

## The success of AdaBoost and its application in portfolio management

Yijian Chuan\*, Chaoyi Zhao\*, Zhenrui He\* and Lan Wu<sup>†,‡</sup>

*\*School of Mathematical Sciences, Peking University  
Beijing, P. R. China*

*†LMEQF, School of Mathematical Sciences  
Peking University, Beijing, P. R. China*

Received: 26 December 2020; Accepted: 4 April 2021

Published: 14 May 2021

---

### Abstract

We develop a novel approach to explain why AdaBoost is a successful classifier. By introducing a measure of the influence of the noise points (ION) in the training data for the binary classification problem, we prove that there is a strong connection between the ION and the test error. We further identify that the ION of AdaBoost decreases as the iteration number or the complexity of the base learners increases. We confirm that it is impossible to obtain a consistent classifier without deep trees as the base learners of AdaBoost in some complicated situations. We apply AdaBoost in portfolio management via empirical studies in the Chinese market, which corroborates our theoretical propositions.

*Keywords:* AdaBoost; interpolation; noise point; base learner; equal-weighted portfolio.

---

### 1. Introduction

Equal-weighted portfolios are one of the most important strategies in portfolio management. They are portfolios with weights equally distributed across the selected securities in the long and/or short positions. In academic research, numerous studies have suggested that equal-weighted portfolios have a better out-of-sample performance than other portfolios (e.g., [Jobson and Korkie, 1981](#); [James, 2003](#); [DeMiguel \*et al.\*, 2007](#)). [Michaud \(1989\)](#) and [DeMiguel \*et al.\* \(2007\)](#) argued that the equal-weighted strategies do not suffer from the estimation error of the

---

<sup>‡</sup>Corresponding author.

*Email address:* [lwu@pku.edu.cn](mailto:lwu@pku.edu.cn)

covariance matrix, which is vulnerable to outliers (Tu and Zhou, 2011). In industry, equal-weighted portfolios are popular across portfolio management in practice, particularly in the hedge funds. The MSCI has issued many equal-weighted indexes, which are “some of the oldest and best-known factor strategies that have aimed to identify specific characteristics of stocks generating excess return.”<sup>1</sup>

The core of constructing equal-weighted portfolios is to forecast the long and/or short positions, which is a classification problem. Machine learning usually outstands when dealing with classification problems, and the research of machine learning on classification is diverse. Devroye et al. (1996) lists numerous traditional pattern recognition research in computer science, where the pattern recognition is an alias for the classification problem. Hastie et al. (2009) analyzes and summarizes many machine learning methods, among which include lots of classification methods, such as the Linear/Quadratic Discriminant Analysis, Support Vector Machine, Boosting, Random Forest, etc. The application of machine learning in classification succeeds in various fields, such as email spam identification, handwritten digit recognition, etc. In portfolio management, López de Prado (2018) explained how to apply machine learning to managing funds for some of the most demanding institutional investors. Specifically, Creamer and Freund (2005, 2010) applied the Boosting method in finance, and revealed its practical value. Creamer (2012) used LogitBoost in high-frequency data for Euro futures, and generated positive returns. Wang et al. (2012) invented the nonlinear adaptive style rotation (N-LASR) model by applying AdaBoost in stock’s factor strategy. They wisely incorporated benefits of different factors by the N-LASR model, and the empirical study on the component stocks in Russel 3000 showed a significant risk-adjusted portfolio return. Fiévet and Sornette (2018) proposed a decision tree forecasting model and applied it to S&P 500, which is capable of capturing arbitrary interaction patterns and generating positive returns. Rasekhschaffe and Jones (2019) provided an example of the machine learning techniques to forecast the cross-section of stock returns. Gu et al. (2020) and D’Hondt et al. (2020) gave comprehensive analysis of machine learning methods for the canonical problem of empirical asset pricing, attributing their predicted gains to the nonlinearity.

AdaBoost is a classification method in machine learning, inspiring a tremendous amount of innovations. AdaBoost has developed for more than two decades since Freund and Schapire (1996). Since it is less prone to overfit, Breiman praised AdaBoost—“the best off-the-shelf classifier in the world”—at the 1996 NeurIPS conference (Friedman et al., 2000). AdaBoost has made a significant impact on

---

<sup>1</sup><https://www.msci.com/msci-equal-weighted-indexes>.

machine learning and statistics. To explain AdaBoost's overfitting resistance, Schapire *et al.* (1998) proposed the "margin" theory. Meanwhile, Breiman (1998, 1999) and Friedman *et al.* (2000) discovered the fact that AdaBoost is equivalent to an optimization algorithm, so Friedman (2001) put forward the Gradient Boosting. Inspired by AdaBoost, Breiman (2001) invented the Random Forest (RF) and believed that there are some similarities between RF and AdaBoost. Subsequently, people generalized the Boosting methods, in which two Boosting algorithms are widely-adopted—the XGBoost (Chen and Guestrin, 2016) and the LightGBM (Ke *et al.*, 2017). Till now, the family of Boosting flourishes, and becomes a considerable part in machine learning.

Although there are many studies explaining why AdaBoost is a successful method, people are still curious about its excellent achievement till now. Wyner (2003) and Mease and Wyner (2008) believed that the available interpretation of AdaBoost is "incomplete", particularly on the explanation of its overfitting resistance property. Wyner *et al.* (2017, p. 10) introduced a novel perspective on AdaBoost and RF, and conjectured that their success could be explained by the "spiked-smooth", where *spiked* is "in the sense that it is allowed to adapt in very small regions to noise points", and *smooth* is "in the sense that it has been produced through averaging". In other words, AdaBoost is a self-averaging interpolating method, localizing the effect of the noise points as the iteration number increases.

Our work is motivated by the questions from the industry: "May machine learning strategies outperform other traditional strategies in quantitative investment? Why and how do they work?" López de Prado (2018) gave a comprehensive and systematic approach to apply machine learning methods, and highly appraised Boosting: "We explored a number of standard black-box approaches. Among machine learning methods, we found gradient tree boosting to be more effective than others." Besides, Wang *et al.* (2012) applied AdaBoost to select and combine factors with consistent and interpretative performance, and Zhang *et al.* (2019) proposed a Boosting method to compose portfolios which performs well. These findings answered the first question. There is limited research concerning the mechanism or the interpretability of machine learning in portfolio management. However, interpretability is essential in investment (Feng *et al.*, 2020). In detail, Harvey *et al.* (2016, p. 37) argued: "... a factor developed from first [economic] principles should have a lower threshold t-statistic than a factor that is discovered as a purely empirical exercise." Harvey (2017) proposed an example. They constructed portfolios based on the first, second, and third letters of the ticker symbols, gaining significant excess returns. Nevertheless, most people would not like to adopt this symbol-based portfolio, as they implied. Thus, without

interpretability, portfolio investment is vulnerable. We must pay more attention to the second/third questions.

To answer the “why and how” questions, we should investigate AdaBoost in the framework of statistics to find a theoretical explanation of its outperformance and apply them in portfolio management. [Wyner et al. \(2017\)](#) pointed out that: “The statistics community was especially confounded by two properties of AdaBoost: (1) interpolation (perfect prediction in sample) was achieved after relatively few iterations, (2) generalization error continues to drop even after interpolation is achieved and maintained.” They innovated the concept of “spiked-smooth” classifier created by a self-averaging and interpolating algorithm. They conjectured that the “spiked-smooth” property renders the success of AdaBoost, and provided many delicate examples by simulation to support their viewpoints. Thus, we would like to narrow the gap between the theory and the simulation by strengthening their work from a statistical perspective. To explain the “spiked-smooth” mathematically, we need to distinguish the signal and the noise within the training set in a statistical framework first. Then, one should connect the relationship between the “spiked-smooth” and the test error, explaining the property of overfitting resistance.

In addition, [Wyner et al. \(2017\)](#) pointed out that “boosting should be used like random forests: with large decision trees, without regularization or early stopping”. The point is that larger and deeper decision trees are preferred to be used as the “weak” classifiers (base learners) of AdaBoost, since they can both “interpolate” the training set and realize the goal of “spiked-smooth”. This point contradicts with the common sense about machine learning and statistics, and statisticians usually believe complexity leads to overfitting. Therefore, we wonder if the AdaBoost method can boost shallow trees when the true model is very complicated, just as we cannot “make bricks without straw”. We try to find out that under what populations will the AdaBoost method be unable to achieve good performance if the base learners are very weak. We want to demonstrate the viewpoints of [Wyner et al. \(2017\)](#) in a mathematical framework.

In this paper, we show that how AdaBoost can dig out more nonlinear information in the training set without increasing the test error. Our work is composed of three parts. First, to concrete the abstract concept “spiked-smooth” into a measurable value, we define a measure of the influence of the noise points in the training set for a given method. The measure can also be regarded as a measure of the localization of the given method. We discover the connection between the measure and the out-of-sample performance. That is, under certain conditions, if the influence of the noise points is not essential, then the test error will be low. A toy example clarifies the theorem, intuitively illustrating the influence of the noise

and explaining why it controls the test error. For AdaBoost, we show that, as the number of iterations increases or the depth of the base learners grows, it becomes more robust to the influence of the noise, and thus lead to a lower test error. Therefore, we give a theoretical explanation about why AdaBoost has a good performance without overfitting in noisy training sets.

Second, we confirm that it is a better choice to use deeper/larger decision trees as base learners of AdaBoost in the sense of digging out complex information. Specifically, we propose several counterexamples that AdaBoost based on shallow decision trees fails to handle, even after iterating infinite times. We generalize the results and indicate that AdaBoost based on shallow decision trees would fail in recognizing a certain kind of information, while the one based on deep decision trees could easily solve out. Therefore, these findings suggest that AdaBoost based on deep decision trees maybe better.

Third, the empirical studies in the Chinese market corroborate our theoretical propositions. The theoretical results about the interpolation and the localization of AdaBoost in the previous parts of this paper is verified by constructing an optimal portfolio strategy. Besides, the result also illustrates the good performance of the equal-weighted portfolio generated by the selected optimal classifier trained by AdaBoost.

The outline of this paper is as follows. Section 2 introduces a measure of “spiked-smooth”, illustrates the relationship between the measure and the test error, and explains the success of AdaBoost. Section 3 identifies that AdaBoost based on deep trees can dig out more information, while the one based on shallow trees fails. Section 4 provides empirical studies of AdaBoost in the Chinese stock market. Section 5 concludes.

## **2. The Influence of the Noise Points and AdaBoost**

In this section, we give a strict definition for the “spiked-smooth” suggested by [Wyner \*et al.\* \(2017\)](#) in the framework of the Bayes classifier. Under the framework, we explain the success of AdaBoost by developing a concrete measure.

First, we describe a background model of the binary classification problem and the Bayes classifier, and define the signal/noise points for a given training set. Based on these concepts, we build a bridge between the Bayes classifier and the interpolating classifier. We define a measure of the influence of the noise points, and specify its property. Second, we explore the connection between the measure and the test error. Last, we explain the success of AdaBoost as the minimization of the influence of the noise points in the sense of the “spiked-smooth”, and reveal its potential applications in portfolio management.

## 2.1. The model of the binary classification

A prediction model consists of an input vector  $X \in \mathcal{X}$ , an output  $Y \in \mathcal{Y} = \{\pm 1\}$ , and a prediction classifier  $f : \mathcal{X} \rightarrow \mathcal{Y}$ . For simplicity, let us assume that the distribution of  $X$  is absolutely continuous with respect to the Lebesgue measure. We restate the definition of the Bayes classifier/error (Devroye et al., 1996, p. 2) in Definition 1 below.

**Definition 1 (Bayes Classifier/Error).** Given the population  $X$  and  $Y \in \mathcal{Y} = \{\pm 1\}$ , the Bayes classifier is

$$f^B := \arg \min_{f: \mathcal{X} \rightarrow \mathcal{Y}} \mathbb{P}\{f(X) \neq Y\},$$

and the minimum is the Bayes error  $q$ , i.e.,

$$q := \min_{f: \mathcal{X} \rightarrow \mathcal{Y}} \mathbb{P}\{f(X) \neq Y\}.$$

According to the definition,  $f^B$  is the classifier minimizing the test error, which can be represented by the conditional distribution  $\mathbb{P}_{Y|X}$  in the population. One can show  $f^B(\mathbf{x}) = \text{sign}(\mathbb{P}\{Y = 1|X = \mathbf{x}\} - 0.5)$ , when the population satisfies certain canonical conditions. There is no classifier having lower test error than  $f^B$ . We give a general representation of the test error of a classifier by the Bayes classifier in Lemma 1 while the noise and  $X$  are independent.

**Lemma 1.** Given the population  $(X, Y)$ , the Bayes classifier  $f^B$ , and the Bayes error  $q$ , if  $\mathbf{1}_{\{Y \neq f^B(X)\}}$  and  $X$  are independent, then, for any classifier  $f$ ,

$$\mathbb{P}\{f(X) \neq Y\} = q\mathbb{P}_X\{f(X) = f^B(X)\} + (1 - q)\mathbb{P}_X\{f(X) \neq f^B(X)\}. \quad (1)$$

**Proof.** We have

$$\begin{aligned} \mathbb{P}\{f(X) \neq Y\} &= \mathbb{P}\{f(X) \neq Y, f(X) = f^B(X)\} + \mathbb{P}\{f(X) \neq Y, f(X) \neq f^B(X)\} \\ &= \mathbb{P}\{f^B(X) \neq Y, f(X) = f^B(X)\} + \mathbb{P}\{f^B(X) = Y, f(X) \neq f^B(X)\} \\ &= \mathbb{P}\{f^B(X) \neq Y\}\mathbb{P}_X\{f(X) = f^B(X)\} + \mathbb{P}\{f^B(X) = Y\}\mathbb{P}_X\{f(X) \neq f^B(X)\} \\ &= q\mathbb{P}_X\{f(X) = f^B(X)\} + (1 - q)\mathbb{P}_X\{f(X) \neq f^B(X)\}. \end{aligned}$$

□

A natural corollary of (1) is  $\mathbb{P}\{f(X) \neq Y\} = q + (1 - 2q)\mathbb{P}_X\{f(X) \neq f^B(X)\}$ . In other words,  $\mathbb{P}\{f(X) \neq Y\}$  is a linear function of  $\mathbb{P}_X\{f(X) \neq f^B(X)\}$ .

Next, we introduce the concept of the signal/noise points of the training set  $T$  in the following Definition 2.

**Definition 2 (Signal/Noise Points).** Given a training set  $T = (\mathbf{x}_i, y_i)_{i=1}^n$  generated from the population  $(X, Y)$  and the Bayes classifier  $f^B$ , a point  $(\mathbf{x}_i, y_i)$  is a signal point, if  $f^B(\mathbf{x}_i) = y_i$ ; and it is a noise point, if  $f^B(\mathbf{x}_i) \neq y_i$ .

In short, the Bayes classifier  $f^B$  distinguishes the signal/noise points of a training set. Heuristically, the signal points are the points that are equal to the output of the Bayes classifier, while the noise points are not.

We recall the definition of the interpolation classifier proposed by [Wyner et al. \(2017\)](#) for coherence.

**Definition 3 (Interpolating Classifier).** A classifier  $f$  is an interpolating classifier on the training set  $T = (\mathbf{x}_i, y_i)_{i=1}^n$ , if  $f(\mathbf{x}_i) = y_i$  for all  $(\mathbf{x}_i, y_i), i = 1, \dots, n$ .

Immediately, we can obtain a property of the interpolating classifier, i.e., its training error is 0 on the training set  $T = (\mathbf{x}_i, y_i)_{i=1}^n$ .

Though the Bayes classifier is the best classifier in the sense of minimizing the test error, it does not necessarily interpolate the given training set  $T = (\mathbf{x}_i, y_i)_{i=1}^n$ . The Bayes classifier  $f^B$  violates interpolation at and only at the noise points, as implied in Definition 3. Thus, in view of the training set, the difference between an interpolating classifier and the Bayes classifier is only on the noise points. So, we propose a definition of a purified training set of  $T$  by converting the noise points into the ‘‘signal’’ points.

**Definition 4 (Purified Training Set).** Given a training set  $T = (\mathbf{x}_i, y_i)_{i=1}^n$  from the population  $(X, Y)$  and the Bayes classifier  $f^B$ , the purified training set of  $T$  is defined as  $T_p := (\mathbf{x}_i, f^B(\mathbf{x}_i))_{i=1}^n$ .

There is no noise point in  $T_p$ . In other words, the Bayes classifier must interpolate the purified training set  $T_p$ . We can also rewrite the definition of the purified training set as

$$T_p = (\mathbf{x}_i, \theta_i)_{i=1}^n, \quad \theta_i = \begin{cases} y_i, & i \text{ is a signal point;} \\ -y_i, & i \text{ is a noise point.} \end{cases}$$

The two training sets  $T$  and  $T_p$  share the same input  $\mathbf{x}_i$ 's but different outputs, and the difference between the outputs of the two sets is only on the noise points. The purpose is to separate out the influence of the noise points from the whole information contained in the training set  $T$ .

Last, based on the previous preparations, we propose a measure of the influence of the noise points (ION) for a given training set  $T$  and a given method  $\mathcal{M}$ . It helps us to compare the properties of different methods, such as one nearest neighbor (1NN) or AdaBoost, on a given training set.

**Definition 5 (ION).** Given the marginal probability measure of  $X$  ( $\mathbb{P}_X$ ), we define the influence of noise (ION), a function of the training set  $T$  and the method  $\mathcal{M}$ :

$$\text{ION}(\mathcal{M}, T) := \mathbb{P}_X\{f_T(X) \neq f_{T_p}(X)\}, \quad (2)$$

where  $f_T$  is the classifier trained on the training set  $T$  using method  $\mathcal{M}$ , and  $f_{T_p}$  is on the purified training set  $T_p$  using method  $\mathcal{M}$ .

We interpret Definition 5.  $\mathcal{M}$  represents a specific method. For instance, for 1NN, one can apply it on the training set  $T$  and  $T_p$ , which generates two classifiers  $f_T$  and  $f_{T_p}$ . Then, by comparing the two classifiers, we can get the value of ION (1NN,  $T$ ). The ION is defined according to two sets: the training set  $T$  and the proxy of the training set generated by  $f^B$ .

Although Definition 5 does not require the interpolation of the classifier  $f$ , ION usually characterizes the performance of the method which generates interpolating classifiers on a given training set. Meanwhile,  $0 \leq \text{ION} \leq 1$ . If ION is low, then the classifier is robust to the noise points on the training set of the given method, and vice versa.

Interpolation is not necessarily bad, if it subjects to some “mechanism” (Wyner et al., 2017). Although some interpolating classifiers “can be shown to be inconsistent and have poor generalization error in environments with noise”, “the claim that all interpolating classifiers overfit is problematic”. The classifiers generated by 1NN or random forest are both interpolating classifiers, but their ION may not be the same. Furthermore, Wyner et al. (2017) suggested: “an interpolated classifier, if sufficiently local, minimizes the influence of noise points in other parts of the data.” The next question is, what is the relationship between the ION and the so-called “spiked-smooth” classifier.

## 2.2. The ION and the test error

In this section, we reveal the connection between the ION and the test error from theoretical and numerical perspectives.

First, we prove that, under certain conditions, the lower the ION, the lower the test error.

**Proposition 1.** Given the population  $(X, Y)$  such that  $\mathbf{1}\{Y \neq f^B(X)\}$  is independent of  $X$ , let  $f_T^{(1)}$  and  $f_T^{(2)}$  denote the classifiers generated from two different methods  $\mathcal{M}^{(1)}$  and  $\mathcal{M}^{(2)}$  on the training set  $T$ , and  $f_{T_p}^{(1)}$  and  $f_{T_p}^{(2)}$  denote the ones generated from  $\mathcal{M}^{(1)}$  and  $\mathcal{M}^{(2)}$  on the purified training set  $T_p$ , respectively. If

$$f_{T_p}^{(1)}(x) = f_{T_p}^{(2)}(x) = f^B(x), \quad \text{a.s.}, \quad (3)$$



and

$$\text{ION}(\mathcal{M}^{(1)}, T) < \text{ION}(\mathcal{M}^{(2)}, T),$$

then

$$\mathbb{P}\{f_T^{(1)}(X) \neq Y\} < \mathbb{P}\{f_T^{(2)}(X) \neq Y\}. \quad (4)$$

**Proof.** Because of (3), we have

$$\begin{aligned} \mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X)\} &= \mathbb{P}_X\{f_T^{(1)}(X) \neq f_{T_p}^{(1)}(X)\} = \text{ION}(\mathcal{M}^{(1)}, T) < \text{ION}(\mathcal{M}^{(2)}, T) \\ &= \mathbb{P}_X\{f_T^{(2)}(X) \neq f_{T_p}^{(2)}(X)\} = \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X)\}. \end{aligned}$$

Thus, by Lemma 1, (4) holds.  $\square$

Proposition 1 shows that the ION controls the test error. Specifically, it means that, if the two methods could reach the Bayes classifier in the purified training set, then the method with lower ION outperforms the others in the sense of the test error. For instance,  $\mathcal{M}^{(1)}$  might indicate 1NN, while  $\mathcal{M}^{(2)}$  indicate AdaBoost.

However, the condition (3) is slightly unnatural. It is so strong that it could only hold in several particular training sets. We therefore weaken the original condition (3) and establish a more natural condition in Theorem 1 below.

**Theorem 1.** *Given the population  $(X, Y)$  such that  $\mathbf{1}_{\{Y \neq f^B(X)\}}$  is independent of  $X$ , let  $f_T^{(1)}$  and  $f_T^{(2)}$  denote the classifiers generated from two different methods  $\mathcal{M}^{(1)}$  and  $\mathcal{M}^{(2)}$  on the training set  $T$ , and  $f_{T_p}^{(1)}$  and  $f_{T_p}^{(2)}$  denote the ones generated from  $\mathcal{M}^{(1)}$  and  $\mathcal{M}^{(2)}$  on the purified training set  $T_p$ , respectively. The size of the training set  $T$  or  $T_p$  is  $n$ . If*

$$\lim_{n \rightarrow \infty} \mathbb{P}\{f_{T_p}^{(j)}(X) \neq Y\} = q, \quad j = 1, 2, \quad (5)$$

and

$$\limsup_{n \rightarrow \infty} [\text{ION}(\mathcal{M}^{(1)}, T) - \text{ION}(\mathcal{M}^{(2)}, T)] < 0, \quad (6)$$

then

$$\limsup_{n \rightarrow \infty} [\mathbb{P}\{f_T^{(1)}(X) \neq Y\} - \mathbb{P}\{f_T^{(2)}(X) \neq Y\}] < 0. \quad (7)$$

**Proof.** To begin with, by Lemma 1, (5) is equivalent to

$$\lim_{n \rightarrow \infty} \mathbb{P}_X\{f_{T_p}^{(j)}(X) \neq f^B(X)\} = 0, \quad j = 1, 2, \quad (8)$$

Then,

$$\begin{aligned}
 & \mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X)\} - \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X)\} \\
 & \quad - (\mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X), f_T^{(2)}(X) \neq f_{T_p}^{(2)}(X)\}) \\
 & \quad + \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X), f_T^{(2)}(X) = f_{T_p}^{(2)}(X)\}) \\
 & = \mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X), f_T^{(1)}(X) \neq f_{T_p}^{(1)}(X)\} \\
 & \quad - \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X), f_T^{(2)}(X) \neq f_{T_p}^{(2)}(X)\} \\
 & \quad + \mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X), f_T^{(1)}(X) = f_{T_p}^{(1)}(X)\} \\
 & \quad - \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X), f_T^{(2)}(X) = f_{T_p}^{(2)}(X)\} \\
 & = \mathbb{P}_X\{f_{T_p}^{(1)}(X) = f^B(X), f_T^{(1)}(X) \neq f_{T_p}^{(1)}(X)\} \\
 & \quad - \mathbb{P}_X\{f_{T_p}^{(2)}(X) = f^B(X), f_T^{(2)}(X) \neq f_{T_p}^{(2)}(X)\} \\
 & \quad + \mathbb{P}_X\{f_{T_p}^{(1)}(X) \neq f^B(X), f_T^{(1)}(X) = f_{T_p}^{(1)}(X)\} \\
 & \quad - \mathbb{P}_X\{f_{T_p}^{(2)}(X) \neq f^B(X), f_T^{(2)}(X) = f_{T_p}^{(2)}(X)\} \\
 & =: A_n^{(1)} - A_n^{(2)} + B_n^{(1)} - B_n^{(2)}.
 \end{aligned}$$

By (8), we have

$$\lim_{n \rightarrow \infty} \mathbb{P}_X\{f_{T_p}^{(j)}(X) \neq f^B(X), f_T^{(j)}(X) = f_{T_p}^{(j)}(X)\} = 0, \quad j = 1, 2,$$

and thus  $\lim_{n \rightarrow \infty} B_n^{(j)} = 0, j = 1, 2$ . By (8), we also have

$$\lim_{n \rightarrow \infty} (A_n^{(j)} - \mathbb{P}_X\{f_T^{(j)}(X) \neq f_{T_p}^{(j)}(X)\}) = 0, \quad j = 1, 2,$$

so, by (2), one can show that  $A_n^{(j)}$  and  $\text{ION}(\mathcal{M}^{(j)}, T)$  share the same limit,  $j = 1, 2$ . Therefore,

$$\begin{aligned}
 & \lim_{n \rightarrow \infty} [(\mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X)\} - \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X)\}) \\
 & \quad - (\text{ION}(\mathcal{M}^{(1)}, T) - \text{ION}(\mathcal{M}^{(2)}, T))] = 0.
 \end{aligned}$$

Because of (6), we have

$$\limsup_{n \rightarrow \infty} [\mathbb{P}_X\{f_T^{(1)}(X) \neq f^B(X)\} - \mathbb{P}_X\{f_T^{(2)}(X) \neq f^B(X)\}] < 0.$$

Further, by Lemma 1, (7) holds. □

We interpret Theorem 1. First, (5) is a very weak condition. It assumes the two methods are consistent on the purified training set  $T_p$ . In fact, many classical

methods have been proved to be consistent. Furthermore, because there is no noise point in  $T_p$ , the consistency on  $T_p$  is easier to achieve than on  $T$ . Even the notoriously easy-to-overfit method, 1NN, is consistent in such a good training set  $T_p$  but not necessarily consistent in  $T$ , according to the Cover–Hart inequality (Cover and Hart, 1967).

Second, (6) is about the property of some methods regarding a certain training set. Instead of subjectively describing the property of the methods, it measures the influence of the noise points in the particular training set objectively.

Third, according to Theorem 1, the decrease of ION implies that the method is minimizing the influence of the noise points and thus enhancing the generalization ability. It means that, for most methods, the ION is a good indicator of the test error.

Fourth, there is no natural conflict between interpolation and lower ION. For a classifier, the purpose of interpolating is to take all information contained in the signal points as much as possible, while the goal of lowering ION is to reduce the impact attributed to the noise points.

In order to have a concrete understanding of Proposition 1 and Theorem 1, we give a 2-dim toy example. First, the population is

$$\mathbb{P}\{Y = 1|X = \mathbf{x}\} = \begin{cases} 0.1, & \text{if } x_1 < 0, \\ 0.9, & \text{if } x_1 \geq 0, \end{cases}$$

where  $X$  is uniformly distributed in  $(-1, 1]^2$ . In other words, only the first dimension of  $X$  is relevant to  $Y$ , while the second dimension contributes no information. One can easily solve the Bayes classifier  $f^B(\mathbf{x}) = \text{sign}\{x_1\}$  and the Bayes error  $q = 0.1$ .

Second, we randomly generate a training set  $T$  with a size  $n = 500$ , as in Fig. 1. The training set  $T$  is composed of 460 signal points and 40 noise points. Roughly speaking, the yellow triangles to the left and the blue circles to the right are all noise points. Particularly, on the left and bottom side of the graph in Fig. 1, there is a solid triangle  $\mathbf{x}_{i_0} = (-0.79, -0.41)$ , which is the noise point that would be discussed later.

Third, we apply two methods,  $\mathcal{M}^{(1)}$  (1NN) and  $\mathcal{M}^{(2)}$  (AdaBoost<sup>2</sup>), to generate interpolating classifiers on the training set  $T$ , respectively. Figure 2(a) is the classifier  $f_T^{\text{1NN}}$  generated by method 1NN, while 2(b)  $f_T^{\text{AdaBoost}}$ . The purple vertical dotted line ( $x_1 = 0$ ) is the watershed of  $f^B$ , while the black solid lines are the decision boundaries of  $f_T^{\mathcal{M}^{(i)}}$ . Both classifiers are interpolating, which means that  $f_T^{\mathcal{M}^{(i)}}(\mathbf{x}_i) = y_i, \forall i$ , even though they are from different methods.

---

<sup>2</sup>To be more specific, for the AdaBoost we used, the base learners are decision trees with a maximum depth 4, and the number of iterations is 50.




Fig. 1. The training set  $T$ .

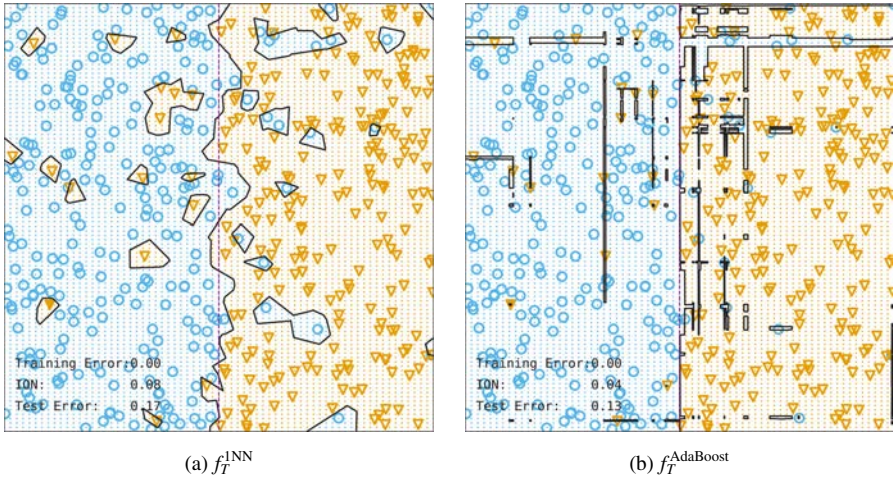


Fig. 2. The training set  $T$  and the classifiers: AdaBoost has lower ION than 1NN.

Fourth, we argue that the ION of AdaBoost is lower than that of 1NN on the training set  $T$ . The classifiers in Fig. 2 are different: The decision boundary of 1NN in Fig. 2(a) is smooth and natural, while that of AdaBoost in Fig. 2(b) is sharp and uneven. However, we argue that the sharp and uneven is better than the smooth and natural in the sense of minimizing and localizing the influence of the noise points. For the isolated noise points, the regions surround them in 2(b) are smaller and

Table 1. The ION and the test error.

Method	ION	Training error	Test error
$f^B$	—	—	0.10
$\mathcal{M}^{(1)} = \text{1NN}$	0.08	0	0.17
$\mathcal{M}^{(2)} = \text{AdaBoost}$	0.04	0	0.13

narrower than that in 2(a). In detail, we focus on the particular noise point lies at about  $\mathbf{x}_{i_0} = (-0.79, -0.41)$ , which is the solid triangle. The area around  $\mathbf{x}_{i_0}$  of  $f_T^{\text{AdaBoost}}$  in 2(b) is very small, while that of  $f_T^{\text{1NN}}$  in 2(a) is a big irregular polygon—the influence of the noise point  $\mathbf{x}_{i_0}$  seems to be lower for AdaBoost than 1NN.<sup>3</sup>

Last, we calculate the ION and the test error, summarized in Table 1, where  $\text{ION}(\text{1NN}, T) \approx 0.08$  and  $\text{ION}(\text{AdaBoost}, T) \approx 0.04$ ,  $\mathbb{P}\{f_T^{\text{1NN}}(X) \neq Y\} \approx 0.17$  and  $\mathbb{P}\{f_T^{\text{AdaBoost}}(X) \neq Y\} \approx 0.13$ . We can observe that the results are in line with our theorem, i.e., the lower the ION, the lower the test error.

Overall, this section connects the ION and the test error. Both the theoretical derivation and the toy example of simulation demonstrate the importance of ION. Particularly, the toy example explains why 1NN is easy to overfit, while AdaBoost not. However, AdaBoost is only a general term for a class of methods, since both the base learners and the number of iterations need to be specified. By choosing different kinds of base learners and different numbers of iterations, we can generate a tremendous amount of specific methods. In Sec. 2.3, we take a close look at the performance of AdaBoost with different hyperparameters from our new perspective: ION.

### 2.3. The ION and AdaBoost

AdaBoost mainly has two hyperparameters. One of them is the complexity of the base learners. The decision trees are one of the most popular base learners of AdaBoost, which is the classical base learner in the monograph [Hastie et al. \(2009\)](#). In this paper, we use the maximum depth of decision trees to indicate the complexity of the base learners. The deeper the decision trees, the more complex the base learners, and the more complex the AdaBoost. The other is the number of iterations, which is the number of the base learners added in total. The higher the number of iterations, the more complex the AdaBoost.

---

<sup>3</sup>It is noteworthy that the influence of the noise points are acting jointly rather than individually, but it does not matter in this heuristic case.

This section corroborates the conclusion of [Wyner et al. \(2017\)](#) with our newly defined concept ION. We show their conclusion that AdaBoost based on large decision trees without early stopping is better. We want to show that AdaBoost generates interpolating classifiers, and both the ION and the test error decrease as the depth of the base learners and the number of iterations increase. Instead of comparing AdaBoost with 1NN, we digest AdaBoost itself with different parameters in details via high-dimensional population of simulation.

The simulation population is

$$\mathbb{P}\{Y = 1|X = \mathbf{x}\} = \begin{cases} 0.1, & \text{if } x_1 \cdot x_2 \cdot x_3 < 0, \\ 0.9, & \text{if } x_1 \cdot x_2 \cdot x_3 \geq 0, \end{cases}$$

where  $X$  is uniformly distributed in  $(-1, 1]^6$ . We randomly generate a training set  $T$  with  $n = 500$ , and compare the results of AdaBoost with different hyperparameters.

In order to explain the reason why AdaBoost without early stopping might be better, we compare the results of AdaBoost with different numbers of iterations  $m = 1, 2, \dots, 250$  but the same maximum depth of the base decision trees. The maximum depth is set as 5. We denote the corresponding classifiers by  $f_T^{(m)}$ ,  $m = 1, 2, \dots, 250$ .

The results are in Fig. 3. The  $x$ -axis in the figure represents the number of iterations  $m$ . We clarify the three lines in detail. The green dashed line is the training error of  $f_T^{(m)}$  on its training set  $T$ , the red dashed-dotted line is the test error of  $f_T^{(m)}$ , and the blue solid line is the ION of AdaBoost in the training set:

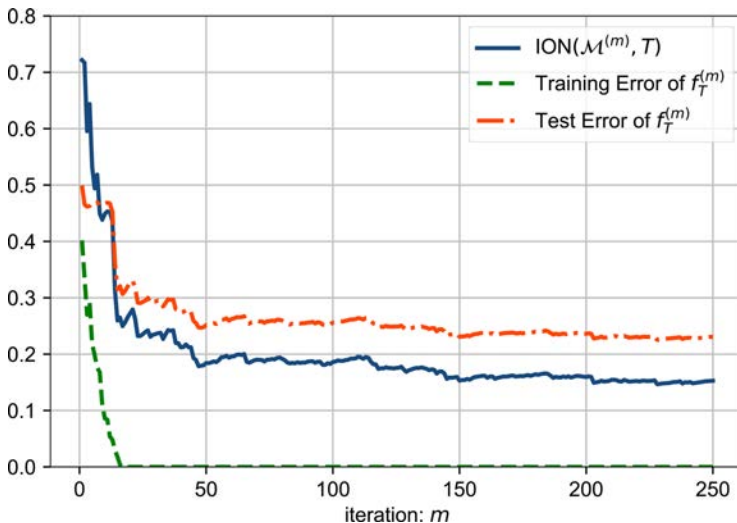


Fig. 3. The performance of AdaBoost regarding to  $m = 1, 2, \dots, 250$ .

$\text{ION}(\mathcal{M}^{(m)}, T)$ . All the three lines decrease sharply when  $n < 20$ . When  $n \geq 20$ , the training error remains 0, but the test error and ION keep decreasing.

From Fig. 3, we have the following observations. First, AdaBoost is minimizing the influence of the noise points. When  $m \geq 20$ , the test error decreases but the training error remains 0. A natural question is, what is AdaBoost doing? There are many explanations. [Wyner et al. \(2017\)](#) believed that AdaBoost is self-averaging and generating a “spiked-smooth” classifier by minimizing the influence of the noise points. We corroborate their work with the blue solid line ION. When  $m \geq 20$ , although the training error remains 0, the ION continues to decrease, which reflects the decrease of the influence of the noise points. Thus, as the number of iterations increases, AdaBoost keeps interpolating, and simultaneously minimizes the influence of the noise points. Second, the iteration of AdaBoost can be divided into two stages: The first stage is the sharp decrease of the training error ( $m < 20$ ), and the second stage is the decrease of ION ( $m \geq 20$ ). The first stage can be considered as the formation of the rough skeleton of the classifier, while the second stage can be treated as the process of the details with the minimization of the influence of the noise points.<sup>4</sup>

For another, we show that AdaBoost based on deep/large decision trees is better, and explain it by ION. Specifically, we apply AdaBoost based on different decision trees but the same number of iterations  $m = 250$ , where the base decision trees have different maximum depths from 1 to 8. In other words, the number of the terminal leaves of the base decision trees varies from 2 to 256. We denote the corresponding classifiers by  $f_T^{(j)}, j = 1, 2, \dots, 8$ .

The results are presented in Fig. 4. The  $x$ -axis in the figure represents the maximum depth, i.e.,  $j$  for  $f_T^{(j)}$ . The three lines are the same as those in Fig. 3, and so are the interpretations.

Overall, AdaBoost based on large decision trees without early stopping is better, which can be explained as the decrease of the ION. Given the condition that the training error is 0, the influence of the noise points decreases as the depth of the base decision trees and the number of iterations increase.

Now, return to the main line of the paper, we show that AdaBoost would not overfit even interpolating, when digging out complex structures of factors in constructing equal-weighted portfolios. As it was emphasized by [López de Prado \(2018, p. 15\)](#), the linear methods are awfully simplistic and useless, and would “fail to recognize the complexity of the data”. The academia and industry shift their focus to the nonlinear ones. There are a tremendous amount of machine learning methods applied in various data and fields in finance. Many of the

---

<sup>4</sup>However, the two stages cannot be divided arbitrarily, because ION may also play a role in the first stage.

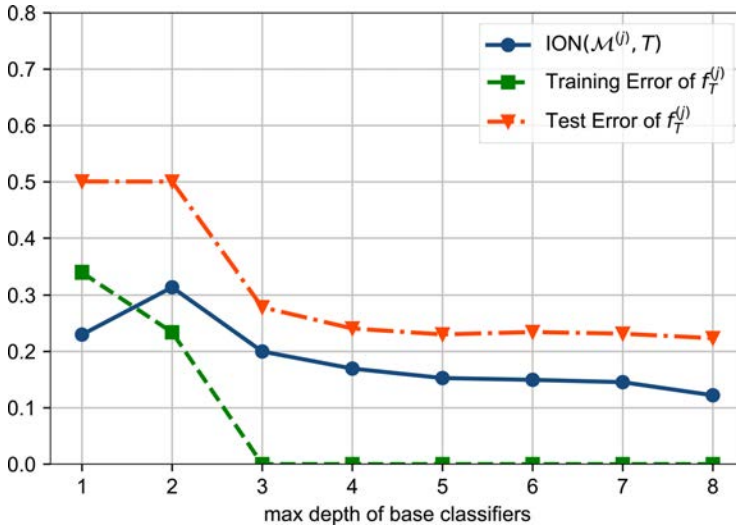


Fig. 4. The performance of AdaBoost regarding to  $j = 1, 2, \dots, 8$ .

machine learning methods suffer from the overfitting and the low interpretation. However, AdaBoost is not heavily affected by them as illustrated above.

In Sec. 4, we give empirical studies about specific factors or strategies, and prove the advantage of AdaBoost in portfolio management. But we want to clarify what kind of nonlinear information can AdaBoost dig out first.

### 3. Base Learners of AdaBoost

AdaBoost is a boosting method. It boosts the performance of a series of base learners, or “weak classifiers”. People usually choose shallow trees (such as “stumps”, i.e., decision trees with only one layer) as base learners since they are “weak” enough and thus can avoid overfitting.

However, in many fields, especially in the area of portfolio management, using stumps as base learners may not capture the nature of the population, since the population is usually rather complicated. *Wyner et al. (2017)* proposed that the deep and large trees will allow the base learners to interpolate the data without overfitting, and it is a better choice to use deep trees as base learners. We have already shown the result mathematically in Sec. 2.3 from the perspective of the ION.

In this section, we discuss the shortcomings of AdaBoost based on stumps. We first show that stumps cannot deal with the “XOR” classification problem. Then, we generalize the result and demonstrate that AdaBoost based on stumps cannot



deal with populations without “comonotonicity”. These kinds of populations are common in finance, since the investment activities in the financial world are usually rather complicated and interactive.

### 3.1. The “XOR” population

In this section, we use a toy example to show that the shallow trees (especially stumps) are not always capable to capture the patterns of the population. We introduce the Boolean operator “exclusive OR” (XOR) first.

**Definition 6 (2-XOR).** The 2-XOR function is defined as  $\text{XOR}_2 : \mathbb{R}^2 \rightarrow \{\pm 1\}$  such that

$$\text{XOR}_2(x_1, x_2) = \begin{cases} -1, & \text{if } x_1 \cdot x_2 \geq 0, \\ 1, & \text{if } x_1 \cdot x_2 < 0. \end{cases}$$

**Definition 7 ( $k$ -XOR).** For  $k > 2$ , the  $k$ -XOR function, denoted as  $\text{XOR}_k$ , is defined recursively as

$$\text{XOR}_k(x_1, x_2, \dots, x_k) = \text{XOR}_2(\text{XOR}_{k-1}(x_1, x_2, \dots, x_{k-1}), x_k).$$

The Boolean operator  $k$ -XOR is an important function in computer science (for instance, the parity check), and it is also a classical example in [Hastie et al. \(2009\)](#). It can also bring insights into portfolio management, since the  $k$ -XOR can characterize the interaction among different factors. There are many studies focusing on the interaction among various factors ([Asness et al., 2018](#)). Figure 5 shows intuitive illustrations of the 2-XOR and the 3-XOR functions. The outputs are not the same in adjacent quadrants (or octants), which is a common pattern of interaction.




Fig. 5. Intuitions of 2-XOR and 3-XOR functions.

Now we show that the stumps cannot deal with the classification problems with the Bayes classifier  $f^B = \text{XOR}_k$ , even in the case that the Bayes error is 0. For instance, if we use a stump classifier  $f$  to classify the 2-XOR function, we can easily show that  $\mathbb{P}_{x_1, x_2} \{f(x_1, x_2) \neq \text{XOR}_2(x_1, x_2)\}$  is always 50% no matter how the stump is trained. That is because, a stump is equivalent to a partition of  $\mathbb{R}^2$  along the direction of one axis. After the partition, both half-spaces still contain values of 1 (accounts for 50%) and  $-1$  (50%), which leads to a test error 50%.

The conclusion can also be generalized to high-dimensional spaces. Let  $f^{(\leq k)}$  denote a decision tree whose depth is not more than  $k$ , and  $f^{(k)}$  denote a decision tree whose depth is  $k$ . We have the following result in Theorem 2.

**Theorem 2.** *Applying a decision tree  $f^{(\leq k)}$  on the  $(k + 1) - \text{XOR}$  classification problem will always lead to  $\mathbb{P}_X \{f^{(\leq k)}(X) = \text{XOR}_{k+1}(X)\} = 50\%$ , where  $X = (x_1, \dots, x_{k+1})$ .*

**Proof.** We prove the theorem by induction. The case of  $k = 1$  has already been proved. Now we assume that our conclusion holds for  $f^{(\leq k-1)}$ . We want to prove that it also holds for  $f^{(k)}$ .

Without loss of generality, we assume that the splitting variable of  $f^{(k)}$ 's first layer is the first feature  $x_1$ , then

$$f^{(k)} = \mathbf{1}_{\{x_1 \leq c\}} f_1^{(k-1)} + \mathbf{1}_{\{x_1 > c\}} f_2^{(k-1)},$$

where  $f_1^{(k-1)}$  and  $f_2^{(k-1)}$  represent the left subtree and the right subtree of  $f^{(k)}$ 's top node respectively, and  $c$  is the splitting value. Let  $X_{[-1]} = (x_2, \dots, x_{k+1})$ , then

$$\begin{aligned} \mathbb{P}_X \{f^{(k)}(X) = \text{XOR}_{k+1}(X)\} &= \mathbb{P}_X \{f_1^{(k-1)}(X) = \text{XOR}_{k+1}(X) | x_1 \leq c\} \mathbb{P}_X \{x_1 \leq c\} \\ &\quad + \mathbb{P}_X \{f_2^{(k-1)}(X) = \text{XOR}_{k+1}(X) | x_1 > c\} \mathbb{P}_X \{x_1 > c\}. \end{aligned}$$

Assuming  $c > 0$  without loss of generality, then we have

$$\begin{aligned} &\mathbb{P}_X \{f_1^{(k-1)}(X) = \text{XOR}_{k+1}(X) | x_1 \leq c\} \\ &= \mathbb{P}_X \{f_1^{(k-1)}(X) = \text{XOR}_{k+1}(X) | x_1 \leq c, x_1 \leq 0\} \mathbb{P}_X \{x_1 \leq 0\} \\ &\quad + \mathbb{P}_X \{f_1^{(k-1)}(X) = \text{XOR}_{k+1}(X) | 0 < x_1 \leq c\} \mathbb{P}_X \{x_1 > 0\} \\ &= \mathbb{P}_X \{f_1^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\} \mathbb{P}_X \{x_1 \leq 0\} \\ &\quad + \mathbb{P}_X \{-f_1^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\} \mathbb{P}_X \{x_1 > 0\}, \end{aligned}$$

and

$$\mathbb{P}_X \{f_2^{(k-1)}(X) = \text{XOR}_{k+1}(X) | x_1 > c\} = \mathbb{P}_X \{-f_2^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\}.$$

Our inductive assumption tells us that, for the  $\text{XOR}_k$  classification problem, both  $f_1^{(k-1)}$  and  $f_2^{(k-1)}$  will have a 50% error. Hence, the three probabilities  $\mathbb{P}_X\{f_1^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\}$ ,  $\mathbb{P}_X\{-f_1^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\}$  and  $\mathbb{P}_X\{-f_2^{(k-1)}(X_{[-1]}) = \text{XOR}_k(X_{[-1]})\}$  are all equal to 50%. Finally,

$$\begin{aligned} & \mathbb{P}_X\{f^{(k)}(X) = \text{XOR}_{k+1}(X)\} \\ &= [50\%\mathbb{P}_X\{x_1 \leq 0\} + 50\%\mathbb{P}_X\{x_1 > 0\}]\mathbb{P}_X\{x_1 \leq c\} + 50\%\mathbb{P}_X\{x_1 > c\} \\ &= 50\%\mathbb{P}_X\{x_1 \leq c\} + 50\%\mathbb{P}_X\{x_1 > c\} \\ &= 50\%. \end{aligned}$$

□

In the proof above, we suppose that each component of  $X$  would be split just only one time. In other words, once the CART algorithm (Hastie *et al.*, 2009, p. 305) split a decision tree at  $x_1$ , it will not split at  $x_1$  again in other layers. It is just for clarity and conciseness, because one can use the total probability formula to deal with more complicated situations.

Although the  $k$ -XOR is a special case that each factor interacts with other factors, it is enough to demonstrate that shallow decision trees (especially for one-layer stumps) may be unable to deal with factors that are not independent of each others.

### 3.2. The population without “comonotonicity”

In this section, we show the shortcomings of AdaBoost based on stumps by introducing the concept of “comonotonicity”. The conclusion in this section can be regarded as an extension of Sec. 3.1, since the XOR function do not have the property of comonotonicity, as we will discuss later.

**Definition 8 (Comonotonic Population).** A population  $X \in \mathcal{X} = \mathbb{R}^k$ ,  $Y \in \mathcal{Y} = \{\pm 1\}$  is comonotonic, if its Bayes classifier  $f^B$  satisfies: for any constant  $c$  and any  $i = 1, \dots, k$ , there exists an  $\varepsilon > 0$  such that for each  $a \in (c - \varepsilon, c)$  and  $b \in (c, c + \varepsilon)$ , the elements in

$$\{f^B(x_1, \dots, x_{i-1}, a, x_{i+1}, \dots, x_k) - f^B(x_1, \dots, x_{i-1}, b, x_{i+1}, \dots, x_k) : X_{[-i]} \in \mathbb{R}^{k-1}\}$$

are all non-positive or all non-negative, where  $X_{[-i]} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$ .

To give an intuition of comonotonicity, in Fig. 6, we give three examples of populations which are not comonotonic. For Figs. 6(a)–6(c), the decision boundaries of their Bayes classifiers form shapes of an XOR, a ring, and a diagonal band, respectively. The yellow (light) region takes a value of +1, and the blue (dark) region –1. Note that, in each figure, there both exist arrows from

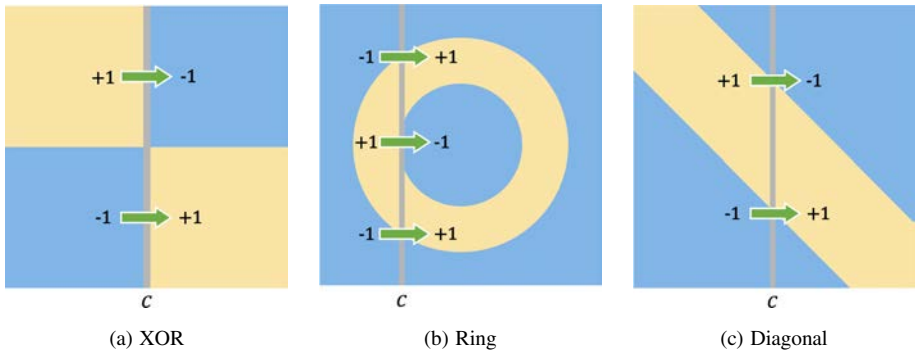


Fig. 6. The populations WITHOUT comonotonicity.

values of  $-1$  to  $+1$ , and arrows from values of  $+1$  to  $-1$ . These arrows tell us that the populations are not comonotonic.

The following theorem shows that AdaBoost based on stumps cannot deal with populations without comonotonicity.

**Theorem 3.** For a population in  $\mathbb{R}^k$  with a Bayes classifier  $f^B(x_1, x_2, \dots, x_k)$ , a necessary condition for the classifier trained by AdaBoost based on stumps can converge to  $f^B$  as the number of iterations  $m \rightarrow \infty$  is: the population is comonotonic.

**Proof.** The AdaBoost.M1 algorithm in [Hastie et al. \(2009, p. 339\)](#) shows that, if the number of iterations is  $m$ , the final strong classifier  $f^{(m)}$  must take the form

$$f^{(m)}(x_1, x_2, \dots, x_k) = \text{sign} \left[ \sum_{s=1}^m \alpha_s g_s(x_1, x_2, \dots, x_k) \right],$$

where  $g_s(\cdot), s = 1, \dots, m$  are base learners (stumps). In other words,  $f^{(s)}(\cdot)$  must be a linear combination of base learners.

A stump “A” with  $k$  variables can be expressed as

$$\text{stump}_A(x_1, x_2, \dots, x_k) = \text{sign}(x_i \bowtie c) := \begin{cases} 1, & x_i \bowtie c, \\ -1, & \text{otherwise,} \end{cases}$$

where  $\bowtie \in \{\geq, \leq, >, <\}$ , and the splitting variable of the stump is the  $i$ th feature.

Without loss of generality, we require that  $\bowtie$  can only be  $\leq$  or  $>$ . Then the linear combination of stumps trained by AdaBoost can be represented as

$$f^{(m)}(x_1, x_2, \dots, x_k) = \text{sign} \left[ L + \sum_{s=1}^m \alpha_s \text{sign}(x_{i_s} \bowtie c_s) \right].$$

where  $L$  is a constant,  $x_{i_s}$  is the splitting variable of the  $s$ th stump, and  $c_s$  is the splitting value of the  $s$ th stump. Since  $\text{sign}(x_i > c) = -\text{sign}(x_i \leq c)$ , we can adjust all inequality signs to the same direction

$$f^{(m)}(x_1, x_2, \dots, x_k) = \text{sign} \left[ L + \sum_{s=1}^m \alpha_s \text{sign}(x_{i_s} \leq c_s) \right].$$

For simplicity, let us consider the 2-dim case ( $k = 2$ ). One can generalize the following conclusions to high-dimensional spaces similarly. According to the splitting variable of each stump, we can separate the  $m$  stumps into two groups as

$$f^{(m)}(x_1, x_2) = \text{sign} \left[ L + \left( \sum_{i=1}^{m_1} \beta_i \text{sign}(x_1 \leq a_i) \right) + \left( \sum_{j=1}^{m_2} \gamma_j \text{sign}(x_2 \leq b_j) \right) \right],$$

where  $m_1$  is the number of stumps with  $x_1$  as the splitting variable,  $m_2$  is the number of stumps with  $x_2$  as the splitting variable, and  $m_1 + m_2 = m$ . Without loss of generality, we assume that  $a_1 \leq a_2 \leq \dots \leq a_{m_1}$ , and  $b_1 \leq b_2 \leq \dots \leq b_{m_2}$ .

Recall the definition of comonotonicity (Definition 8). For any constant  $c$ , take any  $\varepsilon > 0$ ,  $a \in (c - \varepsilon, c)$  and  $b \in (c, c + \varepsilon)$ . We sort  $a, b, c$  and  $a_1, \dots, a_{m_1}$  together as

$$a_{m_1, a} \leq a \leq a_{m_1, a+1} \leq \dots \leq a_{m_1, c} \leq c \leq a_{m_1, c+1} \leq \dots \leq a_{m_1, b} \leq b \leq a_{m_1, b+1}.$$

Then, from the expression of  $f^{(m)}(x_1, x_2)$ , we have

$$f^{(m)}(a, x_2) - f^{(m)}(b, x_2) \equiv \sum_{i=m_1, b}^{m_1, a} \beta_i, \quad \forall x_2,$$

i.e., it does not depend on  $x_2$ . Similarly, we also have that  $f^{(m)}(x_1, a) - f^{(m)}(x_1, b)$  does not depend on  $x_1$ . Let  $m \rightarrow \infty$ , and if the algorithm will converge, then

$$\lim_{m \rightarrow \infty} [f^{(m)}(a, x_2) - f^{(m)}(b, x_2)]$$

and

$$\lim_{m \rightarrow \infty} [f^{(m)}(x_1, a) - f^{(m)}(x_1, b)]$$

will also be constants which do not depend on  $x_2$  and  $x_1$ , respectively. Therefore, according to the definition of comonotonicity,  $f^{(m)}$  cannot converge to  $f^B$  if the population is not comonotonic.  $\square$

To show the intuition of  $f^{(m)}$  in the proof above, let  $f(x_1, x_2) = L + (\sum_{i=1}^{m_1} \beta_i \text{sign}(x_1 \leq a_i)) + (\sum_{j=1}^{m_2} \gamma_j \text{sign}(x_2 \leq b_j))$ . Figure 7 illustrates the property of the final strong classifier  $f$ .

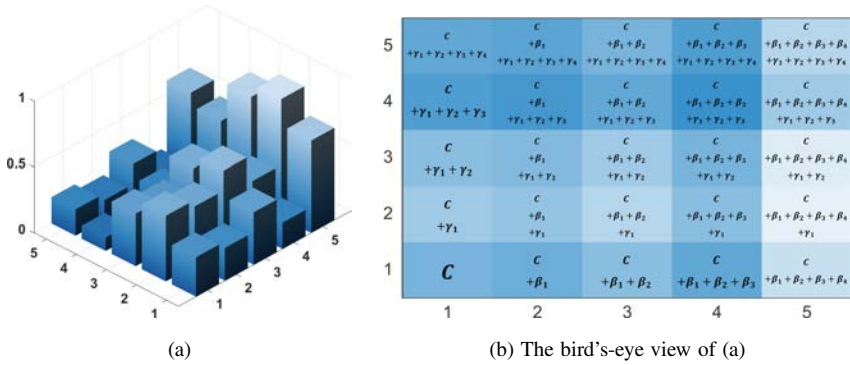


Fig. 7. An example of the strong classifier  $f$ .

Figure 7 is a toy example of the strong classifier  $f$  with  $M_1 = 5$  and  $M_2 = 5$ . Figure 7(a) is the graph of the function  $f(x_1, x_2)$ , and 7(b) is the bird's-eye view of 7(a). The darker the color is, the smaller the value of  $f$  takes. The value of  $f$  is written explicitly in 7(b), which shows that, the values in row 2 are  $\gamma_1$  greater than row 1, and the values in row 3 are  $\gamma_2$  greater than row 2, and so on. Similarly, the values in column 2 are  $\beta_1$  greater than column 1, and the values in column 3 are  $\beta_2$  greater than column 2. All numbers in the grids increase or decrease *the same* values, from left to right, and from bottom to top. It is the pattern of comonotonicity.

We have already shown in Fig. 6(a) that the XOR function is not comonotonic. Therefore, if the Bayes classifier of a population is the XOR function, it is impossible to give a good answer to the classification problem by training AdaBoost based on stumps. The conclusion in this section can be regarded as a generalization of Sec. 3.1.

In portfolio management, it is very common that factors may have interactions among each others. Hence, non-comonotonic populations are not rare. Although AdaBoost based on stumps can achieve good results in some areas, in financial studies, just based on stumps is far from reaching the desired goal. In Sec. 4, we use empirical studies to show that, using deeper trees as base learners of AdaBoost is usually a better choice in portfolio management.

### 4. Empirical Studies

In this section, we use the data of the Chinese A-share market to give empirical studies about the factor investing strategy based on AdaBoost. How to construct a stock factor strategy is an open problem with long history in portfolio management. From Wang et al. (2012) invented the N-LASR to Fiévet and

Sornette (2018) proposed a decision tree forecasting model, and Gu *et al.* (2020) and D'Hondt *et al.* (2020) gave a comprehensive analysis of machine learning methods for the canonical problem of empirical asset pricing, all of them agree that it may improve the strategy performance if the prediction model can dig out nonlinear and complex information.

Our empirical studies have two goals. On the one hand, by selecting an optimal portfolio management strategy based on AdaBoost, we want to verify the general theoretical results about the interpolation and localization of AdaBoost in Secs. 2 and 3. On the other hand, we want to illustrate the good performance of the equal-weighted strategy based on AdaBoost.

In order to achieve the first goal, we give a sensitivity analysis about the depth of the base learners (decision trees) and the number of iterations of AdaBoost on the training set and the test set. We specifically explain the performance of AdaBoost that it can dig out useful information efficiently, as well as decrease the test error.

#### **4.1. Data**

The empirical data starts in June 2002 and ends in June 2017, 181 months in total. All stocks traded in the Chinese A-share market are included. 60 factors are used in our strategy. The data of the factor exposures and the monthly returns are downloaded from the Wind Financial Terminal.<sup>5</sup> The 60 factors include not only the fundamental factors, but also the technical factors, such as the momentum and the turnover. All 60 factors are listed in Table 2.

The original data has been preliminarily cleaned, but we still need to do some preprocessing before training. We remove all stocks which are not traded (or cannot be traded due to the limit-up or limit-down in the Chinese market) during the period we study. We remove the factors with over 10% missing data, and fill in the missing data of other factors with 0. For each month, we assign the response variables  $Y$  of all stocks according to their ranks of the next-month returns cross-sectionally. The response variables of the top 50% stocks are +1, and that of the bottom 50% stocks are -1.

We divide all data into a training set and a test set manually. Total 181 months' data is divided into two sets: the first 127 months' data (June 2002–December 2012) is taken as the training set, and the last 54 months' data (January 2013–June 2017) is taken as the test set. Then, the size of the training set is 1,93,455 (sum of the stock numbers in all months), and the size of the test set is 1,33,277. We use

---

<sup>5</sup><https://www.wind.com.cn/en/Default.html>.

Table 2. The 60 factors.

alr	IR_bps_252	IR_netasset_126	IR_net_profit_63	IR_roe_252	net_assets
amount_21	IR_bps_63	IR_netasset_252	IR_oper_rev_126	IR_roe_63	oper_rev_ttm
avg_volume_63	IR_eps_126	IR_netasset_63	IR_oper_rev_252	IR_totasset_126	pb
bps	IR_eps_252	IR_net_profit_126	IR_oper_rev_63	IR_totasset_252	q_eps
IR_bps_126	IR_eps_63	IR_net_profit_252	IR_roe_126	IR_totasset_63	q_grossprofitmargin
q_netprofitmargin	q_ps	rt_252	tot_assets	tmm_pcf	turnover_126
q_oper_rev	q_roa	rt_63	tmm_eps	tmm_pe	turnover_21
q_orps	q_roe	shr_float2tot	tmm_grossprofitmargin	tmm_ps	turnover_252
q_pcf	rt_126	s_dq_mv	tmm_netprofitmargin	tmm_roa	turnover_63
q_pe	rt_21	s_val_mv	tmm_orps	tmm_roe	val_float2tot



the training set to fit models, and then use the test set to evaluate the models and verify our conclusions in Secs. 2 and 3.

## **4.2. The performance of the AdaBoost classifiers**

In this section, we analyze how the performance of the classifiers trained by AdaBoost varies with the two hyperparameters: the depth of the base learners (decision trees), and the number of iterations. Both hyperparameters are typically the source of overfitting in common sense. More specifically, we consider the following hyperparameters:

- **Max\_Depth:** The maximum depth of the base learners (decision trees), takes 2, 4, 6, 8, and 10, respectively;
- **N\_Steps:** The number of iterations, takes 10, 20, 30, 40, 50, 100, 500, and 1000, respectively.

In order to analyze the influence of these two hyperparameters on the fitting ability of AdaBoost, we fix other parameters, and set the learning rates of all models as 0.1. Both in the training set and in the test set, we use the AUC (area under the ROC curve) to measure the performance of the models, which is a supplement to the usual error calculation.

The performance results of all models we studied are summarized in Table 3. Based on the results, we observe that

- The training/test AUC and the training/test error are consistent, since if the AUC is high, the error will be low in almost every scenarios. For example, when  $\text{Max\_Depth} = 2$  and  $\text{N\_Steps} = 10$  (the first model), the training AUC is 0.5412 while the training error is 0.4701; and when  $\text{Max\_Depth} = 2$  and  $\text{N\_Steps} = 20$  (the second model), the training AUC is 0.5436 while the training error is 0.4700—They vary in different directions.
- The training AUC increases monotonically as the complexity of the model increases. Specifically, from the first model to the last model, the complexity increases. Meanwhile, the training AUC increases from 0.5412 to 0.6828; The training error decreases too.
- The test AUC also almost increases monotonically as the complexity of the model increases. For instance, when  $\text{Max\_Depth} = 2$  and  $\text{N\_Steps} = 10, \dots, 50$  (the first–fifth models), the test AUC increases from 0.5433 to 0.5480; When  $\text{N\_Steps} = 20$  and  $\text{Max\_Depth} = 2, 4, 6$  (the second, seventh and twelfth models), the test AUC also increases from 0.5462 to 0.5490.
- The changes of the test AUC are relatively small and stable with regard to that of the training AUC. For example, for the first 15 models, the test AUC changes

Table 3. Model performance results.

Model no.	Hyperparameters		Training set		Test set	
	Max_Depth	N_Steps	Training AUC	Training error	Test AUC	Test error
1	2	10	0.5412	0.4701	0.5433	0.4713
2	2	20	0.5436	0.4700	0.5462	0.4741
3	2	30	0.5499	0.4665	0.5468	0.4728
4	2	40	0.5511	0.4656	0.5476	0.4716
5	2	50	0.5533	0.4636	0.5480	0.4714
6	4	10	0.5628	0.4545	0.5463	0.4671
7	4	20	0.5682	0.4515	0.5487	0.4697
8	4	30	0.5699	0.4500	0.5489	0.4681
9	4	40	0.5713	0.4505	0.5498	0.4669
10	4	50	0.5723	0.4498	0.5500	0.4669
11	6	10	0.5818	0.4418	0.5458	0.4715
12	6	20	0.5870	0.4392	0.5490	0.4683
13	6	30	0.5913	0.4353	0.5502	0.4675
14	6	40	0.5930	0.4346	0.5506	0.4676
15	6	50	0.5946	0.4338	0.5513	0.4670
16	8	100	0.6300	0.4108	0.5519	0.4663
17	8	500	0.6356	0.4071	0.5531	0.4659
18	8	1000	0.6358	0.4071	0.5532	0.4659
19	10	100	0.6731	0.3805	0.5486	0.4679
20	10	500	0.6769	0.3780	0.5502	0.4677
21	10	1000	0.6828	0.3740	0.5499	0.4681

from 0.5433 to 0.5513, while the training AUC changes from 0.5412 to 0.5946. It suggests that the test AUC around 54% may be a stable threshold of the model, which reflects the ability of our methods to dig out the market information contained in our dataset. It is noteworthy that, 54% is not a bad result in the Chinese stock market, according to the experience of the industry.


- In the training set, the performance is more sensitive to Max\_Depth than to N\_Steps. In detail, given Max\_Depth = 2, the training AUC changes from 0.5412 to 0.5533 for N\_Steps = 10, . . . , 50 (the first–fifth models); However, given N\_Steps = 10, the training AUC changes from 0.5412 to 0.5818 for Max\_Depth = 2, 4, 6.

We find that, as the depth of the trees and the number of iterations increases, the AUC for the test set increases stably without significant change. We can conclude that, in these cases, the more iteration steps, the better the classifier, and the more complex the base learner trees, the better the classifier.

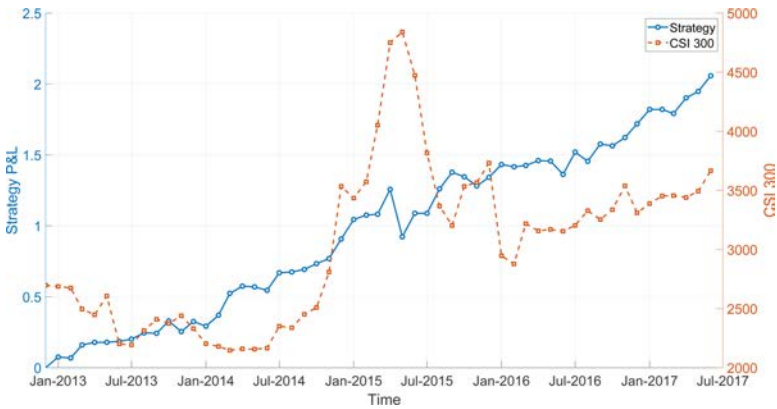
### 4.3. Strategy results

In this section, we use the classifier with good performance in the training set to establish equal-weighted strategies in the test set (January 2013–June 2017). We choose the 16th classifier in Table 3: Max\_Depth = 8, and N\_steps = 100.

Our strategy is: At the beginning of each month, we treat the last month's exposures of the 60 factors for all stocks as the input of the classifier. Then, we use the output of the classifier to determine whether to take a long or a short position. For the long-short strategy, we long the top 50 stocks with the highest probability of going up, and short the bottom 50 stocks with the highest probability of going down; For the long-only strategy, we only long the top 50 stocks with the highest probability of going up.



(a) Long-only strategy



(b) Long-short strategy

Fig. 8. The P&L of two strategies.

Table 4. Strategy performance results.

Strategy	Win rate	Sharpe ratio	Average return (%)	Std. (%)	Maximum drawdown (%)
Long-only	0.704	1.665	47.62	28.60	24.18
Long-short	0.704	1.568	45.74	29.17	33.41
CSI 300	0.537	0.243	7.99	32.83	72.80

In Fig. 8, we use the CSI 300 to denote the benchmark of the buy-and-hold strategy which we want to compare with. Figures 8(a) and 8(b) show the P&Ls of the long-only strategy and the long-short strategy, respectively. Both the transaction cost rates of taking a long position and a short position are set as 0.15%. Table 4 shows the performance summaries of our strategies, including the win rate, the Sharpe ratio, the average return, the standard deviation of return, and the maximum drawdown as different measures of the performance.

From the strategy performance, we have the following results:

- The performance of our strategies is remarkable, although the AdaBoost classifier is fixed for the whole test period (January 2013–June 2017). The average annualized return for both the long-only and the long-short strategy are over 45% during the test period. Particularly, the Chinese A-share market experienced a dramatic fluctuation in 2015, but our strategies (the blue solid line) still have a relatively robust return comparing with the CSI 300 benchmark (the orange dashed line).
- The performance of our strategies is relatively stable, and even the maximum drawdown of the long-only strategy is much lower than the maximum drawdown of the benchmark CSI 300 (72.80%). The maximum drawdown 24.18% of the long-only strategy occurred in August 2015, and 33.41% of the long-short strategy occurred in May 2015.

Once again, we should note that the goal of these empirical studies is to find an optimal strategy by AdaBoost which is a complicated classifier in-sample, as well as to demonstrate that the performance of the complicated model is not bad out-of-sample.

Overall, we can conclude from the empirical studies:

- For the equal-weighted factor investing strategy, increasing the depth of the base learners (decision trees) and the number of iterations of AdaBoost do not significantly reduce the out-of-sample performance;
- For factor investing strategies, it is a feasible way for investors to use the complicated AdaBoost method to learn in the training set.

## 5. Conclusion

In order to implement machine learning in constructing equal-weighted portfolios interpretatively, our paper explains the success of AdaBoost and applies it in portfolio management.

We prove that a “complex” method, even interpolating, does not necessarily result in poor out-of-sample performance if it can lower the influence of the noise points. AdaBoost based on large trees without early stopping could be one of such methods, and we explain its success with our newly defined ION.

We illustrate the shortcomings of AdaBoost based on shallow trees from the perspective of its ability to dig out nonlinear information, and emphasize the benefits of AdaBoost based on large trees.

Our empirical studies not only corroborate the theoretical results, but also give an effective approach to construct equal-weighted portfolios via machine learning.

Conclusively, we show that AdaBoost can minimize the influence of the noise points while interpolating the training set, and thus have a good out-of-sample performance. We confirm the conjectures in [Wyner et al. \(2017\)](#) under a mathematical framework. The empirical studies verify the potential applications of AdaBoost in portfolio management.

## References

- Asness, C, A Frazzini, R Israel, TJ Moskowitz and LH Pedersen (2018). Size matters, if you control your junk, *Journal of Financial Economics*, 129(3), 479–509.
- Breiman, L (1998). Arcing classifier (with discussion and a rejoinder by the author), *The Annals of Statistics*, 26(3), 801–849.
- Breiman, L (1999). Prediction games and arcing algorithms, *Neural Computation*, 11(7), 1493–1517.
- Breiman, L (2001). Random forests, *Machine Learning*, 45(1), 5–32.
- Chen, T and C Guestrin (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, New York, NY, USA, pp. 785–794. ACM.
- Cover, T and P Hart (1967). Nearest neighbor pattern classification, *IEEE Transactions on Information Theory*, 13(1), 21–27.
- Creamer, G and Y Freund (2005). Using adaboost for equity investment scorecards. Howe School Research Paper. NIPS Workshop Machine Learning in Finance, 2005, Whistler, British Columbia, Canada.
- Creamer, G and Y Freund (2010). Automated trading with boosting and expert weighting, *Quantitative Finance*, 10(4), 401–420.
- Creamer, G (2012). Model calibration and automated trading agent for euro futures, *Quantitative Finance*, 12(4), 531–545.

- DeMiguel, V, L Garlappi and R Uppal (2007). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, 22(5), 1915–1953.
- Devroye, L, L Györfi and G Lugosi (1996). *A Probabilistic Theory of Pattern Recognition*, Stochastic Modelling and Applied Probability. Vol. 31. New York: Springer-Verlag.
- D’Hondt, C, R De Winne, E Ghysels and S Raymond (2020). Artificial intelligence alter egos: Who might benefit from robo-investing? *Journal of Empirical Finance*, 59, 278–299.
- Feng, G, S Giglio and D Xiu (2020). Taming the factor zoo: A test of new factors, *The Journal of Finance*, 75(3), 1327–1370.
- Fiévet, L and D Sornette (2018). Decision trees unearth return sign predictability in the S&P 500, *Quantitative Finance*, 18(11), 1797–1814.
- Freund, Y and RE Schapire (1996). Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on International Conference on Machine Learning*, ICML’96, San Francisco, CA, USA, pp. 148–156. Morgan Kaufmann Publishers Inc.
- Friedman, JH (2001). Greedy function approximation: A gradient boosting machine, *The Annals of Statistics*, 29(5), 1189–1232.
- Friedman, J, T Hastie and R Tibshirani (2000). Additive logistic regression: A statistical view of boosting (with discussion and a rejoinder by the authors), *The Annals of Statistics*, 28(2), 337–407.
- Gu, S, B Kelly and D Xiu (2020). Empirical asset pricing via machine learning, *The Review of Financial Studies*, 33(5), 2223–2273.
- Harvey, CR, Y Liu and H Zhu (2016). And the cross-section of expected returns, *The Review of Financial Studies*, 29(1), 5–68.
- Harvey, CR (2017). Presidential address: The scientific outlook in financial economics, *The Journal of Finance*, 72(4), 1399–1440.
- Hastie, T, R Tibshirani and JH Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. 2nd edition. New York: Springer.
- James, GM (2003). Variance and bias for general loss functions, *Machine Learning*, 51(2), 115–135.
- Jobson, JD and BM Korkie (1981). Performance hypothesis testing with the Sharpe and Treynor measures, *The Journal of Finance*, 36(4), 889–908.
- Ke, G, Q Meng, T Finley, T Wang, W Chen, W Ma, Q Ye and T-Y Liu (2017). LightGBM: A highly efficient gradient boosting decision tree. In I Guyon, UV Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 30, pp. 3146–3154. New York: Curran Associates, Inc.
- López de Prado, M (2018). *Advances in Financial Machine Learning*. New York: Wiley.
- Mease, D and A Wyner (2008). Evidence contrary to the statistical view of boosting, *Journal of Machine Learning Research*, 9, 131–156.

- Michaud, RO (1989). The Markowitz optimization enigma: Is “optimized” optimal? *Financial Analysts Journal*, 45(1), 31–42.
- Rasekhschaffe, KC and RC Jones (2019). Machine learning for stock selection, *Financial Analysts Journal*, 75(3), 70–88.
- Schapire, RE, Y Freund, P Bartlett and WS Lee (1998). Boosting the margin: A new explanation for the effectiveness of voting methods, *The Annals of Statistics*, 26(5), 1651–1686.
- Tu, J and G Zhou (2011). Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies, *Journal of Financial Economics*, 99(1), 204–215.
- Wang, S, Y Luo, R Cahan, M-A Alvarez, J Jussa and Z Chen (2012). Signal processing: The rise of the machines, *Deutsche Bank Quantitative Strategy* (5 June).
- Wyner, AJ, M Olson, J Bleich and D Mease (2017). Explaining the success of AdaBoost and random forests as interpolating classifiers, *Journal of Machine Learning Research*, 18(48), 1–33.
- Wyner, AJ (2003). On boosting and the exponential loss. In *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics, AISTATS 2003*, Key West, Florida, USA, January 3–6.
- Zhang, H, B Jacobsen, F de Roon and F Jiang (2019). Boosting portfolio choice in the big data era. SSRN 3495901.