathbb{P},\lambda}(z_x))g_z(z_x) \\ &\stackrel{(ii)}{=} -2\lambda \langle f_{\mathbb{P},\lambda}, g_z \rangle_H - \phi'(z_y, f_{\mathbb{P},\lambda}(z_x))g_z(z_x) \\ &\stackrel{(iii)}{\leq} 2\lambda\|f_{\mathbb{P},\lambda}\|_H\|g_z\|_H - \phi'(z_y, f_{\mathbb{P},\lambda}(z_x))g_z(z_x) \\ &\stackrel{(iv)}{\leq} 2\sqrt{\lambda C_\phi}\|g_z\|_H + |\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))||g_z(z_x)| \\ &= 2\sqrt{\lambda C_\phi}\|g_z\|_H + |\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))|\langle g_z, k(z_x, \cdot) \rangle_H \\ &\stackrel{(v)}{\leq} 2\sqrt{\lambda C_\phi}\|g_z\|_H + |\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))|\sqrt{\langle g_z, g_z \rangle_H} \sqrt{\langle k(z_x, \cdot), k(z_x, \cdot) \rangle_H} \\ &= 2\sqrt{\lambda C_\phi}\|g_z\|_H + |\phi'(z_y, f_{\mathbb{P},\lambda}(z_x))|\|g_z\|_H |k(z_x, z_x)|. \end{aligned}$$

where (i) is due to (34); (ii) due to (35); (iii) is due to the Cauchy-Schwartz inequality; (iv) is due to (36); (v) is again due to the Cauchy-Schwartz inequality.

Since  $k$  is a bounded kernel, there exists a  $M_k > 0$  such that  $|k(x_1, x_2)| \leq M_k$  for all  $x_1, x_2 \in \mathcal{X}$ . Hence,  $|\phi'(z_y, f_{\mathbb{P}, \lambda}(z_x))| \|g_z\|_H |k(z_x, z_x)| \leq M_k |\phi'(z_y, f_{\mathbb{P}, \lambda}(z_x))| \|g_z\|_H$ , which in turn leads to

$$2\lambda \|g_z\|_H^2 \leq 2\sqrt{\lambda C_\phi} \|g_z\|_H + M_k |\phi'(z_y, f_{\mathbb{P}, \lambda}(z_x))| \|g_z\|_H$$

and hence

$$\|g_z\|_H \leq \sqrt{\frac{C_\phi}{\lambda}} + \frac{M_k |\phi'(z_y, f_{\mathbb{P}, \lambda}(z_x))|}{2\lambda}$$

Hence, for small  $\lambda$ , we obtain  $\|g_z\|_H = O(\lambda^{-1/2})$ , whereas for large lambda it is  $\|g_z\|_H = O(\lambda^{-1})$ .

## A.11 Proof of Theorem 4

First, we prove Theorem 4 (a). The result follows from the following Lemma 6, which builds on the Lemma 4 in Bartlett and Traskin (2007) supplemented by the uniform convergence of the truncated loss to  $\phi$ -loss presented in Lemma 7.

**Lemma 6.** *For an Arch boosting loss function  $\phi$  satisfying Assumption 2, let  $L_\phi$  be the Lipschitz constant of  $\phi$ , that is  $L_\phi = \inf\{L > 0 : |\phi(x) - \phi(y)| \leq L|x - y|, x, y \in \mathbb{R}\}$ , and  $M_\phi$  be the maximum value of  $\phi$ , i.e.  $M_\phi = \sup_{x \in \mathbb{R}} \phi(x)$ . Then for  $V = d_{VC}(\mathcal{H})$ ,  $c = 24 \int_0^1 \sqrt{\log \frac{8e}{\mu^2}} d\mu$  and sequences  $T_n \rightarrow \infty$  and  $\zeta_n \rightarrow \infty$  as  $n \rightarrow \infty$ , and  $\delta_n \rightarrow 0$  as  $n \rightarrow \infty$ , with probability at least  $1 - \delta_n$ ,*

$$\sup_{f \in \pi_{\zeta_n} \circ \mathcal{F}^{T_n}} |\widehat{R}_{\phi, n}(f) - R_\phi(f)| \leq c\zeta_n L_\phi \sqrt{\frac{(V+1)(T_n+1) \log_2\left(\frac{2(T_n+1)}{\log 2}\right)}{n}} + M_\phi \sqrt{\frac{\log \frac{1}{\delta_n}}{2n}}, \quad (37)$$

where the truncation function  $\pi_l(\cdot)$  is defined as  $\pi_l(x) = l\mathbb{1}\{x > l\} + x\mathbb{1}\{|x| \leq l\} - l\mathbb{1}\{x < -l\}$ . Furthermore, if we choose  $\delta_n$ ,  $T_n$  and  $\zeta_n$  such that  $\sum_{n=1}^{\infty} \delta_n < \infty$ , and right hand side of (37) converges to zero as  $n \rightarrow \infty$ , then  $\sup_{f \in \pi_{\zeta_n} \circ \mathcal{F}^{T_n}} |\widehat{R}_{\phi, n}(f) - R_\phi(f)| \rightarrow 0$  a.s. as  $n \rightarrow \infty$ .

*Proof of Lemma 6.* The result mainly follows from Lemma 3 and 4 in Bartlett and Traskin (2007). Since  $\phi$  satisfies Assumption 2, we know  $L_\phi$  and  $M_\phi$  both exist and are finite, and  $\max_{v \in [-\zeta_n, \zeta_n]} |\phi(v)| = M_\phi$  for all  $n$ . The almost surely convergence follows from Borel-Cantelli Lemma.  $\square$

In order to make the RHS of (37) converging to zero, we can choose  $\delta_n = \frac{1}{n^2}$ ,  $\zeta_n = \kappa \log n$  and  $T_n = n^{1-\epsilon}$  with  $\kappa > 0$ ,  $\epsilon \in (0, 1)$  and  $\kappa < \frac{1}{2}\epsilon$ . Then by Borel-Cantelli lemma, since  $\sum_{n=1}^{\infty} \delta_n = \sum_{n=1}^{\infty} \frac{1}{n^2} < \infty$ , we obtain  $\sup_{f \in \pi_{\zeta_n} \circ \mathcal{F}^{T_n}} |\widehat{R}_{\phi, n}(f) - R_\phi(f)| \rightarrow 0$  a.s. as  $n \rightarrow \infty$ . In order to get rid of the truncation  $\pi_{\zeta_n}$ , we will need the next Lemma 7.

**Lemma 7.** *Let  $\phi$  be an Arch boosting loss function satisfying Assumption 2. Then for any increasing positive sequence  $\zeta_n \rightarrow \infty$ , we have*

$$\sup_{f \in \mathcal{F}^{T_n}} \left| |\widehat{R}_{\phi, n}(f) - R_\phi(f)| - |\widehat{R}_{\phi, n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)| \right| \rightarrow 0 \text{ a.s.}$$

when  $n \rightarrow \infty$ .

*Proof of Lemma 7.* First note that

$$\begin{aligned} & \left| |\widehat{R}_{\phi, n}(f) - R_\phi(f)| - |\widehat{R}_{\phi, n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)| \right| \\ & \leq \left| \widehat{R}_{\phi, n}(f) - R_\phi(f) - (\widehat{R}_{\phi, n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)) \right| \\ & \leq \left| R_\phi(f) - R_\phi(\pi_{\zeta_n} \circ f) \right| + \left| \widehat{R}_{\phi, n}(f) - \widehat{R}_{\phi, n}(\pi_{\zeta_n} \circ f) \right|. \end{aligned} \quad (38)$$

And for any Arch boosting loss function  $\phi$  satisfying Assumption 2, since  $\lim_{v \rightarrow \infty} \phi(v)$  and  $\lim_{v \rightarrow -\infty} \phi(v)$  both exist in  $\mathbb{R}$ , given any  $\epsilon > 0$ , we know there exists  $M > 0$  such that  $|\phi(u) - \phi(w)| < \epsilon/2$  for all  $|u|, |w| > M$  and  $uw > 0$  by the definition of convergence.

Now for any  $(X, Y) \in \mathcal{X} \times \mathcal{Y}$  and any  $f \in \mathcal{F}^{T_n}$ , if we choose  $\zeta > M$ , then when  $|f(X)| < \zeta$ ,  $|Y\pi_\zeta \circ f(X)| = |Yf(X)|$  and hence  $|\phi(Yf(X)) - \phi(Y\pi_\zeta \circ f(X))| = 0 < \epsilon/2$ . When  $|f(X)| \geq \zeta$ , we have  $|Y\pi_\zeta \circ f(X)| = |Y\zeta| = \zeta > M$ ,  $|Yf(X)| = |f(X)| \geq \zeta > M$ , and  $Yf(X)Y\pi_\zeta \circ f(X) = f(X)\pi_\zeta \circ f(X) > 0$ , so again  $|\phi(Yf(X)) - \phi(Y\pi_\zeta \circ f(X))| < \epsilon/2$  by the discussion in previous paragraph. In summary, we have

$$|\phi(Yf(X)) - \phi(Y\pi_\zeta \circ f(X))| < \epsilon/2 \quad (39)$$

for all  $\zeta > M$ ,  $(X, Y) \in \mathcal{X} \times \mathcal{Y}$  and  $f \in \mathcal{F}^{T_n}$ . Now in (38),

$$|R_\phi(f) - R_\phi(\pi_{\zeta_n} \circ f)| \leq \mathbb{E}_{X,Y} |\phi(Yf(X)) - \phi(Y\pi_{\zeta_n} \circ f(X))|.$$

Since  $\zeta_n$  is increasing and  $\zeta_n \rightarrow \infty$  as  $n \rightarrow \infty$ , there exists  $N \in \mathbb{N}$  such that for all  $n > N$ , we have  $\zeta_n > M$ , and hence  $\mathbb{E}_{X,Y} |\phi(Yf(X)) - \phi(Y\pi_{\zeta_n} \circ f(X))| < \epsilon/2$  for all  $n > N$  by (39). For the empirical risk part, we have

$$|\widehat{R}_{\phi,n}(f) - \widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f)| \leq \sum_{i=1}^n \frac{1}{n} |\phi(Y_i f(X_i)) - \phi(Y_i \pi_{\zeta_n} \circ f(X_i))|.$$

But for each  $i = 1, \dots, n$ , we know  $(X_i, Y_i) \in \mathcal{X} \times \mathcal{Y}$ , and hence by (39),  $|\phi(Y_i f(X_i)) - \phi(Y_i \pi_{\zeta_n} \circ f(X_i))| < \epsilon/2$  for all  $n > N$  and for all  $i$ . Therefore, the right hand side of the above equation is smaller than  $\sum_{i=1}^n \frac{1}{n} \frac{\epsilon}{2} = \epsilon/2$  for all  $n > N$ , and this is true for any i.i.d. observations  $(X_1, Y_1), \dots, (X_n, Y_n)$  from  $\mathcal{X} \times \mathcal{Y}$  and any  $f \in \mathcal{F}^{T_n}$ . Therefore, we have for all  $n > N$ ,  $f \in \mathcal{F}^{T_n}$  and any i.i.d. observations  $(X_1, Y_1), \dots, (X_n, Y_n)$ ,

$$|R_\phi(f) - R_\phi(\pi_{\zeta_n} \circ f)| + |\widehat{R}_{\phi,n}(f) - \widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f)| < \epsilon,$$

that is,  $\sup_{f \in \mathcal{F}^{T_n}} |R_\phi(f) - R_\phi(\pi_{\zeta_n} \circ f)| + |\widehat{R}_{\phi,n}(f) - \widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f)| \leq \epsilon$  for all  $n > N$  and  $(X_1, Y_1), \dots, (X_n, Y_n) \in \mathcal{X} \times \mathcal{Y}$ . Since  $\epsilon$  is arbitrarily chosen, we know  $\sup_{f \in \mathcal{F}^{T_n}} |R_\phi(f) - R_\phi(\pi_{\zeta_n} \circ f)| + |\widehat{R}_{\phi,n}(f) - \widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f)| \rightarrow 0$  as  $n \rightarrow \infty$ , and hence  $\sup_{f \in \mathcal{F}^{T_n}} ||\widehat{R}_{\phi,n}(f) - R_\phi(f)| - |\widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)|| \rightarrow 0$ .  $\square$

By Lemma 7,

$$\begin{aligned} 0 &\leq \left| \sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| - \sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)| \right| \\ &\leq \sup_{f \in \mathcal{F}^{T_n}} \left| |\widehat{R}_{\phi,n}(f) - R_\phi(f)| - |\widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)| \right| \rightarrow 0 \text{ a.s. as } n \rightarrow \infty. \end{aligned}$$

Therefore, if  $\sup_{f \in \pi_{\zeta_n} \circ \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| \rightarrow 0$  .a.s., that is,  $\sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(\pi_{\zeta_n} \circ f) - R_\phi(\pi_{\zeta_n} \circ f)| \rightarrow 0$  a.s., then  $\sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| \rightarrow 0$  a.s. as  $n \rightarrow \infty$ . Theorem 4 (a) is proved.

Here we want to emphasize on the boundedness of loss function  $\phi$  as a key prerequisite for Theorem 4 (a). For an unbounded function (e.g. exponential loss), without truncating on the classifier  $f$ , we may not have the above conclusion about the uniform deviation between  $\widehat{R}_{\phi,n}(f)$  and  $R_\phi(f)$  in  $\mathcal{F}^{T_n}$  even when  $T_n \rightarrow \infty$  slow enough.

Before proving Theorem 4 (b), we need the following inequality.



**Lemma 8.** For any family of functions  $\mathcal{F}$  and loss function  $\phi$ , and for any sample  $\mathcal{S}_n$ , let  $f_n^* = \operatorname{argmin}_{f \in \mathcal{F}} \widehat{R}_{\phi,n}(f)$ , we have

$$0 \leq R_\phi(f_n^*) - \inf_{f \in \mathcal{F}} R_\phi(f) \leq 2 \sup_{f \in \mathcal{F}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)|.$$

*Proof of Lemma 8.* By the choice of  $f_n^*$ , we know  $|\widehat{R}_{\phi,n}(f_n^*) - R_\phi(f_n^*)| \leq \sup_{f \in \mathcal{F}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)|$ . Then

$$\begin{aligned} R_\phi(f_n^*) - \inf_{f \in \mathcal{F}} R_\phi(f) &\leq R_\phi(f_n^*) - \widehat{R}_{\phi,n}(f_n^*) + \widehat{R}_{\phi,n}(f_n^*) - \inf_{f \in \mathcal{F}} R_\phi(f) \\ &\leq \sup_{f \in \mathcal{F}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| + \sup_{f \in \mathcal{F}} \left( \widehat{R}_{\phi,n}(f_n^*) - R_\phi(f) \right) \\ &\leq 2 \sup_{f \in \mathcal{F}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)|. \end{aligned}$$

□

Let  $\mathcal{F} = \mathcal{F}^{T_n}$  in the above lemma, then  $f_n^* = \operatorname{argmin}_{f \in \mathcal{F}^{T_n}} \widehat{R}_{\phi,n}(f)$ , and hence we can bound  $R_\phi(f_n^*) - \inf_{f \in \mathcal{F}^{T_n}} R_\phi(f)$  by

$$0 \leq R_\phi(f_n^*) - \inf_{f \in \mathcal{F}^{T_n}} R_\phi(f) \leq 2 \sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)|. \quad (40)$$

Therefore, if we can show that

$$\sup_{f \in \mathcal{F}^{T_n}} |\widehat{R}_{\phi,n}(f) - R_\phi(f)| \rightarrow 0 \text{ a.s.} \quad (41)$$

as  $n \rightarrow \infty$ , then  $R_\phi(f_n^*) \rightarrow R_\phi^*$  a.s. by inequality (40) and Assumption 1. But (41) is just Theorem 4 (a).

## A.12 Proof of Lemma 5

For part (a), we will use Hoeffding's inequality for bounded random variables to obtain

$$\mathbb{P}(\widehat{R}_{\phi,n}(\tilde{f}_n) - R_\phi(\tilde{f}_n) \geq t_n) \leq \exp\left(-\frac{2nt_n^2}{M_\phi^2}\right) = \delta_n,$$

where  $M_\phi = \sup_{x \in \mathbb{R}} \phi(x)$ . Since  $\tilde{f}_n$  only depends on  $n$ , we know  $\{\phi(Y_i \tilde{f}_n(X_i))\}_{i=1}^n$  are independent. We require  $t_n \rightarrow 0$  as  $n \rightarrow \infty$  and  $\sum_{n=1}^{\infty} \delta_n < \infty$ . For example, we take  $t_n = \frac{1}{n^{1/4}}$ . Then by Borel-Contelli Lemma, we have Lemma 5 (a).

For part (b), by the numerical convergence Theorem 1, we know  $\widehat{R}_{\phi,n}(F_T) \rightarrow R_{\phi,n}^*$  a.s. as  $T \rightarrow \infty$ , and the convergence rate only depends on  $T$ . Now let  $\{T_n\}_{n=1}^{\infty}$  be a sequence with  $T_n \rightarrow \infty$ , we then have  $\widehat{R}_{\phi,n}(F_{T_n}) - R_{\phi,n}^* \rightarrow 0$  as  $n \rightarrow \infty$ . But  $\widehat{R}_{\phi,n}(f_n) \geq R_{\phi,n}^*$ , therefore,  $(\widehat{R}_{\phi,n}(F_{T_n}) - \widehat{R}_{\phi,n}(\tilde{f}_n))_+ \rightarrow 0$  a.s. when  $n \rightarrow \infty$ .

## A.13 Proof of Corollary 1

We will now adapt the method of Bartlett and Traskin (2007) (see Theorem 1 therein) to incorporate non-convex losses. For almost every  $w \in (w, \mathcal{S}, \mathbb{P})$ , we can have sequences  $\epsilon_n^1(w) \rightarrow 0$ ,  $\epsilon_n^2(w) \rightarrow 0$  and  $\epsilon_n^3(w) \rightarrow 0$  such that

$$R_\phi(F_{T_n}) \leq R_{\phi,n}(F_{T_n}) + \epsilon_n^1(w) \quad (42)$$

$$\leq R_{\phi,n}(\tilde{f}_n) + \epsilon_n^1(w) + \epsilon_n^2(w) \quad (43)$$

$$\leq R_\phi(\tilde{f}_n) + \epsilon_n^1(w) + \epsilon_n^2(w) + \epsilon_n^3(w). \quad (44)$$

(42) follows from Theorem 4 (a), (43) follows from Theorem 5 (b), and (44) follows from Theorem 5 (a).

Finally by Theorem 4 (b) and how we chosen the reference sequence  $\{\tilde{f}_n\}$ , we have  $R_\phi(\tilde{f}_n) \rightarrow R_\phi^*$  as  $n \rightarrow \infty$ . It follows from (44) that  $R_\phi(F_{T_n}) \rightarrow R_\phi^*$  a.s. as  $n \rightarrow \infty$ . Since  $\phi$  is an Arch boosting function, we know  $\phi$  is classification-calibrated, and hence by Theorem 3 in Bartlett et al. (2006),

$$L(\text{sign}(F_{T_n})) \rightarrow L^* \text{ a.s.}$$

as  $n \rightarrow \infty$ .

## References

- P. Bartlett, M. Jordan, and J. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138-156, 2006.
- P. Bartlett and M. Traskin. Adaboost is consistent. *Journal of Machine Learning Research*, 8: 2347-2368, 2007.
- A. Ben-Israel and B. Mond. What is invexity? *The Journal of the Australian Mathematical Society. Series B. Applied Mathematics*, 28(01), 1-9, 1986.
- S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. *Advanced Lectures on Machine Learning. Springer Berlin Heidelberg*, 208-240, 2004.
- L. Breiman. Population theory for boosting ensembles. *Annals of Statistics*, 32 (??), 1–11, 2004.
- N. Cesa-Bianchi, S. Shalev-Shwartz, and O. Shamir. Online Learning of Noisy Data. *Information Theory, IEEE Transactions on*, 57(12):7907-7931, 2011.
- A. Christmann and I. Steinwart. On robustness properties of convex risk minimization methods for pattern recognition. *The Journal of Machine Learning Research*, 5:1007-1034, 2004.
- A. Christmann and I. Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 13(3):799-819, 2007.
- M. Collins, RE. Schapire, and Y. Singer. Logistic regression, AdaBoost and Bregman distances. *Machine Learning*, 48(1-3):253-285, 2002.
- B. Craven and B. Glover. Invex functions and duality. *J. Austral. Math. Soc. Ser. A*, 39, 1-20, 1985.
- P. L. Davies and U. Gather. Breakdown and groups *Annals of Statistics*, 38,3, 977-988, 2005.
- A.G. Deshwar and Q. Morris. PLIDA: cross-platform gene expression normalization using perturbed topic models. *Bioinformatics*, 30:(7):956-961, 2014.
- T.G. Dietterich. An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning*, 40(2):, 139–158, 2000.
- C. Domingo and O. Watanabe. Madaboost: a modified version of adaboost. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, 180–189, 2000.
- D. L. Donoho and P. J. Huber. The notion of breakdown point. In *A Festschrift for Erich L. Lehmann* (P. J. Bickel, K. Doksum and J. L. Hodges, Jr., eds.), p 157–184, 1983.
- B. Efron and C. Stein. The jackknife estimate of variance. *The Annals of Statistics*, 586-596, 1981.
- Y. Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121 (2):256-285, 1995.
- Y. Freund. An adaptive version of the boost-by-majority algorithm. *Machine learning*, 43(3): 293-318, 2001.
- Y. Freund and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- Y. Freund and R. Schapire. Large margin classification using the perceptron algorithm. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, pages 209–217., 1998.

- Y. Freund. A more robust boosting algorithm. *arXiv preprint* arXiv:0905.2138, 2009
- J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of statistics*, 28(2):337-407, 2000.
- J. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189-1232, 2001.
- J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*. Vol. 1. Springer, Berlin: Springer series in statistics, 2001.
- C. Gentile and N. Littlestone. The robustness of the  $p$ -norm algorithms. In *Proceedings of the 12th Annual Conference on Computational Learning Theory*, pages 1–11, 1999.
- M. G. Genton, and A. Lucas. Comprehensive definitions of breakdown points for independent and dependent observations. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 65, 81–94, 2003.
- P.D. Grünwald and A. P. Dawid Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics* 32 (4), 1367–1433, 2004
- F. R. Hampel, F. R. Contributions to the theory of robust estimation. *Ph.D. dissertation*, Univ. California, Berkeley, 1968.
- F. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69:383-393, 1974.
- F. Hampel, E. Ronchetti, P. Rousseeuw, and W. Stahel. *Robust Statistics. The approach based on influence functions*. Wiley, New York, 1986.
- C. Hennig. Breakdown points for maximum likelihood estimators of location?scale mixtures. *Annals of Statistics* 32,, no. 4, 1313–1340, 2004.
- P. Huber. *Robust Statistics*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, pp. 308 pp, 1981.
- W. Jiang. Process consistency for AdaBoost. *The Annals of Statistics* 32 (1), 13–29. 2004.
- T. Kanamori, T. Takenouchi, S. Eguchi and N. Murata, Robust loss functions for boosting, *Neural Computation* 19, 2183–2244, 2007.
- M. Kearns and M. Li. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 22(4):807–837, 1993.
- Kobetski, Miroslav, and J. Sullivan. Improved Boosting Performance by Exclusion of Ambiguous Positive Examples. In *ICPRAM*, 11-21, 2013.
- V. Koltchinskii and D. Panchenko. Empirical Margin Distributions and Bounding the Generalization Error of Combined Classifiers. *The Annals of Statistics* 30 (1), 1–50. 2002.
- N. Littlestone. Redundant noisy attributes, attribute errors, and linear-threshold learning using winnow. In *Proceedings of the Fourth Annual Workshop on Computational Learning Theory*, pages 147–156, 1991.
- P. M. Long and R.A. Servedio Random Classification Noise Defeats All Convex Potential Boosters. *Machine Learning*, 78(3): 287–304, 2010.
- G. Lugosi and N. Vayatis. On the Bayes-risk consistency of regularized boosting methods. *The Annual of Statistics*, 32(1): 30–55, 2004.
- L. Mason, J. Baxter, P. Bartlett and M. Frean. Functional gradient techniques for combining hypotheses. *Advances in Neural Information Processing Systems*, pages 221-246, 1999.
- P. McCullagh and J.A. Nelder. *Generalized Linear Models*. Chapman & Hall/CRC, 1989.
- H. Masnadi-Shirazi and N. Vasconcelos. On the design of loss functions for classification: theory, robustness to outliers, and savageboost. *Advances in Neural Information Processing Systems* 21, pp. 1049–1056, 2009.
- Mishra, S. Kant, and G. Giorgi. Invex Functions (The Smooth Case). *Invexity and Optimization*, 11-38, 2008.
- C. H. Müller and S. Uhlig. Estimation of variance components with high breakdown point and high efficiency. *Biometrika*, 88, 353–366, 2001.

- N. Natarajan, J. Dhillon, P. Ravikumar and A. Tewari. Learning with Noisy Labels. In *Advances in Neural Information Processing Systems 26*, pages 1196–1204, 2013.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, pages 2825-2830, 2011.
- A.F. Ruckstuhl, and A.H, Welsh. Robust fitting of the binomial model. *Annals of Statistics*, 29, 1117–1136, 2001.
- P. J. Rousseeuw. Least median of squares regression. *J. Amer. Statist. Assoc.* , 79, 871–880, 1984.
- L. Savage. The elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783-801, 1971.
- R. Schapire, Y. Freund, P. Bartlett, and W. Lee. Boosting the margin: a new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26(5):1651–1686, 1998.
- R. Schapire. Theoretical views of boosting and applications. In *Proceedings of the Tenth International Conference on Algorithmic Learning Theory*, pages 12–24, 1999.
- C. Scott, G. Blanchard, G. Handy, S. Pozzi, and M. Flaska. Classification with Asymmetric Label Noise: Consistency and Maximal Denoising *ArXiv e-prints 1303.1208*, 2013.
- R. Servedio. Smooth Boosting and Learning with Malicious Noise. *Journal of Machine Learning Research*, 4, pages 633-648, 2003.
- L.A. Stefanski, Y. Wu and K. White. Variable Selection in Nonparametric Classification via Measurement Error Model Selection Likelihoods *Journal of American Statistical Association*, 109: (506):574–589, 2014.
- W. Street, W. Wolberg, and O. Mangasarian. Nuclear feature extraction for breast tumor diagnosis. In *Proc. Int. Symp. Electronic Imaging: Science and Technology, San Jose*, 861-870.
- A. J. Stromberg and D. Ruppert. Breakdown in nonlinear regression. *J. Amer. Statist. Assoc.* 87, 991–997, 1992.
- D.E. Tyler. Finite sample breakdown points of projection based multivariate location and scatter statistics. *Annals of Statistics*, 22, 1024–1044, 1994.
- X. Zhang, Y. Wu, L. Wang, and R. Li. Variable selection for support vector machines in moderately high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2014.
- T. Zhang and B. Yu, Boosting with early stopping: Convergence and consistency. *Annals of Statistics* 33 (4) 1538–1579, (2005).