

# Gradient and Newton Boosting for Classification and Regression

Fabio Sigrist\*

Lucerne University of Applied Sciences and Arts

October 21, 2020

## Abstract

Boosting algorithms are frequently used in applied data science and in research. To date, the distinction between boosting with either gradient descent or second-order Newton updates is often not made in both applied and methodological research, and it is thus implicitly assumed that the difference is irrelevant. The goal of this article is to clarify this situation. In particular, we present gradient and Newton boosting, as well as a hybrid variant of the two, in a unified framework. We compare these boosting algorithms with trees as base learners using various datasets and loss functions. Our experiments show that Newton boosting outperforms gradient and hybrid gradient-Newton boosting in terms of predictive accuracy on the majority of datasets. We also present evidence that the reason for this is not faster convergence of Newton boosting. In addition, we introduce a novel tuning parameter for tree-based Newton boosting which is interpretable and important for predictive accuracy.

*Keywords:* boosting, supervised learning, ensembles, trees

## 1 Introduction

Boosting (Freund et al., 1996; Friedman et al., 2000; Friedman, 2001) refers to a type of supervised learning algorithms that enjoy high popularity in applied data science and research, among other things, due to their high predictive accuracy (Chen and Guestrin, 2016). This is reflected in statements such as “[i]n general ‘boosted decision trees’ is regarded as the most effective off-the-shelf nonlinear learning method for a wide range of application problems” (Johnson and Zhang, 2013). Boosting iteratively adds so-called base learners to an ensemble of learners. Broadly speaking, there exist three different versions for selecting a base learner in every boosting iteration: functional gradient descent, a functional version of

---

\*Email: fabio.sigrist@hslu.ch. Address: Lucerne University of Applied Sciences and Arts, Suurstoffi 1, 6343 Rotkreuz, Switzerland. Phone: +41 41 757 67 61.

Newton’s method, and a combination of the two. We refer to these three different versions of boosting as *gradient boosting*, *Newton boosting*, and *hybrid gradient-Newton boosting*; see Section 2 for more information.

In both methodological and applied research, the distinction between gradient and Newton boosting is often not made and/or it is not declared which version of boosting is used (e.g. Ahamad et al., 2020; Djeundje et al., 2020; Moscatelli et al., 2020). It is thus implicitly assumed that the difference is not important. For instance, the two recent popular boosting libraries **LightGBM** and **TF Boosted Trees** do not distinguish in their companion articles (Ke et al., 2017; Ponomareva et al., 2017) between gradient and Newton boosting, and it is unclear to the reader which version is used. Similarly, Prokhorenkova et al. (2018) briefly mention in their article on **CatBoost** that the minimization for finding a boosting update can be done using the Newton method or with a gradient step, and then continue to write that “[b]oth methods are kinds of functional gradient descent”. However, Newton’s method is different from gradient descent. Further, Bühlmann and Hothorn (2007) state that for gradient boosting “an additional line search ... seems unnecessary for achieving a good estimator.” For trees as base learners, the additional line search is often done for each leaf separately by using a Newton step (Friedman, 2001). I.e., this corresponds to what we denote as hybrid gradient-Newton boosting which is different from plain gradient boosting also in terms of predictive accuracy. Besides, particular software implementations of boosting such as **XGBoost** (Chen and Guestrin, 2016) are sometimes presented as if they were separate boosting algorithms (e.g. Xia et al., 2017; Ahamad et al., 2020; Djeundje et al., 2020) when, in fact, they implement a particular version of boosting.

The novel contributions of this article are the following ones. First, we show how gradient, Newton, as well as hybrid gradient-Newton boosting can be derived in a unified framework. Further, we systematically compare gradient, Newton, and hybrid gradient-Newton boosting on a large set of both real-world and simulated classification and regression datasets. In our experiments, using trees as base learners, we find that Newton boosting achieves lower test errors than both gradient boosting and hybrid gradient-Newton boosting, and hybrid gradient-Newton boosting often has higher predictive accuracy than gradient boosting. Interestingly, we find that Newton boosting results in both lower in-sample training losses, which are essentially zero for most classification datasets, and lower out-of-sample test errors for most datasets. We also present evidence that the higher pre-

dictive accuracy is not due to a faster convergence speed of Newton boosting. In addition, we introduce a novel tuning parameter for Newton boosting with trees as base learners. We argue that this minimum *equivalent sample size per leaf* parameter is a natural and interpretable tuning parameter which is important for predictive accuracy. In particular, we present evidence that the unnormalized version of this tuning parameter currently adopted in popular software implementations such as `XGBoost` is difficult to tune and thus likely results in lower predictive accuracy.

## 1.1 Related work

The first boosting algorithms for classification, including the well known AdaBoost algorithm, were introduced by Schapire (1990), Freund and Schapire (1995), and Freund et al. (1996). Later, several authors (Breiman, 1998, 1999; Friedman et al., 2000; Mason et al., 2000; Friedman, 2001) presented the statistical view of boosting as a stagewise optimization approach. See Schapire (2003), Bühlmann and Hothorn (2007), Schapire and Freund (2012), Mayr et al. (2014a), and Mayr et al. (2014b) for reviews on boosting algorithms in both the machine learning and statistical literature.

To the best of our knowledge, a systematic comparison concerning the predictive accuracy of gradient, Newton, and hybrid gradient-Newton boosting for various choices of loss functions, including regression and classification losses, has not been done so far. The  $L_K$ -TreeBoost algorithm (Friedman, 2001) is compared in Friedman (2001) with  $K$ -class LogitBoost (Friedman et al., 2000) for classification with five classes in a simulation study for one type of random functions. In our terminology,  $L_K$ -TreeBoost is a version of hybrid gradient-Newton boosting, and  $K$ -class LogitBoost corresponds to Newton boosting for the Bernoulli likelihood. Friedman (2001) finds that the algorithms perform “nearly the same” with “LogitBoost perhaps having a slight advantage”. In addition, it is mentioned that “it is likely that when the shrinkage parameter is carefully tuned for each of the three methods [ $L_K$ -TreeBoost,  $K$ -class LogitBoost, AdaBoost], there would be little performance differential between them.” Our empirical evidence is not in line with this statement. Saberian et al. (2011) also briefly compare variants of boosting with gradient and second-order updates using three different binary classification datasets and Haar wavelets as base learners. However, their boosting approach is different from the one usually adopted in practice and research in the sense that they assume normed based learners, find base learners as maxima

of inner products of gradients and base learners, and then have to perform an additional line search to find the step size. Further, tuning parameters such as the learning rate and the number of boosting iterations are not chosen using cross-validation, and only 25 boosting iterations are performed. Nonetheless, they come to the same conclusion as we do, i.e., they find that their version of second-order boosting performs better than gradient boosting. The closest to our empirical analysis are Li (2010) and Zheng and Liu (2012). Li (2010) compares Newton boosting (“logitboost”) with hybrid gradient-Newton boosting (“mart”) for several multi-class classification datasets and also finds that Newton results in lower test errors than hybrid gradient-Newton boosting. Further, Zheng and Liu (2012) compare gradient and Newton boosting when using the probit link function in a logistic regression model and find that Newton boosting results in lower error rates than gradient boosting for several classification applications including face detection, cancer classification, and handwritten digit recognition. However, both Li (2010) and Zheng and Liu (2012) consider only specific classification tasks, tuning parameters are not chosen using validation data in their experiments, and it is not investigated whether the observed differences are statistically significant. Finally, Sun et al. (2014) compare Newton and gradient boosting for binary classification using the logistic loss. Their focus is on the convergence rate and their empirical comparison only considers training errors, though.

## **2 The statistical view of boosting: three approaches for stage-wise optimization**

In this section, we present the statistical view of boosting as finding the minimizer of a risk functional in a function space using a stagewise, or greedy, optimization approach. We distinguish between gradient and Newton boosting as well as a hybrid version of the two and show how these can be presented in a unified framework. Note that these boosting algorithms have been proposed in prior research (Breiman, 1998, 1999; Friedman et al., 2000; Mason et al., 2000; Friedman, 2001; Saberian et al., 2011), but, to the best of our knowledge, the presentation below in a unified framework and the extension to the multivariate case is novel.

## 2.1 Population versions

We assume that there is a response variable  $Y \in \mathbb{R}$  and a vector of  $p$  predictor variables  $X \in \mathbb{R}^p$ .<sup>1</sup> Our goal is to predict the response variable using the predictor variables, where predictions can be both deterministic or probabilistic. We assume that  $(Y, X)$  are random variables on  $\mathbb{R} \times \mathbb{R}^p$ , and both the distribution of  $X$  and the conditional distribution  $Y|X$  are absolutely continuous with respect to either the Lebesgue measure, a counting measure, a mixture of both, or a product measure of the former measures. In particular, this covers both regression and classification tasks or mixtures of the two such as Tobit regression (Sigrist and Hirnschall, 2019).

The goal of boosting is to find a minimizer  $F^*$  of the risk  $R(F)$  which is defined as the expected loss

$$R(F) = E_{Y,X}(L(Y, F(X))), \quad (1)$$

where  $F(\cdot)$  is a function in a Hilbert space  $\mathcal{H}$  with inner product  $\langle \cdot, \cdot \rangle$  given by

$$\langle F, F \rangle = E_X (F(X)^2),$$

and  $L(Y, F)$  is a loss function. See below and Appendix A for examples of loss functions. For notational simplicity, we often denote a function  $F(\cdot)$  shortly by  $F$  in this article. In general,  $F$  can also be a multivariate function in a direct sum Hilbert space. However, for notational simplicity, we assume in the following that  $F$  is univariate. In Section 2.3, we extend this to the multivariate case.

Boosting assumes that the minimizer  $F^* \in \Omega_{\mathcal{S}}$  lies in the span  $\Omega_{\mathcal{S}} = \text{span}(\mathcal{S})$  of a set  $\mathcal{S}$  of base learners  $f_j : \mathbb{R}^p \rightarrow \mathbb{R}$ :

$$F^* = \underset{F \in \Omega_{\mathcal{S}}}{\operatorname{argmin}} R(F). \quad (2)$$

If the risk  $R(F)$  is convex in  $F$ , then (2) is a convex optimization problem since  $\Omega_{\mathcal{S}}$  is also convex. Boosting finds  $F^*$  in a stagewise way by sequentially adding an update  $f_m$  to the current estimate  $F_{m-1}$ ,

$$F_m(x) = F_{m-1}(x) + f_m(x), \quad f_m \in \mathcal{S}, \quad m = 1, \dots, M, \quad (3)$$

---

<sup>1</sup>For the sake of simplicity, we focus on univariate  $Y \in \mathbb{R}$ . The extension to the case of a multivariate response variable  $Y$  is straightforward. See also Section 2.3 where we present multivariate versions of boosting.

such that the risk is minimized

$$f_m = \operatorname{argmin}_{f \in \mathcal{S}} R(F_{m-1} + f). \quad (4)$$

This minimization can often not be done analytically and an approximation has to be used.

Different boosting algorithms vary in the way the minimization in (4) is done, the loss function  $L$  used in (1), and in the choice of base learners  $f_j \in \mathcal{S}$ . Concerning loss functions, potential choices include the squared loss  $L(y, F) = (y - F)^2/2$  for regression, the negative Gaussian log-likelihood where both the mean and the scale parameter depend on predictor variables (see, e.g., Mayr et al., 2012), the negative log-likelihood  $-yF + \log(1 + e^F)$  of a binomial model with a logistic link function for binary classification, or the entropy loss with a softmax function for multiclass classification. Under appropriate regularity assumptions, one can use the negative log-likelihood of any statistical model as loss function:

$$L(y, F) = -\log(f_{F, \theta}(y)),$$

where  $f_{F, \theta}(y)$  is the density of  $Y$  given  $X$  with respect to some reference measure,  $F$  is linked to one or several, possibly transformed, parameters of this density, and  $\theta$  are additional parameters. See Appendix A for various examples of loss functions and, in particular, the ones we consider in the empirical evaluation of this article. As shown by Friedman et al. (2000), AdaBoost algorithms are versions of Newton boosting for classification with an exponential loss function.

Concerning base learners, regression trees (Breiman et al., 1984) is the most frequently adopted choice. Other potential base learners include splines or linear functions (Bühlmann and Yu, 2003; Bühlmann et al., 2006; Schmid and Hothorn, 2008). In this article, we focus on trees:

$$f(x) = w_{s(x)},$$

where  $s : \mathbb{R}^p \rightarrow \{1, \dots, J\}$ ,  $w \in \mathbb{R}^J$ , and  $J \in \mathbb{N}$  denotes the number of terminal nodes, or leaves, of the tree  $f(x)$ . The function  $s$  represents the structure of the tree, i.e., the partition of the space  $\mathbb{R}^p$ , and  $w$  contains the values of the leaves. As in Breiman et al. (1984), we assume that the partition of the space made by  $s$  is a binary tree where each cell in the partition is a rectangle of the form  $R_j = (l_1, u_1] \times \dots \times (l_p, u_p] \subset \mathbb{R}^p$  with  $-\infty \leq l_m < u_m \leq \infty$  and  $s(x) = j$  if  $x \in R_j$ .

For finding an update in (3), either a form of gradient descent, Newton's method, or a hybrid variant is used to obtain an approximate solution to the minimization problem in (4). In the following, we describe these approaches.

### 2.1.1 Gradient boosting

Assuming that the risk  $R(F)$  is Gâteaux differentiable for all  $F \in \Omega_{\mathcal{S}}$ , we denote the Gâteaux derivative by

$$\begin{aligned} dR(F, f) &= \left. \frac{d}{d\epsilon} R(F + \epsilon f) \right|_{\epsilon=0} \\ &= \lim_{\epsilon \rightarrow 0} \frac{R(F + \epsilon f) - R(F)}{\epsilon}, \quad F, f \in \Omega_{\mathcal{S}}. \end{aligned}$$

Gradient boosting then works by choosing  $f_m$  as the minimizer of a first-order Taylor approximation around  $F_{m-1}$  with a penalty on the norm of the base learner:

$$\begin{aligned} f_m &= \operatorname{argmin}_{f \in \mathcal{S}} R(F_{m-1}) + dR(F_{m-1}, f) + \frac{1}{2} \langle f, f \rangle \\ &= \operatorname{argmin}_{f \in \mathcal{S}} dR(F_{m-1}, f) + \frac{1}{2} \langle f, f \rangle. \end{aligned} \tag{5}$$

Note that we add the penalty  $\frac{1}{2} \langle f, f \rangle$  since the functions  $f$  are not necessarily normed and  $\langle f, f \rangle$  is not assumed to be constant.

If we assume that  $L(Y, F)$  is differentiable in  $F$  for P-almost all  $X$  and that the derivative is integrable with respect to the measure of  $(Y, X)$ , then  $dR(F_{m-1}, f)$  is given by

$$dR(F_{m-1}, f) = E_{Y,X} (g_m(Y, X) f(X)),$$

where  $g_m(Y, X)$  denotes the gradient of the loss function  $L(Y, F)$  with respect to  $F$  at the current estimate  $F_{m-1}$ :

$$g_m(Y, X) = \left. \frac{\partial L(Y, F)}{\partial F} \right|_{F=F_{m-1}(X)}. \tag{6}$$

Consequently, (5) can be written as

$$\begin{aligned} f_m &= \operatorname{argmin}_{f \in \mathcal{S}} E_{Y,X} \left( g_m(Y, X) f(X) + \frac{1}{2} f(X)^2 \right) \\ &= \operatorname{argmin}_{f \in \mathcal{S}} E_{Y,X} \left( (-g_m(Y, X) - f(X))^2 \right). \end{aligned} \tag{7}$$

This shows that  $f_m$  is the  $L^2$  approximation to the negative gradient  $-g_m(Y, X)$  of the loss function  $L(Y, F)$  with respect to  $F$  evaluated at the current estimate  $F_{m-1}(X)$ .

If the following expression is well defined for P-almost all  $X$ , then the minimization in (7) can also be done pointwise

$$f_m(X) = \operatorname{argmin}_{f \in \mathcal{S}} E_{Y|X} \left( (-g_m(Y, X) - f(X))^2 \right).$$

### 2.1.2 Newton boosting

For Newton boosting, we assume that  $R(F)$  is two times Gâteaux differentiable and denote the second Gâteaux derivative by

$$d^2 R(F, f) = \frac{d^2}{d\epsilon^2} R(F + \epsilon f) \Big|_{\epsilon=0}, \quad F, f \in \Omega_{\mathcal{S}}.$$

Newton boosting chooses  $f_m$  as the minimizer of a second-order Taylor approximation around  $F_{m-1}$ :

$$f_m = \operatorname{argmin}_{f \in \mathcal{S}} R(F_{m-1}) + dR(F_{m-1}, f) + \frac{1}{2} d^2 R(F_{m-1}, f). \quad (8)$$

If we assume the P-almost all existence and integrability of the second derivative of  $L(Y, F)$  with respect to  $F$ , then (8) can be written as

$$\begin{aligned} f_m &= \operatorname{argmin}_{f \in \mathcal{S}} E_{Y,X} \left( g_m(Y, X) f(X) + \frac{1}{2} h_m(Y, X) f(X)^2 \right) \\ &= \operatorname{argmin}_{f \in \mathcal{S}} E_{Y,X} \left( h_m(Y, X) \left( -\frac{g_m(Y, X)}{h_m(Y, X)} - f(X) \right)^2 \right), \end{aligned} \quad (9)$$

where the gradient  $g_m(Y, X)$  is defined in (6) and  $h_m(Y, X)$  is the second derivative of  $L(Y, F)$  with respect to  $F$  at  $F_{m-1}$ :

$$h_m(Y, X) = \frac{\partial^2 L(Y, F)}{\partial F^2} \Big|_{F=F_{m-1}(X)}. \quad (10)$$

The last line in Equation (9) shows that  $f_m$  is the weighted  $L^2$  approximation to negative ratio of the gradient over the Hessian  $-\frac{g_m(Y, X)}{h_m(Y, X)}$  and the weights corresponds to the second derivative  $h_m(Y, X)$ .

If the following expression is well defined for P-almost all  $X$ , we can again calculate the pointwise minimizer of (9) as:

$$f_m(X) = \operatorname{argmin}_{f \in \mathcal{S}} E_{Y|X} \left( h_m(Y, X) \left( -\frac{g_m(Y, X)}{h_m(Y, X)} - f(X) \right)^2 \right).$$

Note that gradient boosting can be seen as a special case of Newton boosting. If the second derivative of the loss function  $h_m(Y, X)$  exists and is constant,  $h_m(Y, X) = c \in$



$\mathbb{R} \setminus \{0\}$ , for P-almost all  $X$ , then the Newton boosting update in (9) essentially equals the gradient update in (7). Specifically, they are exactly equal if  $h_m(Y, X) = 1$ . Since in practice the update is usually damped, see Equation (16) in Section 2.4, and the shrinkage parameter  $\nu$  is considered a tuning parameter, the two approaches are essentially also equivalent for  $h_m(Y, X) = c \neq 1$ .

### 2.1.3 Hybrid gradient-Newton boosting

A hybrid variant of gradient and Newton boosting proposed in Friedman (2001) is obtained by first learning part of the parameters of the base learner using a gradient step and the remaining part using a Newton update. For instance, for trees as bases learners, the structure  $s$  of a tree is learned using a gradient update:

$$s_m = \operatorname{argmin}_{s: f=w_s \in \mathcal{S}} E_{Y,X} \left( (-g_m(Y, X) - f(X))^2 \right),$$

and then, conditional on this, one finds the weights  $w$  using a Newton step:

$$w_m = \operatorname{argmin}_{\substack{w: f=w_s \in \mathcal{S} \\ s=s_m}} E_{Y,X} \left( g_m(Y, X)f(X) + \frac{1}{2}h_m(Y, X)f(X)^2 \right).$$

### 2.1.4 Line search

The update step in (3) is sometimes presented in the form  $F_m(x) = F_{m-1}(x) + \rho_m f_m(x)$  with  $\rho_m \in \mathbb{R}$ , where  $\rho_m$  is found by doing an additional line-search  $\rho_m = \operatorname{argmin}_{\rho \in \mathbb{R}} R^e(F_{m-1} + \rho f_m)$ . For gradient boosting, this has the advantage that the length of the gradient does not depend on the scaling of the loss function. However, we are not considering this approach explicitly here since, first, we assume that the set of base learners  $\mathcal{S}$  is rich enough to include not just normalized base learners but base learners of any norm and, second, the line-search often cannot be done analytically and a second-order Taylor approximation is used instead. I.e., the latter case corresponds to a version of hybrid gradient-Newton or Newton boosting.

### 2.1.5 Applicability of Newton boosting

As mentioned in Section 2.1.2, there is no difference between the three above presented optimization approaches for loss functions with non-zero and constant second derivatives in  $F$ . In particular, this holds true for the squared loss function. Further, for loss functions

where the second derivative is zero on a non-null set of the support of  $X$ , such as the least absolute deviation (LAD), any other quantile regression loss function, and the Huber loss, Newton and also hybrid gradient-Newton boosting are not applicable. In these cases, the above-mentioned line-search can be useful in addition to a pure gradient step. Similarly, if a loss function is not  $P$ -almost everywhere twice differentiable in  $F$ , Newton boosting is also not applicable. However, the majority of commonly used loss functions are twice differentiable.

## 2.2 Sample versions

In the following, we assume that we observe  $n$  samples  $(y_i, x_i), i = 1, \dots, n$ , from the same distribution as the one of  $(Y, X)$ , and approximate the risk  $R(F)$  in (1) with the empirical risk  $R^e(F)$  obtained by replacing the population distribution with the empirical distribution:

$$R^e(F) = \frac{1}{n} \sum_{i=1}^n L(y_i, F(x_i)). \quad (11)$$

For gradient boosting, the sample version of (7) can be written as

$$\begin{aligned} f_m &= \operatorname{argmin}_{f \in \mathcal{S}} \sum_{i=1}^n g_{m,i} f(x_i) + \frac{1}{2} f(x_i)^2 \\ &= \operatorname{argmin}_{f \in \mathcal{S}} \sum_{i=1}^n (-g_{m,i} - f(x_i))^2, \end{aligned} \quad (12)$$

where  $g_{m,i}$  is the gradient of the loss function for observation  $i$

$$g_{m,i} = \left. \frac{\partial}{\partial F} L(y_i, F) \right|_{F=F_{m-1}(x_i)}.$$

This means that the stagewise minimizer  $f_m$  can be found as the least squares approximation to the negative gradient  $-g_{m,i}$ .

Similarly, the sample version of the Newton update in (9) is given by

$$\begin{aligned} f_m &= \operatorname{argmin}_{f \in \mathcal{S}} \sum_{i=1}^n g_{m,i} f(x_i) + h_{m,i} \frac{1}{2} f(x_i)^2 \\ &= \operatorname{argmin}_{f \in \mathcal{S}} \sum_{i=1}^n h_{m,i} \left( -\frac{g_{m,i}}{h_{m,i}} - f(x_i) \right)^2, \end{aligned} \quad (13)$$

where  $h_{m,i}$  is the Hessian of the loss function for observation  $i$ :

$$h_{m,i} = \left. \frac{\partial^2}{\partial F^2} L(y_i, F) \right|_{F=F_{m-1}(x_i)}. \quad (14)$$

I.e.,  $f_m$  can be found as the weighted least squares approximation to the ratio of the negative gradient over the Hessian  $-\frac{g_{m,i}}{h_{m,i}}$  with weights given by  $h_{m,i}$ .

The sample version of the hybrid gradient-Newton algorithm first finds the structure  $s$  of a tree using a gradient step:

$$s_m = \operatorname{argmin}_{s: f=w_s \in \mathcal{S}} \sum_{i=1}^n (-g_{m,i} - f(x_i))^2,$$

and then determines the weights  $w$  using a Newton step:

$$w_m = \operatorname{argmin}_{w: f=w_s \in \mathcal{S}, s=s_m} \sum_{i=1}^n h_{m,i} \left( -\frac{g_{m,i}}{h_{m,i}} - f(x_i) \right)^2.$$

### 2.3 Multivariate case

In this section, we briefly present gradient and Newton boosting when the function  $\mathbf{F}$  is multivariate. In this case,

$$\mathbf{F}(X) = (F^1(X), F^2(X), \dots, F^d(X))^T$$

is assumed to be a function in a direct sum Hilbert space  $\mathcal{H} = \mathcal{H}^1 \oplus \mathcal{H}^2 \oplus \dots \oplus \mathcal{H}^d$ , where the  $\mathcal{H}^k$ 's are Hilbert spaces with inner products  $\langle \cdot, \cdot \rangle_k$  given by  $\langle F^k, F^k \rangle_k = E_X (F^k(X)^2)$ , and the inner product for  $\mathcal{H}$  is given by  $\langle \mathbf{F}, \mathbf{F} \rangle = \sum_{k=1}^d \langle F^k, F^k \rangle_k$ . For the sake of readability, we use boldface in this subsection to distinguish vector-valued functions from scalar-valued functions. Examples of loss functions where  $\mathbf{F}$  is multivariate include the entropy loss with a softmax function for multiclass classification or generalized additive models for location, scale, and shape (GAMLSS) where location, scale, and shape parameters are modeled as functions of predictor variables  $X$  (Rigby and Stasinopoulos, 2005; Mayr et al., 2012).

A gradient boosting update  $\mathbf{f}_m \in \mathcal{S} \oplus \dots \oplus \mathcal{S}$  in Equation (3) is then obtained as

$$\mathbf{f}_m = \operatorname{argmin}_{\mathbf{f} \in \mathcal{S} \oplus \dots \oplus \mathcal{S}} dR(\mathbf{F}_{m-1}, \mathbf{f}) + \frac{1}{2} \langle \mathbf{f}, \mathbf{f} \rangle. \quad (15)$$

Under appropriate regularity conditions,  $dR(\mathbf{F}_{m-1}, \mathbf{f})$  is given by

$$dR(\mathbf{F}_{m-1}, \mathbf{f}) = E_{Y,X} (\mathbf{g}_m(Y, X)^T \mathbf{f}(X)),$$

where

$$\mathbf{g}_m(Y, X) = \left( \frac{\partial}{\partial F^1} L(Y, \mathbf{F}), \dots, \frac{\partial}{\partial F^d} L(Y, \mathbf{F}) \right)^T \Big|_{\mathbf{F}=\mathbf{F}_{m-1}(X)}.$$

It follows that the solution in (15) can be determined for each  $k$ ,  $k = 1, \dots, d$ , separately as

$$f_m^k = \operatorname{argmin}_{f^k \in \mathcal{S}} E_{Y,X} \left( \left( -g_m^k(Y, X) - f^k(X) \right)^2 \right),$$

where

$$g_m^k(Y, X) = \frac{\partial}{\partial F^k} L(Y, \mathbf{F}) \Big|_{\mathbf{F}=\mathbf{F}_{m-1}(X)}.$$

The sample version of this gradient boosting update is given by

$$f_m^k = \operatorname{argmin}_{f^k \in \mathcal{S}} \sum_{i=1}^n \left( -g_{m,i}^k - f^k(x_i) \right)^2,$$

where  $g_{m,i}^k = g_m^k(y_i, x_i)$ .

Newton boosting obtains an update  $\mathbf{f}_m \in \mathcal{S} \oplus \dots \oplus \mathcal{S}$  as

$$\mathbf{f}_m = \operatorname{argmin}_{\mathbf{f} \in \mathcal{S} \oplus \dots \oplus \mathcal{S}} dR(\mathbf{F}_{m-1}, \mathbf{f}) + \frac{1}{2} d^2 R(\mathbf{F}_{m-1}, \mathbf{f}),$$

where, again under appropriate conditions, this can also be written as

$$\mathbf{f}_m = \operatorname{argmin}_{\mathbf{f} \in \mathcal{S} \oplus \dots \oplus \mathcal{S}} E_{Y,X} \left( \mathbf{g}_m(Y, X)^T \mathbf{f}(X) + \frac{1}{2} \mathbf{f}(X)^T \mathbf{h}_m(Y, X) \mathbf{f}(X) \right)$$

with  $\mathbf{h}_m(Y, X) = [h_m(Y, X)]_{k,l}$ ,  $k, l = 1, \dots, d$ , and

$$[h_m(Y, X)]_{k,l} = \frac{\partial^2}{\partial F^k \partial F^l} L(Y, \mathbf{F}) \Big|_{\mathbf{F}=\mathbf{F}_{m-1}(X)}.$$

The sample version of the Newton update is given by

$$\mathbf{f}_m = \operatorname{argmin}_{\mathbf{f} \in \mathcal{S} \oplus \dots \oplus \mathcal{S}} \sum_{i=1}^n \mathbf{g}_{m,i}^T \mathbf{f}(x_i) + \frac{1}{2} \mathbf{f}(x_i)^T \mathbf{h}_{m,i} \mathbf{f}(x_i),$$

where  $\mathbf{g}_{m,i} = \mathbf{g}_m(y_i, x_i)$  and  $\mathbf{h}_{m,i} = \mathbf{h}_m(y_i, x_i)$ . In practice, one often approximates  $\mathbf{h}_{m,i}$  by a diagonal matrix

$$\mathbf{h}_{m,i} \approx \operatorname{diag} \left( \frac{\partial^2}{\partial F^{k^2}} L(y_i, \mathbf{F}) \Big|_{\mathbf{F}=\mathbf{F}_{m-1}(x_i)} \right).$$

In this case, the updates can be determined independently as

$$f_m^k = \operatorname{argmin}_{f^k \in \mathcal{S}} \sum_{i=1}^n h_{m,i}^k \left( -\frac{g_{m,i}^k}{h_{m,i}^k} - f^k(x_i) \right)^2,$$

where  $h_{m,i}^k = h_m^k(y_i, x_i)$ .

## 2.4 Tuning parameters and regularization

It has been empirically observed that damping the update in (3) results in increased predictive accuracy (Friedman, 2001). This means that the update in (3) is replaced with

$$F_m(x) = F_{m-1}(x) + \nu f_m(x), \quad \nu > 0, \quad (16)$$

where  $\nu$  is a shrinkage parameter or learning rate. The parameter  $\nu$  can be thought of as a regularization parameter. Under additional assumptions, one can show for linear base learners that when the parameter  $\nu$  goes to zero, the obtained solutions correspond to the set of Lasso solutions (Efron et al., 2004; Zhao and Yu, 2007).

The main tuning parameters of boosting algorithms are thus the number of boosting iterations  $M$  and the shrinkage parameter  $\nu$ . These tuning parameters and also the ones for the base learners presented in the following can be chosen by minimizing a performance measure on a validation dataset, using cross-validation, or using an appropriate model selection criterion.

### 2.4.1 The minimum equivalent sample size per leaf parameter

Depending on the choice of base learners, there are additional tuning parameters. For instance, if trees are used as base learners, the depth of the trees  $L$  and the minimum number of samples per leaf are tuning parameters. Since Newton boosting solves the weighted least squares problem in (13) in each update step, the raw number of samples per leaf is not meaningful, and we argue that instead, one should consider what we denote as the *equivalent sample size per leaf* per leaf. As we show below on real-world and simulated data, this parameter can be important for predictive accuracy.

Specifically, we first normalize the weights

$$\tilde{w}_{m,i} = n \cdot \frac{h_{m,i}}{\sum_{j=1}^n h_{m,j}},$$

such that the sum of all normalized weights  $\tilde{w}_{m,i}$  equals the number of data points  $n$ . We then denote the sum of all normalized weights  $\sum_{i \in L_j} \tilde{w}_{m,i}$  per leaf  $L_j$  as the equivalent sample size per leaf, or equivalent number of weighted data points, and require that this is larger than a certain constant  $S$ :

$$\sum_{i \in L_j} \tilde{w}_{m,i} \geq S. \quad (17)$$

The constant  $S$  is considered as a tuning parameter analogous to the minimum sample size per leaf in gradient boosting.

To the best of our knowledge, other software implementations that use Newton boosting such as **XGBoost** (Chen and Guestrin, 2016) and **LightGBM** (Ke et al., 2017) handle this tuning parameter differently by requiring that the sum of all raw weights  $h_{m,i}$  per leaf is larger than a certain constant which is by default one.<sup>2</sup> According to the authors of **XGBoost**, the motivation for this is that for linear regression, “this simply corresponds to minimum number of instances needed to be in each node”.<sup>3</sup> We argue that this is not a good choice for the following reasons.

First, the second derivative  $h_{m,i}$  of the loss function of a linear regression model with Gaussian noise  $L(Y, F) = \frac{(Y-F)^2}{2\sigma^2}$  equals one only if the noise variance  $\sigma^2$  equals one  $\sigma^2 = 1$ . Otherwise, the second derivative  $h_{m,i}$  equals  $\sigma^{-2}$ . This means that the analogy to the linear regression case does not hold true in general. In contrast, our proposed normalized weights  $\tilde{w}_{m,i}$  do indeed equal one for the linear regression case no matter what the noise variance is, and thus the sum of normalized weights  $\sum_{i \in L_j} \tilde{w}_{m,i}$  equals the number of samples per leaf for the linear regression model also when  $\sigma^2 \neq 1$ . In general, the sum of normalized weights  $\tilde{w}_{m,i}$  corresponds to the number of weighted samples, and one has thus good intuition concerning reasonable candidate values or ranges for this. If the raw weights are not normalized, this is not the case. I.e., the sum of raw weights cannot be interpreted as the number of weighted samples, and its interpretation changes depending on the loss function used. Consequently, the minimum sum of raw weights  $\sum_{i \in L_j} h_{m,i}$  is a parameter that is difficult to tune in practice and we obtain inferior predictive accuracy for the large majority of datasets in our experiments in Section 3. Further, in Section 4.2 we provide empirical evidence that the minimum number of weighted samples per leaf is an important tuning parameter and that the unnormalized version of this tuning parameter is difficult to tune.

In addition to the above-presented tuning parameters, one can consider further tuning parameters such as  $L^1$  and/or  $L^2$  regularization penalties on the tree weights, or an  $L^0$  penalty on the number of leaves. Finally, boosting algorithms can also be made stochastic

---

<sup>2</sup>This constant is denoted by `min_child_weight` in **XGBoost** (as of September 10, 2020).

<sup>3</sup>Unfortunately, this is not documented in the corresponding companion article (Chen and Guestrin, 2016) We gather this information from the online documentation <https://xgboost.readthedocs.io/en/latest/parameter.html> (retrieved on September 10, 2020).

(Friedman, 2002) by (sub-)sampling data points in each boosting iteration and variables in the tree algorithm as it is done for random forests.

## 2.5 Numerical stability and computational cost

Friedman et al. (2000) observed for the LogitBoost algorithm, i.e. Newton boosting for a Bernoulli likelihood with a logistic link function, that numerical stability can be an issue for Newton boosting. Similarly as in Friedman et al. (2000), we enforce a lower bound on the second derivatives  $h_{m,i}$  at  $10^{-20}$  such that they are always strictly positive in our implementation of Newton boosting.<sup>4</sup>

Concerning computational cost, the main cost of a boosting algorithm with trees as base learners results from growing the regression trees (Ke et al., 2017). Consequently, the differences in computational times are marginal for the three versions of boosting presented in this article. Tree boosting implementations that are designed to scale to large data use computational efficient algorithms for growing trees; see, e.g., Chen and Guestrin (2016).

## 2.6 Software implementations

The methodology presented in this article, i.e., gradient, Newton, and hybrid gradient-Newton boosting is implemented in the Python package **KTBoost**, which is openly available from the Python Package Index (PyPI) repository.<sup>5</sup>

We briefly summarize which types of boosting algorithms are used by existing software implementations. The R package **gbm** (Ridgeway, 2007) and the Python library **scikit-learn** (Pedregosa et al., 2011) follow the approach of Friedman (2001) and use gradient descent steps for finding the structures of trees with Newton updates for the tree leaves (if applicable, see Section 2.1.5). **XGBoost** (Chen and Guestrin, 2016) uses Newton boosting with Newton steps for finding both the tree structure and the tree leaves. The R package **mboost** (Hothorn et al., 2010) uses gradient boosting. In addition to trees, it also supports other base learners which include linear functions, one- and two-dimensional smoothing splines, spatial terms, as well as user-defined ones. Other recent implementa-

---

<sup>4</sup>We have not done a comprehensive study on the impact of this lower bound. However, when we choose the bounds at  $10^{-16}$  and  $10^{-30}$ , we have not observed any noticeable differences in the outcomes (results not tabulated).

<sup>5</sup>The parameter `update_step` of the functions `BoostingClassifier` and `BoostingRegressor` takes as arguments `gradient`, `hybrid`, or `newton`. See <https://github.com/fabsig/KTBoost> for more information.

tions such as **LightGBM** (Ke et al., 2017), **TF Boosted Trees** (Ponomareva et al., 2017) and **Spark MLlib** (Meng et al., 2016), do not explicitly mention in their companion articles (Ke et al., 2017; Ponomareva et al., 2017) or in their online documentation<sup>6</sup> whether gradient descent or Newton updates are used in the stagewise boosting updates. We infer from the corresponding source code that **LightGBM** uses Newton boosting. To the best of our knowledge, none of the existing solutions allows the user to explicitly choose between a gradient or a Newton step for calculating the boosting updates.

### 3 Empirical evaluation and comparison

In the following, we compare the three different boosting algorithms presented in the previous section for different loss functions on various datasets using regression trees as base learners.<sup>7</sup> Specifically, we use the CART version of Breiman et al. (1984) with the mean squared error as splitting criterion. Note that we use trees (Breiman et al., 1984) as base learners as these are the most widely adopted base learners in applied data science and machine learning research (Ridgeway, 2007; Pedregosa et al., 2011; Chen and Guestrin, 2016; Meng et al., 2016; Ke et al., 2017; Ponomareva et al., 2017). Besides Newton boosting with the novel equivalent sample size per leaf parameter, we also consider Newton boosting as implemented in **XGBoost** for which the sum of Hessians in each leaf acts as tuning parameter.<sup>8</sup>

#### 3.1 Real-world data

We consider the following datasets: adult, bank, (breast) cancer, ijcnn, ionosphere, titanic, sonar, car, covtype, digits, glass, letter, satimage, smartphone, usps, insurance, birthweight, and (childhood) malnutrition. Poisson regression is used for the insurance dataset. For the birthweight and malnutrition datasets, we use mean-scale regression assuming a normal likelihood where both the mean and the log-transformed scale parameter, i.e. the log-transformed standard deviation, are modeled as functions of the predictor variables; see

---

<sup>6</sup><https://spark.apache.org/docs/latest/mllib-ensembles.html#gradient-boosted-trees-gbts> (retrieved on September 10, 2020).

<sup>7</sup>The code to reproduce the results can be found on <https://github.com/fabsig/GradientNewtonBoosting>

<sup>8</sup>We use **XGBoost** version number 0.7 in Python with the options `tree_method='exact'`, `updater='grow_colmaker'`, `lambda=0`, and all other parameters at the default values unless otherwise mentioned.



Appendix A for more details. Note that the mean-scale regression model is an example of a GAMLSS model (Rigby and Stasinopoulos, 2005; Mayr et al., 2012). For the remaining datasets, binary or multiclass classification is used. The insurance dataset is obtained from Kaggle<sup>9</sup>. The birthweight (Schild et al., 2008) and malnutrition (Fenske et al., 2011) datasets are obtained from the `tbm` R package<sup>10</sup>. The covtype, ijcn, and usps datasets are LIBSVM datasets<sup>11</sup>. All other datasets are obtained from the UCI Machine Learning Repository<sup>12</sup>. A summary of the datasets can be found in Table 1. If a dataset contains categorical predictor variables, these are converted to binary dummy variables using one-hot encoding.

| Data         | Type / nb. classes | Nb. samples | Nb. features |
|--------------|--------------------|-------------|--------------|
| adult        | 2                  | 48842       | 108          |
| bank         | 2                  | 41188       | 62           |
| cancer       | 2                  | 699         | 9            |
| ijcn         | 2                  | 141691      | 22           |
| ionosphere   | 2                  | 351         | 34           |
| sonar        | 2                  | 208         | 60           |
| car          | 4                  | 1728        | 21           |
| covtype      | 7                  | 581012      | 54           |
| digits       | 10                 | 5620        | 64           |
| glass        | 7                  | 214         | 9            |
| letter       | 26                 | 20000       | 16           |
| satimage     | 6                  | 6438        | 36           |
| smartphone   | 6                  | 10299       | 561          |
| usps         | 10                 | 9298        | 256          |
| insurance    | Poisson regr.      | 50999       | 117          |
| birthweight  | Mean-scale regr.   | 150         | 5            |
| malnutrition | Mean-scale regr.   | 24166       | 42           |

Table 1: Summary of datasets.

<sup>9</sup><https://www.kaggle.com/apex51/poisson-regression>

<sup>10</sup>Available on <https://r-forge.r-project.org>

<sup>11</sup><https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

<sup>12</sup><http://archive.ics.uci.edu/ml/datasets/>

We randomly split the data into three equally sized datasets: training, validation, and test data. Learning is done on the training data, tuning parameters are chosen on the validation data, and model comparison is done on the test data. For the two largest datasets (ijcnn and covtype) we limit the size of the training, validation, and test data to 20000 data points. This is done for computational reasons. We note that there are various strategies so that tree-based boosting scales to large data (Chen and Guestrin, 2016; Ke et al., 2017), but this is not the scope of this article. To quantify variability in the results, we use several different random splits of the data. The number of sample splits is 100 for datasets with less than 1500 samples (less than 500 training samples), 20 for datasets with a size between 1500 and 7500 (number of training samples between 500 and 2500), and 10 for datasets with more than 7500 samples (more than 2500 training samples).

Concerning tuning parameters, we select the number of boosting iterations  $M$  from  $\{1, 2, \dots, 1000\}$ , the learning rate  $\nu$  from  $\{1, 10^{-1}, 10^{-2}, 10^{-3}\}$ , and the minimum number of samples per leaf from  $\{1, 5, 25, 100\}$ . For Newton boosting, the latter is replaced by the equivalent sample size per leaf in Equation (17), and for the XGBoost implementation, the minimum sum of Hessians per leaf (`min_child_weight`) is used. Further, for the mean-scale regression datasets, the minimum number of samples per leaf is chosen from  $\{25, 100\}$  only for gradient and hybrid boosting since a very small number of samples can lead to identifiability problems when modeling both the mean and the scale. Tuning parameters are chosen for each sample split such that they minimize the error rate for classification and the negative log-likelihood for regression on the validation data. The maximal tree depth is set to five for all methods. We are not considering the maximal tree depth as an additional tuning parameter for computational reasons. However, additional results for a subset of the datasets reported in Section 4.3.1 and Appendix C show that similar findings are obtained for other tree depths. Further, we note that in Section 4.2 and Appendix B, we also consider the case when the minimum number of samples per leaf parameter is not chosen by minimizing the test error on the validation data but is instead set to the default value of one.

In Figure 1 and Table 2, we report test error rates for classification and test negative log-likelihoods for regression datasets. Figure 1 visualizes the results using boxplots. In Table 2, we additionally report average test errors and test negative log-likelihoods as well as approximate standard deviations. Further, we report the average rank of every

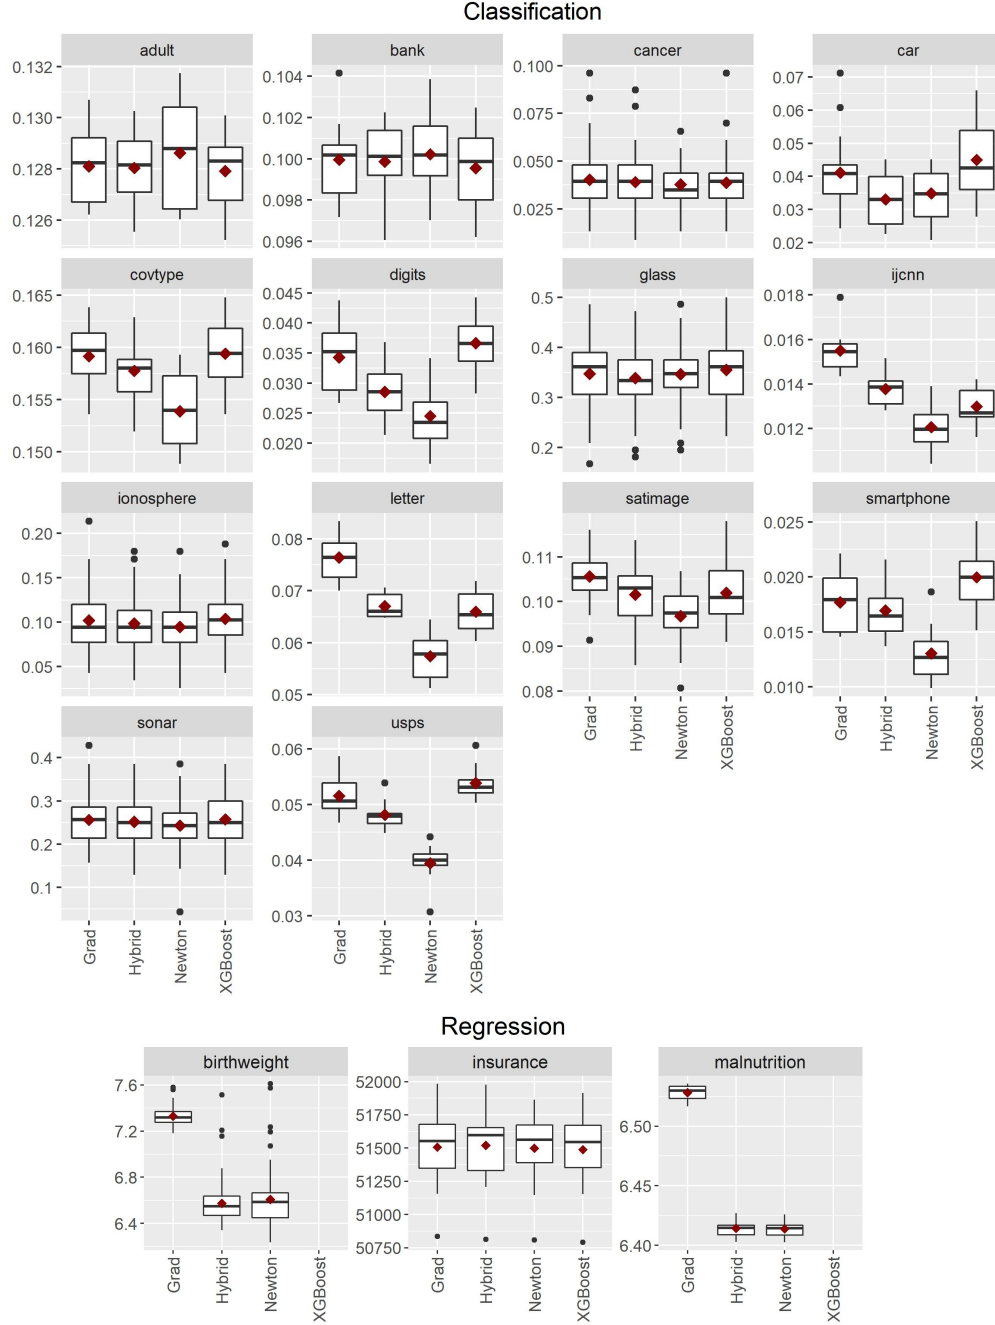


Figure 1: Comparison of boosting methods using test error rate for classification and test negative log-likelihood for regression. The red rhombi represent means.

method over the different datasets. Since **XGBoost** does not support mean-scale regression, we only consider the datasets for which all four methods can be run when calculating average ranks. Overall, we find that Newton boosting with the novel equivalent sample

| Data                     | Grad             | Hybrid                  | Newton                  | XGBoost                 |
|--------------------------|------------------|-------------------------|-------------------------|-------------------------|
| adult                    | 0.128 (0.00158)  | 0.128 (0.00143)         | 0.129 (0.00222)         | <b>0.128</b> (0.00158)  |
| bank                     | 0.1 (0.00213)    | 0.0999 (0.00191)        | 0.1 (0.00208)           | <b>0.0996</b> (0.00198) |
| cancer                   | 0.0402 (0.0147)  | 0.039 (0.0132)          | <b>0.0378</b> (0.0108)  | 0.0387 (0.0124)         |
| ijcnn                    | 0.0155 (0.00101) | 0.0138 (0.000791)       | <b>0.0121</b> (0.00109) | 0.013 (0.000838)        |
| ionosphere               | 0.102 (0.0317)   | 0.0984 (0.0293)         | <b>0.0945</b> (0.0277)  | 0.104 (0.0288)          |
| sonar                    | 0.256 (0.0556)   | 0.252 (0.0517)          | <b>0.243</b> (0.0499)   | 0.257 (0.055)           |
| car                      | 0.0411 (0.0114)  | <b>0.0331</b> (0.00759) | 0.0349 (0.00764)        | 0.045 (0.0112)          |
| covtype                  | 0.159 (0.00331)  | 0.158 (0.00326)         | <b>0.154</b> (0.00384)  | 0.159 (0.00371)         |
| digits                   | 0.0343 (0.00547) | 0.0285 (0.00464)        | <b>0.0245</b> (0.00467) | 0.0367 (0.0044)         |
| glass                    | 0.347 (0.0618)   | <b>0.339</b> (0.0612)   | 0.346 (0.0595)          | 0.355 (0.0607)          |
| letter                   | 0.0764 (0.00449) | 0.067 (0.00233)         | <b>0.0574</b> (0.00438) | 0.066 (0.00409)         |
| satimage                 | 0.106 (0.00622)  | 0.102 (0.00801)         | <b>0.0968</b> (0.00673) | 0.102 (0.00687)         |
| smartphone               | 0.0177 (0.00272) | 0.017 (0.00266)         | <b>0.013</b> (0.00267)  | 0.02 (0.00288)          |
| usps                     | 0.0516 (0.00371) | 0.0482 (0.00261)        | <b>0.0395</b> (0.00364) | 0.0539 (0.00311)        |
| insurance                | 51500 (341)      | 51500 (349)             | 51500 (324)             | <b>51500</b> (332)      |
| malnutrition             | 6.53 (0.00649)   | 6.41 (0.00792)          | <b>6.41</b> (0.00733)   |                         |
| birthweight              | 7.33 (0.0702)    | <b>6.57</b> (0.177)     | 6.61 (0.226)            |                         |
| Av. rank                 | 3.27             | 2.2                     | 1.6                     | 2.93                    |
| p-val Friedman test      | 0.000746         |                         |                         |                         |
| Adj. p-val Wilcoxon test | 0.000229         | 0.0714                  |                         | 0.0302                  |

Table 2: Results for real-world data: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

size per leaf parameter has clearly the lowest generalization error among the four methods. Its average rank is 1.6. The second best method with an average rank of 2.2 is hybrid gradient-Newton boosting. Gradient boosting often has the lowest predictive accuracy with an average rank of 3.27. In addition, Newton boosting with the novel number of weighted samples parameter performs substantially better than the XGBoost variant of Newton boosting with a minimum sum of unnormalized Hessians parameter which has an

average rank of 2.93. We observe particularly striking differences with large outperformance in the predictive accuracy of Newton boosting for several classification datasets (ijcnn, digits, letter, satimage, smartphone, and usps). For the two mean-scale regression datasets (birthweight and malnutrition), we also observe that gradient boosting performs worse than Newton and hybrid gradient-Newton boosting, but no notable difference among the latter two is found. For the Poisson regression dataset (insurance), gradient, hybrid, and Newton boosting perform equally well.

Concerning statistical significance, we note that when using a resampling approach, standard statistical tests, such as a paired t-test, cannot be used to do a pairwise comparison of the different algorithms separately per dataset since training and test datasets in different splits are dependent due to overlap (Dietterich, 1998; Bengio and Grandvalet, 2004; Demšar, 2006), and this can result in biased standard error estimates for the generalization error. Following Demšar (2006), we compare the different methods across all datasets using a Friedman test with an Iman and Davenport correction (Iman and Davenport, 1980). This gives a p-value of 0.000746 which shows that the differences in the four methods are highly significant. We next use a Wilcoxon signed-rank test to investigate whether the pairwise differences in accuracy between Newton boosting with the novel number of weighted samples parameter and the three alternative methods are statistically significant. To account for the fact that we do multiple tests, we apply a Holm-Bonferroni correction (Holm, 1979). Comparing Newton boosting with gradient and hybrid gradient-Newton boosting, we obtain adjusted p-values of 0.000229 and 0.074. I.e., Newton boosting performs significantly better than gradient boosting and the difference between Newton boosting and hybrid gradient-Newton boosting is marginally not significant at a 5% level. However, the sample size for performing these tests is relatively small (17) and, consequently, the tests likely have low power.

### 3.2 Simulated data

In the following, we compare the performance of the different boosting approaches on simulated data for both classification and regression. Concerning regression, we consider two extensions of generalized linear models, boosted Poisson and Gamma regression, as well as the mean-scale regression model used in Section 3. For classification, we consider both binary and multiclass classification. In addition, we consider the boosted Tobit model

(Sigrist and Hirnschall, 2019), which can be interpreted as a hybrid regression-classification model. See Section A in the appendix for more details on these models.

For classification, we use the `scikit-learn` function `make_classification`, which simulates from an algorithm that is adapted from Guyon (2003) and was designed to generate the ‘Madelon’ dataset. We use this for both simulating binary data and a multiclass data with five classes. Further, we assume ten (informative) features and no redundant and repeated features; see Guyon (2003) for more details. These two datasets are denoted by ‘bin\_classif’ and ‘multi\_classif’ in the following. In addition, we simulate binary data according to the following specification introduced in Friedman et al. (2000):

$$F(X) = 10 \sum_{j=1}^6 X_j \left( 1 + \sum_{l=1}^6 (-1)^l X_l \right), \quad X \sim N(0, I_{10}),$$

$$Y|X \sim \text{Bernoulli}(p), \quad p = \left( 1 + e^{-F(X)} \right)^{-1}.$$

This data is denoted by ‘bin\_classif\_fht’ in the following. Finally, we also simulate multiclass data with five classes according to the following specification (Friedman et al., 2000):

$$R^2 = \sum_{j=1}^{10} X_j^2, \quad X \sim N(0, I_{10}),$$

$$Y = k \quad \text{if} \quad t_k \leq R^2 < t_{k+1},$$

where the thresholds  $t_k$  are chosen such that the labels are approximately equally distributed among the different classes. We denote this data by ‘multi\_classif\_fht’.

For Poisson and Gamma regression, the boosted Tobit model, as well as mean-scale regression (‘msr’), we consider two non-linear functions. First, we use a function of Friedman (1991) given by:

$$F(X) = 5 \cdot \tan^{-1} \left( \frac{X_2 X_3 - 1 - \frac{1}{X_2 X_4}}{X_1} \right) + 0.2 \quad (\text{‘f3’}),$$

where  $X = (X_1, X_2, X_3, X_4)'$  with  $X_1 \sim \text{Unif}(0, 100)$ ,  $X_2 \sim \text{Unif}(40\pi, 560\pi)$ ,  $X_3 \sim \text{Unif}(0, 1)$ , and  $X_4 \sim \text{Unif}(1, 11)$ . In contrast to the original function of Friedman (1991), we multiply the function by 5 and add 0.2 such that function also attains larger values and that all values are positive. We use the `scikit-learn` function `make_friedman3` for simulation and denote datasets generated by this function using the suffix ‘\_f3’ in the following. Further, we also consider a function introduced in Ridgeway (1999):

$$F(X) = \exp \left( 2 \sin(3X_1 + 5X_1^2) - 2 \sin(3(X_2 + 0.1) + 5(X_2 + 0.1)^2) \right) \quad (\text{‘r’}),$$

where  $X = (X_1, X_2)'$ ,  $X_j \sim \text{Unif}(0, 1)$ , independent. Datasets generated using this function are denoted by the suffix ‘\_r’.<sup>13</sup> For Poisson and Gamma regressions, the above functions are used to model the mean, and for Tobit regression the functions model the mean of the latent variable. For the mean-scale regression model, we simulate  $2n$  variables and relate both the mean and the logarithmic standard deviation to half of the variables. Both Tobit regression and a regression model where both the mean and the scale depend on predictor variables are not supported in **XGBoost** and, consequently, no comparison can be done for these. For Gamma regression, we set the shape parameter to  $\gamma = 10$  and consider this as a known parameter.<sup>14</sup> For the Tobit model, we use  $\sigma = 1$  and also consider this as a known parameter. Further, we set the lower and upper censoring thresholds  $y_l$  and  $y_u$  in such a way that approximately one-third of all data points are lower and upper censored.

We simulate 10 times datasets with 15000 samples. In each run, 5000 samples are used as training, validation, and test data. As in Section 3, we calculate the p-value of a Friedman test with an Iman and Davenport correction (Iman and Davenport, 1980) to check whether there are significant differences among the methods across all datasets. Further, we calculate Holm-Bonferroni corrected (Holm, 1979) p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the three other approaches.

The results are reported in Figure 2 and in Table 3. See Section 3 for more details on the plot and table. We find again that Newton boosting has the highest predictive accuracy for the large majority of datasets, followed by hybrid gradient-Newton, with gradient boosting having the lowest predictive accuracy. The Friedman test with an Iman and Davenport correction shows that there are statistically significant differences among the different boosting approaches. Further, Newton boosting performs significantly better in terms of predictive accuracy than both gradient and hybrid gradient-Newton boosting despite the relatively small sample size and the multiple testing correction. Finally, Newton boosting with the new equivalent sample size per leaf tuning parameter has higher predictive accuracy than the **XGBoost** implementation with the unnormalized number of weighted samples per leaf parameter.

---

<sup>13</sup>For instance, the dataset ‘msr\_r’ is simulated from a mean-scale regression model with both the mean and the standard deviation given by the above function introduced by Ridgeway (1999).

<sup>14</sup>We note that **XGBoost** only supports Gamma regression for  $\gamma = 1$ . However, this slight miss-specification seems to have no detrimental impact as our results below show.

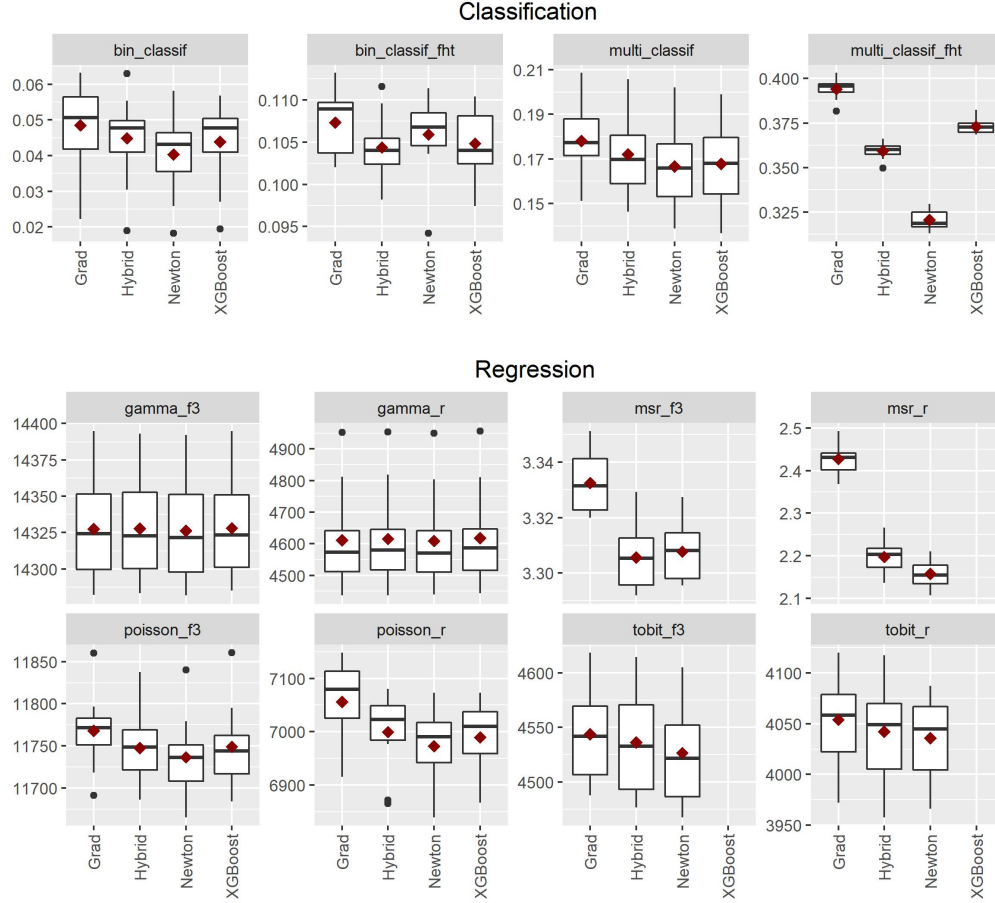


Figure 2: Comparison of boosting methods on simulated datasets using test error rate for classification and test negative log-likelihood for regression. The red rhombi represent means.

## 4 Discussion

### 4.1 Does Newton boosting show higher predictive accuracy than gradient boosting due to faster convergence?

In the previous sections, we have empirically shown that Newton boosting often results in higher predictive accuracy than gradient and also hybrid gradient-Newton boosting. A potential explanation for the observed phenomenon is that Newton boosting converges faster than both gradient and hybrid gradient-Newton boosting, and that hybrid boosting also converges faster than gradient boosting. This, in turn, could allow for using a smaller shrinkage parameter  $\nu$ , and smaller shrinkage parameters usually lead to increased predic-



| Data                     | Grad            | Hybrid               | Newton                 | XGBoost         |
|--------------------------|-----------------|----------------------|------------------------|-----------------|
| bin_classif              | 0.0485 (0.0126) | 0.0449 (0.0126)      | <b>0.0403</b> (0.0116) | 0.0438 (0.0121) |
| bin_classif_fht          | 0.107 (0.00394) | <b>0.104</b> (0.004) | 0.106 (0.00488)        | 0.105 (0.00428) |
| multi_classif            | 0.178 (0.0163)  | 0.172 (0.0176)       | <b>0.167</b> (0.0184)  | 0.168 (0.0185)  |
| multi_classif_fht        | 0.394 (0.00586) | 0.359 (0.00462)      | <b>0.321</b> (0.00597) | 0.373 (0.00415) |
| poisson_r                | 7060 (78.2)     | 7000 (75.4)          | <b>6970</b> (73.7)     | 6990 (68.1)     |
| poisson_f3               | 11800 (45.5)    | 11700 (45.1)         | <b>11700</b> (50)      | 11700 (51.2)    |
| gamma_r                  | 4610 (159)      | 4620 (160)           | <b>4610</b> (158)      | 4620 (157)      |
| gamma_f3                 | 14300 (36.6)    | 14300 (35.5)         | <b>14300</b> (35.7)    | 14300 (35.1)    |
| tobit_r                  | 4050 (45.7)     | 4040 (48.7)          | <b>4040</b> (40.5)     |                 |
| tobit_f3                 | 4540 (45.7)     | 4540 (47.5)          | <b>4530</b> (47.7)     |                 |
| msr_f3                   | 3.33 (0.0107)   | <b>3.31</b> (0.0121) | 3.31 (0.0107)          |                 |
| msr_r                    | 2.43 (0.0408)   | 2.2 (0.0361)         | <b>2.16</b> (0.0327)   |                 |
| Av. rank                 | 3.5             | 2.5                  | 1.25                   | 2.75            |
| p-val Friedman test      | 0.000112        |                      |                        |                 |
| Adj. p-val Wilcoxon test | 0.00146         | 0.00488              |                        | 0.0156          |

Table 3: Results for simulated data: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

tive accuracy. To investigate whether this is the main reason for the differences, we show in Figures 3 and 4 test error rates (classification) and test negative log-likelihoods (regression) as well as training losses versus iteration numbers for several datasets for which we have observed large differences. In order that the results for the different sample splits and also boosting methods are comparable, learning rates are fixed and not tuned. Specifically, we consider the following datasets and learning rates: *ijcnn* ( $\nu = 0.5$ ), *bin\_classif* ( $\nu = 0.5$ ), *digits* ( $\nu = 0.5$ ), *letter* ( $\nu = 0.1$ ), *satimage* ( $\nu = 0.3$ ), *smartphone* ( $\nu = 0.5$ ), *poisson\_r* ( $\nu = 0.03$ ), *malnutrition* ( $\nu = 0.03$ ), and *msr\_r* ( $\nu = 0.05$ ). Note that this list of datasets includes both binary and multiclass classification as well as Poisson and mean-scale regression tasks. The solid lines in Figures 4 and 3 represent means over ten different data splits

into equally sized training, validation, and test data. The lower and upper values of the shaded areas are obtained after point-wise discarding the lowest and largest values. Training losses are shown on a logarithmic scale with a lower cap at  $10^{-5}$  for better visualization.

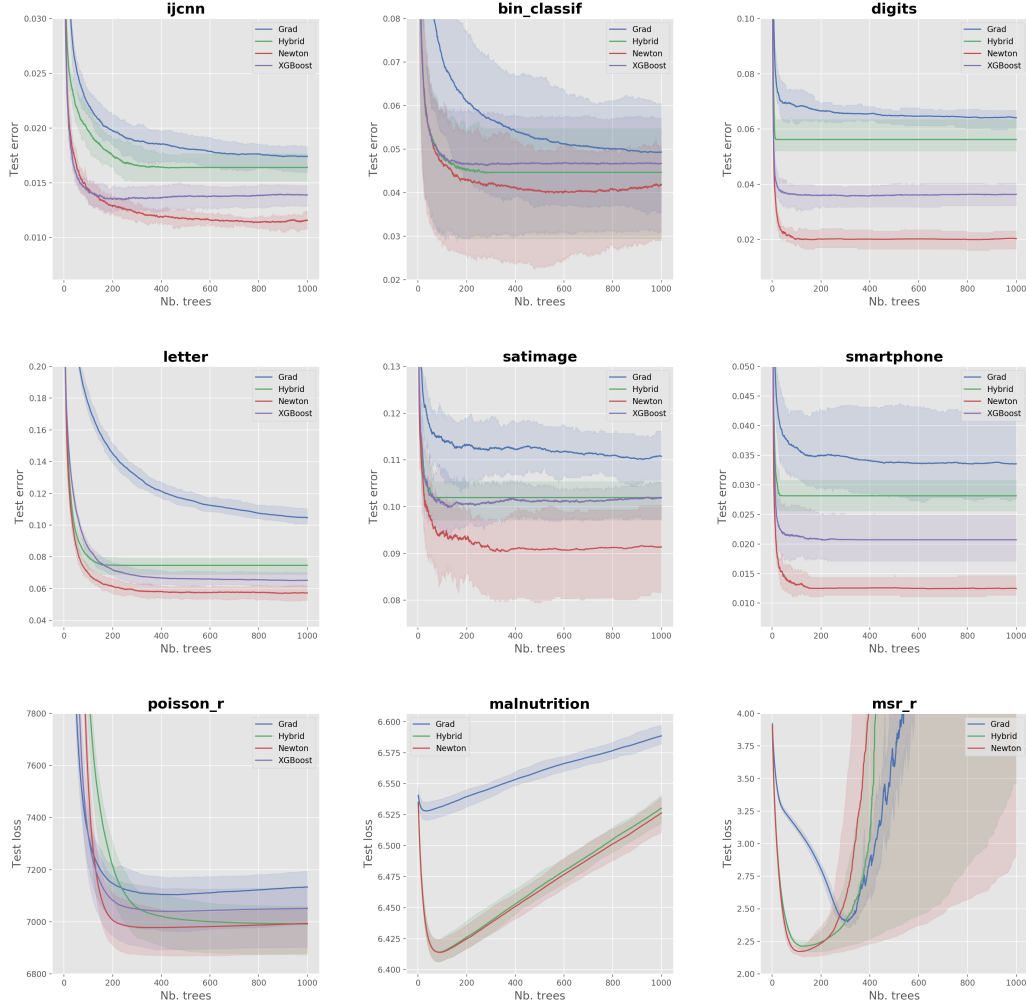


Figure 3: Test error rate (classification) and test negative log-likelihood (regression) versus boosting iteration number.

As expected, Figures 3 and 4 show that gradient boosting often converges slower than Newton and also hybrid gradient-Newton boosting. Concerning the latter two, we visually observe almost no difference in convergence speed. However, the plots also show that the faster convergence speed of Newton boosting is likely not the reason for the observed differences in predictive accuracy since the minima of the test errors for Newton boosting are usually achieved after fewer iterations compared to gradient boosting. Interestingly,

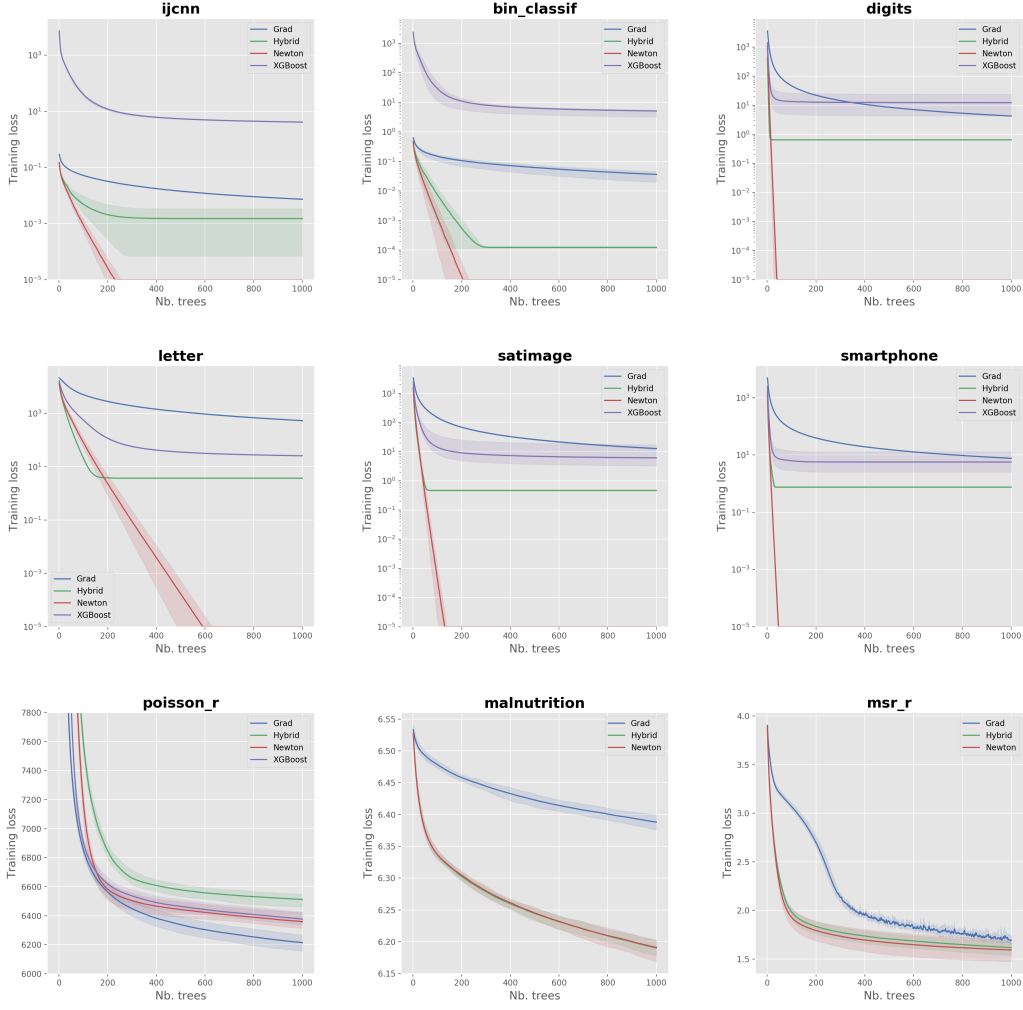


Figure 4: Training loss versus boosting iteration number. The (logarithmic) y-scale is truncated at  $10^{-5}$  for better visualization.

our results indicate that Newton boosting converges to lower training losses, which are essentially zero for the majority of classification datasets, while at the same time having higher test accuracy. The fact that interpolating classifiers with zero training loss generalize well to novel data seems to be in contradiction to the well-known bias-variance trade-off. However, similar results have recently been observed for other datasets and complex models such as deep learning and kernel machines (Zhang et al., 2017; Belkin et al., 2018b,a). For the Poisson regression dataset, we observe that gradient boosting results in a lower training loss but a higher test loss. For the mean-square regression datasets, we find that Newton and hybrid boosting show lower test and training losses than gradient boosting,

and, in contrast to the classification datasets, we find clear signs of overfitting as the test losses start to increase again after a certain number of boosting iterations. Concerning the **XGBoost** implementation, we find that **XGBoost** results in higher training losses and also higher test errors compared to our Newton boosting implementation. In particular, the training losses do not converge to zero for the classification datasets.

Finally, we recall that in cases where the Hessians  $h_{m,i}$  defined in Equation (14) are constant, there is no difference between gradient and Newton boosting. It is thus likely that the more variation there is in the second-order terms  $h_{m,i}$  the larger is the difference between gradient and Newton boosting.

## 4.2 Importance of the minimum number of (weighted) samples per leaf parameter

In Table 4 in Appendix B, we additionally report results for the real-world datasets when the minimum number of (weighted) samples per leaf parameter is not tuned by minimizing a validation loss and simply set to a default value. I.e., for gradient and hybrid gradient-Newton boosting, the minimum number of samples per leaf is one, for Newton boosting with our proposed choice in (17), we set the minimum equivalent sample size per leaf parameter to one, and for the **XGBoost** implementation, we set the minimum sum of Hessians to its default value, i.e., also one.<sup>15</sup> Overall, we find that the difference in predictive accuracy between Newton boosting and gradient as well as hybrid gradient-Newton boosting is more pronounced when the minimum number of samples parameter is not tuned. Interestingly, we find that Newton boosting with the unnormalized sum of Hessians per leaf parameter as implemented in **XGBoost** and Newton boosting with the novel equivalent number of weighted samples per leaf parameters as implemented in **KTBoost** perform now almost equally well, and that the results of Newton boosting with the novel equivalent sample size per leaf parameter are worse compared to the ones in Table 2 when also tuning this parameter. This provides evidence that the number of samples per leaf is an important tuning parameter, and that the unnormalized version of this tuning parameter is difficult to tune.

---

<sup>15</sup>We exclude the mean-scale regression datasets as there is no obvious default value since the minimum number of samples per leaf needs to be larger than one.

### 4.3 Additional results and robustness check

In the following, we report additional results to show that our findings are robust to the choice of the maximal tree depth tuning parameters and also to the sample size.

#### 4.3.1 Maximal tree depth

We additionally consider the following maximal tree depths: 1 (stumps), 3, 8, and 20. The results in Appendix C show that we continue to observe very similar differences among the different boosting versions also for other maximal tree depths. In particular, Newton boosting results in the highest predictive accuracy. For the majority of the datasets, stumps and also trees with maximal depth 3 perform worse than larger trees. However, for one simulated dataset (`multi_classif_fht`), stumps result in clearly increased predictive accuracy. Further, very large trees with a depth of 20 generally perform slightly worse than trees with a depth of 5 or 8.

#### 4.3.2 Simulated data with a smaller sample size

We repeat the simulation study using a smaller sample size. Specifically, we use  $n = 500$  samples and do 100 simulation runs with the same specifications as in Section 3.2. The results for this are reported in Table 9 in Appendix D. Qualitatively, we find similar results as for the larger sample size. When excluding the Tobit and mean-scale regression datasets for which `XGBoost` cannot be run, Newton boosting has an average rank of 1.25, hybrid gradient-Newton boosting has an average rank of 2.38, gradient boosting has an average rank of 3.57, and `XGBoost` has an average rank of 2.88. Further, gradient boosting has significantly lower predictive accuracy than Newton boosting. The differences between hybrid gradient-Newton boosting and Newton boosting are less pronounced and not significant, though. Further, the differences between Newton boosting with the novel equivalent sample size per leaf parameter and the `XGBoost` implementation of Newton boosting with the unnormalized number of weights parameter are also not significant. We note, however, that these statistical tests are done using very small sample sizes of only 12 or 8, respectively, and thus have low power.

## 5 Conclusions

We compare gradient and Newton boosting as well as a hybrid variant of the two with trees as base learners on a wide range of classification and regression datasets. Our empirical results show that Newton boosting outperforms gradient and often also hybrid gradient-Newton boosting in terms of predictive accuracy. Further, we present empirical evidence that this outperformance is not a consequence of a faster convergence speed of Newton boosting. Interestingly, Newton boosting converges to lower values of the empirical risk while at the same time having lower test errors. In addition, we introduce a novel tuning parameter denoted as *equivalent sample size per leaf* parameter which is interpretable, intuitive to tune, and important for predictive accuracy.

We do not have a full explanation for why Newton boosting shows lower generalization errors, and future research should shed light on the reasons for this. Theoretical results such as bounds on the generalization error could help to find an explanation. Further, future research should investigate whether similar results are found for other types of base learners such as splines (Bühlmann and Yu, 2003; Hothorn et al., 2010), when combining different base learners such as trees and kernel machines (Sigrist, 2019), when using a linear model as baseline and adding more flexible base learners in a boosting framework (Costa et al., 2019), or when combining boosting with Gaussian process or mixed effects models (Sigrist, 2020).

## Acknowledgments

This research was partially supported by the Swiss Innovation Agency - Innosuisse (25746.1 PFES-ES). We are grateful to Christoph Hirnschall and Torsten Hothorn for valuable suggestions and discussions.

## References

Ahamad, M. M., Aktar, S., Rashed-Al-Mahfuz, M., Uddin, S., Liò, P., Xu, H., Summers, M. A., Quinn, J. M., and Moni, M. A. (2020). A machine learning model to identify early stage symptoms of sars-cov-2 infected patients. *Expert Systems with Applications*, 160:113661.

- Al-Hadeethi, H., Abdulla, S., Diykh, M., Deo, R. C., and Green, J. H. (2020). Adaptive boost ls-svm classification approach for time-series signal classification in epileptic seizure diagnosis applications. *Expert Systems with Applications*, 161:113676.
- Barsacchi, M., Bechini, A., and Marcelloni, F. (2020). An analysis of boosted ensembles of binary fuzzy decision trees. *Expert Systems with Applications*, page 113436.
- Belkin, M., Hsu, D. J., and Mitra, P. (2018a). Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 2306–2317.
- Belkin, M., Ma, S., and Mandal, S. (2018b). To understand deep learning we need to understand kernel learning. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 541–549.
- Bengio, Y. and Grandvalet, Y. (2004). No unbiased estimator of the variance of k-fold cross-validation. *Journal of machine learning research*, 5(Sep):1089–1105.
- Breiman, L. (1998). Arcing classifiers. *Annals of Statistics*, pages 801–824.
- Breiman, L. (1999). Prediction games and arcing algorithms. *Neural computation*, 11(7):1493–1517.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and regression trees*. CRC press.
- Bühlmann, P. et al. (2006). Boosting for high-dimensional linear models. *The Annals of Statistics*, 34(2):559–583.
- Bühlmann, P. and Hothorn, T. (2007). Boosting algorithms: Regularization, prediction and model fitting. *Statistical Science*, pages 477–505.
- Bühlmann, P. and Yu, B. (2003). Boosting with the  $l_2$  loss: regression and classification. *Journal of the American Statistical Association*, 98(462):324–339.

- Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794. ACM.
- Costa, M. A., Wullt, B., Norrlöf, M., and Gunnarsson, S. (2019). Failure detection in robotic arms using statistical modeling, machine learning and hybrid gradient boosting. *Measurement*, 146:425–436.
- De Menezes, F. S., Liska, G. R., Cirillo, M. A., and Vivanco, M. J. (2017). Data classification with binary response through the boosting algorithm and logistic regression. *Expert Systems with Applications*, 69:62–73.
- Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research*, 7(Jan):1–30.
- Dietterich, T. G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation*, 10(7):1895–1923.
- Djeundje, V. B., Crook, J., Calabrese, R., and Hamid, M. (2020). Enhancing credit scoring with alternative data. *Expert Systems with Applications*, page 113766.
- Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *The Annals of statistics*, 32(2):407–499.
- Fenske, N., Kneib, T., and Hothorn, T. (2011). Identifying risk factors for severe childhood malnutrition by boosting additive quantile regression. *Journal of the American Statistical Association*, 106(494):494–510.
- Freund, Y. and Schapire, R. E. (1995). A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer.
- Freund, Y., Schapire, R. E., et al. (1996). Experiments with a new boosting algorithm. In *ICML*, volume 96, pages 148–156. Bari, Italy.
- Friedman, J., Hastie, T., Tibshirani, R., et al. (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.



- Friedman, J. H. (1991). Multivariate adaptive regression splines. *The annals of statistics*, pages 1–67.
- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232.
- Friedman, J. H. (2002). Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378.
- Guyon, I. (2003). Design of experiments of the nips 2003 variable selection benchmark.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70.
- Hothorn, T., Bühlmann, P., Kneib, T., Schmid, M., and Hofner, B. (2010). Model-based boosting 2.0. *Journal of Machine Learning Research*, 11(Aug):2109–2113.
- Iman, R. L. and Davenport, J. M. (1980). Approximations of the critical region of the fbietkan statistic. *Communications in Statistics-Theory and Methods*, 9(6):571–595.
- Johnson, R. and Zhang, T. (2013). Learning nonlinear functions using regularized greedy forest. *IEEE transactions on pattern analysis and machine intelligence*, 36(5):942–954.
- Kadkhodaei, H. R., Moghadam, A. M. E., and Dehghan, M. (2020). Hboost: A heterogeneous ensemble classifier based on the boosting method and entropy measurement. *Expert Systems with Applications*, page 113482.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, pages 3149–3157.
- Li, P. (2010). Robust logitboost and adaptive base class (abc) logitboost. In *UAI*.
- Lu, H., Wang, H., and Yoon, S. W. (2019). A dynamic gradient boosting machine using genetic optimizer for practical breast cancer prognosis. *Expert Systems with Applications*, 116:340–350.
- Mason, L., Baxter, J., Bartlett, P. L., and Frean, M. R. (2000). Boosting algorithms as gradient descent. In *Advances in neural information processing systems*, pages 512–518.

- Mayr, A., Binder, H., Gefeller, O., and Schmid, M. (2014a). The evolution of boosting algorithms. *Methods of information in medicine*, 53(06):419–427.
- Mayr, A., Binder, H., Gefeller, O., and Schmid, M. (2014b). Extending statistical boosting. *Methods of information in medicine*, 53(06):428–435.
- Mayr, A., Fenske, N., Hofner, B., Kneib, T., and Schmid, M. (2012). Generalized additive models for location, scale and shape for high dimensional data—a flexible approach based on boosting. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 61(3):403–427.
- Meng, X., Bradley, J., Yavuz, B., Sparks, E., Venkataraman, S., Liu, D., Freeman, J., Tsai, D., Amde, M., Owen, S., et al. (2016). Mllib: Machine learning in apache spark. *The Journal of Machine Learning Research*, 17(1):1235–1241.
- Moscatelli, M., Parlapiano, F., Narizzano, S., and Viggiano, G. (2020). Corporate default forecasting with machine learning. *Expert Systems with Applications*, page 113567.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Ponomareva, N., Radpour, S., Hendry, G., Haykal, S., Colthurst, T., Mitrichev, P., and Grushetsky, A. (2017). Tf boosted trees: A scalable tensorflow based framework for gradient boosting. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 423–427. Springer.
- Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., and Gulin, A. (2018). Catboost: unbiased boosting with categorical features. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 6639–6649. Curran Associates, Inc.
- Ridgeway, G. (2007). Generalized boosted models: A guide to the gbm package. *Update*, 1(1):2007.
- Ridgeway, G. K. (1999). *Generalization of boosting algorithms and applications of bayesian inference for massive datasets*. PhD thesis, University of Washington.

- Rigby, R. A. and Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(3):507–554.
- Saberian, M. J., Masnadi-Shirazi, H., and Vasconcelos, N. (2011). Taylorboost: First and second-order boosting algorithms with explicit margin control. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2929–2934. IEEE.
- Schapire, R. E. (1990). The strength of weak learnability. *Machine learning*, 5(2):197–227.
- Schapire, R. E. (2003). The boosting approach to machine learning: An overview. In *Nonlinear estimation and classification*, pages 149–171. Springer.
- Schapire, R. E. and Freund, Y. (2012). *Boosting: Foundations and algorithms*. MIT press.
- Schild, R., Maringa, M., Siemer, J., Meurer, B., Hart, N., Goecke, T., Schmid, M., Hothorn, T., and Hansmann, M. (2008). Weight estimation by three-dimensional ultrasound imaging in the small fetus. *Ultrasound in Obstetrics & Gynecology*, 32(2):168–175.
- Schmid, M. and Hothorn, T. (2008). Boosting additive models using component-wise p-splines. *Computational Statistics & Data Analysis*, 53(2):298–311.
- Sigrist, F. (2019). KTBoost: Combined Kernel and Tree Boosting. *arXiv preprint arXiv:1902.03999*.
- Sigrist, F. (2020). Gaussian process boosting. *arXiv preprint arXiv:2004.02653*.
- Sigrist, F. and Hirnschall, C. (2019). Grabit: Gradient tree-boosted tobit models for default prediction. *Journal of Banking & Finance*, 102:177 – 192.
- Soleymani, R., Granger, E., and Fumera, G. (2018). Progressive boosting for class imbalance and its application to face re-identification. *Expert Systems with Applications*, 101:271–291.
- Sun, P., Zhang, T., and Zhou, J. (2014). A convergence rate analysis for logitboost, mart and their variant. In *ICML*, pages 1251–1259.
- Wang, G., Ma, J., and Yang, S. (2014). An improved boosting based on feature selection for corporate bankruptcy prediction. *Expert Systems with Applications*, 41(5):2353–2361.

- Xia, Y., Liu, C., Li, Y., and Liu, N. (2017). A boosted decision tree approach using bayesian hyper-parameter optimization for credit scoring. *Expert Systems with Applications*, 78:225–241.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*.
- Zhao, P. and Yu, B. (2007). Stagewise lasso. *Journal of Machine Learning Research*, 8(Dec):2701–2726.
- Zheng, S. and Liu, W. (2012). Functional gradient ascent for probit regression. *Pattern Recognition*, 45(12):4428–4437.
- Zhou, Y., Mazzuchi, T. A., and Sarkani, S. (2020). M-adaboost-a based ensemble system for network intrusion detection. *Expert Systems with Applications*, page 113864.

# Appendices

## A Loss functions for regression and classification tasks

In the following, we list the loss functions and corresponding gradients and second derivatives that we consider in this article.

- **Binary classification**

$$Y|X \sim \text{Bernoulli}(p), \quad p = (1 + e^{-F(X)})^{-1}$$

$$\text{Loss: } L(Y, F) = -YF + \log(1 + e^F)$$

$$\text{Gradient: } \frac{\partial L(Y, F)}{\partial F} = -Y + p$$

$$\text{Hessian: } \frac{\partial^2 L(Y, F)}{\partial F^2} = p(1 - p)$$

- **Multiclass classification**

$$Y|X \sim \text{Multinom}(p_1, \dots, p_K), \quad p_k = \frac{e^{F_k(X)}}{\sum_{l=1}^K e^{F_l(X)}}, \quad k = 1, \dots, K$$

Loss:

$$L(Y, F) = \sum_{k=1}^K \left( -\mathbb{1}_{\{Y=k\}} F_k + \log \left( \sum_{l=1}^K e^{F_l(X)} \right) \right), \quad F = (F_1, \dots, F_K)$$

$$\text{Gradient: } \frac{\partial L(Y, F)}{\partial F_k} = -\mathbb{1}_{\{Y=k\}} + p_k$$

$$\text{Hessian: } \frac{\partial^2 L(Y, F)}{\partial F_k^2} = p_k(1 - p_k)$$

As in Friedman et al. (2000), we use  $\frac{\partial^2 L(Y, F)}{\partial F_k \partial F_j} = 0$  for simplicity.

- **Poisson regression**

$$Y|X \sim \text{Pois}(\lambda), \quad \lambda = e^{F(X)}$$

$$\text{Loss: } L(Y, F) = -YF + e^F$$

$$\text{Gradient: } \frac{\partial L(Y, F)}{\partial F} = -Y + e^F$$

$$\text{Hessian: } \frac{\partial^2 L(Y, F)}{\partial F^2} = e^F$$

- **Gamma regression**

$$Y|X \sim \text{Gamma}(\gamma, \lambda) \text{ with shape } \gamma \text{ and rate } \lambda, \quad \lambda = \gamma e^{-F(X)}$$

$$\text{Loss: } L(Y, F) = \gamma (F + e^{-F} Y) - (\gamma - 1) \log(Y) - \gamma \log(\gamma) + \log(\Gamma(\gamma))$$

$$\text{Gradient: } \frac{\partial L(Y, F)}{\partial F} = \gamma (1 - e^{-F} Y)$$

$$\text{Hessian: } \frac{\partial^2 L(Y, F)}{\partial F^2} = \gamma e^{-F} Y$$

- **Tobit model**

$$Y|X \sim \text{Tobit}_{\{y_l, y_u\}}(\mu, \sigma^2), \text{ with mean } \mu, \mu = F(X), \text{ and variance } \sigma^2 \text{ of the latent}$$

variable and lower and upper censoring thresholds  $y_l$  and  $y_u$

*Loss:*

$$\begin{aligned} L(Y, F) = & -\log \left( \Phi \left( \frac{y_l - F}{\sigma} \right) \right) \mathbb{1}_{y_l}(Y) \\ & + \left( \frac{(Y - F)^2}{2\sigma^2} + \log(\sigma) + 0.5 \log(2\pi) \right) \mathbb{1}_{\{y_l < Y < y_u\}} \\ & - \log \left( 1 - \Phi \left( \frac{y_u - F}{\sigma} \right) \right) \mathbb{1}_{y_u}(Y) \end{aligned}$$

*Gradient:*

$$\begin{aligned} \frac{\partial L(Y, F)}{\partial F} = & \frac{\phi \left( \frac{y_l - F}{\sigma} \right)}{\sigma \Phi \left( \frac{y_l - F}{\sigma} \right)} \mathbb{1}_{y_l}(Y) - \frac{Y - F}{\sigma^2} \cdot \mathbb{1}_{\{y_l < Y < y_u\}} \\ & - \frac{\phi \left( \frac{y_u - F}{\sigma} \right)}{\sigma \left( 1 - \Phi \left( \frac{y_u - F}{\sigma} \right) \right)} \mathbb{1}_{y_u}(Y) \end{aligned}$$

*Hessian:*

$$\begin{aligned} & \frac{\partial^2 L(Y, F)}{\partial F^2} \\ = & \frac{\phi \left( \frac{y_l - F}{\sigma} \right)}{\sigma^2 \Phi^2 \left( \frac{y_l - F}{\sigma} \right)} \left( \frac{y_l - F}{\sigma} \Phi \left( \frac{y_l - F}{\sigma} \right) + \phi \left( \frac{y_l - F}{\sigma} \right) \right) \mathbb{1}_{y_l}(Y) \\ & + \frac{1}{\sigma^2} \mathbb{1}_{\{y_l < Y < y_u\}} \\ & - \frac{\phi \left( \frac{y_u - F}{\sigma} \right)}{\sigma^2 \left( 1 - \Phi \left( \frac{y_u - F}{\sigma} \right) \right)^2} \left( \frac{y_u - F}{\sigma} \left( 1 - \Phi \left( \frac{y_u - F}{\sigma} \right) \right) - \phi \left( \frac{y_u - F}{\sigma} \right) \right) \mathbb{1}_{y_u}(Y) \end{aligned}$$

- **Mean-scale regression**

$Y|X \sim N(\mu, \sigma^2)$ , with mean  $\mu = F_1(X)$  and standard deviation  $\sigma = e^{F_2(X)}$

*Loss:*  $L(Y, F) = \frac{(Y - F_1)^2}{2e^{2F_2}} + F_2 + 0.5 \log(2\pi)$

*Gradient:*

$$\begin{aligned} \frac{\partial L(Y, F)}{\partial F_1} &= -\frac{Y - F_1}{e^{2F_2}} \\ \frac{\partial L(Y, F)}{\partial F_2} &= -\frac{(Y - F_1)^2}{e^{2F_2}} + 1 \end{aligned}$$

*Hessian:*

$$\begin{aligned} \frac{\partial^2 L(Y, F)}{\partial F_1^2} &= \frac{1}{e^{2F_2}} \\ \frac{\partial^2 L(Y, F)}{\partial F_2^2} &= 2 \frac{(Y - F_1)^2}{e^{2F_2}} \end{aligned}$$

Similarly as for multiclass classification, we assume for simplicity zero off-diagonals for the Hessian, i.e.,  $\frac{\partial^2 L(Y, F)}{\partial F_1 \partial F_2} = 0$ .

## B Using a default value for the minimum number of (weighted) samples per leaf parameter

| Data                     | Grad             | Hybrid                 | Newton                   | XGBoost                |
|--------------------------|------------------|------------------------|--------------------------|------------------------|
| adult                    | 0.129 (0.00211)  | <b>0.128</b> (0.00174) | 0.128 (0.0019)           | 0.128 (0.00159)        |
| bank                     | 0.101 (0.00196)  | 0.101 (0.00162)        | 0.101 (0.00162)          | <b>0.1</b> (0.00208)   |
| cancer                   | 0.0551 (0.015)   | 0.0504 (0.0163)        | 0.0452 (0.0136)          | <b>0.0381</b> (0.0117) |
| ijcnn                    | 0.0172 (0.00105) | 0.0148 (0.000824)      | <b>0.0128</b> (0.000577) | 0.013 (0.000838)       |
| ionosphere               | 0.121 (0.0325)   | 0.118 (0.036)          | 0.107 (0.0357)           | <b>0.102</b> (0.029)   |
| sonar                    | 0.307 (0.0637)   | 0.308 (0.0595)         | 0.289 (0.0514)           | <b>0.254</b> (0.0548)  |
| car                      | 0.0567 (0.0103)  | 0.0523 (0.0121)        | <b>0.0407</b> (0.0122)   | 0.0445 (0.0114)        |
| covtype                  | 0.16 (0.00355)   | 0.16 (0.00377)         | <b>0.154</b> (0.0035)    | 0.159 (0.00371)        |
| digits                   | 0.0657 (0.00663) | 0.0489 (0.0051)        | <b>0.0295</b> (0.00406)  | 0.0362 (0.0035)        |
| glass                    | 0.381 (0.0585)   | 0.372 (0.0576)         | 0.357 (0.0584)           | <b>0.348</b> (0.0582)  |
| letter                   | 0.0917 (0.00558) | 0.075 (0.00424)        | <b>0.0594</b> (0.00357)  | 0.066 (0.00409)        |
| satimage                 | 0.114 (0.00601)  | 0.105 (0.00594)        | <b>0.0971</b> (0.00614)  | 0.102 (0.00639)        |
| smartphone               | 0.033 (0.00482)  | 0.0238 (0.00318)       | <b>0.0159</b> (0.00242)  | 0.0197 (0.00265)       |
| usps                     | 0.0782 (0.00265) | 0.0613 (0.00364)       | <b>0.0435</b> (0.00271)  | 0.054 (0.00341)        |
| insurance                | 51900 (308)      | 51700 (306)            | <b>51600</b> (300)       | 51700 (305)            |
| Av. rank                 | 3.8              | 2.93                   | 1.6                      | 1.67                   |
| p-val Friedman test      | 2.24e-10         |                        |                          |                        |
| Adj. p-val Wilcoxon test | 0.000366         | 0.00061                |                          | 0.489                  |

Table 4: Results for real-world data when the minimum number of (weighted) samples per leaf parameter is not tuned and set to a default value: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets. Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

## C Results for different maximal tree depths

| Data                     | Grad                 | Hybrid                 | Newton                  | XGBoost          |
|--------------------------|----------------------|------------------------|-------------------------|------------------|
| bin_classif              | <b>0.15</b> (0.0466) | 0.151 (0.0467)         | 0.151 (0.0483)          | 0.151 (0.0476)   |
| multi_classif_fht        | 0.328 (0.00691)      | <b>0.191</b> (0.00528) | 0.211 (0.0223)          | 0.381 (0.0173)   |
| digits                   | 0.0509 (0.00484)     | 0.0426 (0.00563)       | <b>0.0423</b> (0.00373) | 0.0582 (0.0155)  |
| satimage                 | 0.126 (0.00587)      | 0.12 (0.00624)         | <b>0.117</b> (0.00636)  | 0.137 (0.0138)   |
| ijcnn                    | 0.0422 (0.00124)     | 0.0401 (0.00118)       | <b>0.04</b> (0.00127)   | 0.0414 (0.00149) |
| poisson_r                | 6930 (79.2)          | 6950 (62.7)            | <b>6930</b> (69.7)      | 11300 (859)      |
| gamma_r                  | 4600 (159)           | <b>4600</b> (160)      | 4600 (161)              | 5090 (820)       |
| tobit_r                  | 4440 (45.7)          | 4430 (48.1)            | <b>4430</b> (47.9)      |                  |
| msr_r                    | 2.73 (0.0271)        | 2.71 (0.0256)          | <b>2.71</b> (0.0263)    |                  |
| malnutrition             | 6.53 (0.00615)       | 6.41 (0.00739)         | <b>6.41</b> (0.00756)   |                  |
| Av. rank                 | 2.57                 | 2.07                   | 1.79                    | 3.57             |
| p-val Friedman test      | 0.196                |                        |                         |                  |
| Adj. p-val Wilcoxon test | 0.168                | 0.343                  |                         | 0.0938           |

Table 5: Results when the maximal tree depth is one (stumps): Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.



| Data                     | Grad              | Hybrid           | Newton                  | XGBoost          |
|--------------------------|-------------------|------------------|-------------------------|------------------|
| bin_classif              | 0.0558 (0.0132)   | 0.0525 (0.0138)  | <b>0.046</b> (0.013)    | 0.0508 (0.0135)  |
| multi_classif_fht        | 0.38 (0.00335)    | 0.344 (0.00968)  | <b>0.3</b> (0.00918)    | 0.358 (0.00447)  |
| digits                   | 0.0359 (0.00473)  | 0.0299 (0.00469) | <b>0.024</b> (0.00342)  | 0.0338 (0.00573) |
| satimage                 | 0.108 (0.00547)   | 0.104 (0.00618)  | <b>0.102</b> (0.00574)  | 0.106 (0.0059)   |
| ijcnn                    | 0.0185 (0.000857) | 0.0167 (0.00116) | <b>0.0141</b> (0.00113) | 0.0158 (0.00083) |
| poisson_r                | 7030 (60.1)       | 6950 (65.3)      | <b>6940</b> (64.4)      | 6990 (54.3)      |
| gamma_r                  | <b>4600</b> (159) | 4610 (158)       | 4610 (158)              | 4640 (156)       |
| tobit_r                  | 4060 (46)         | 4070 (47.9)      | <b>4060</b> (50.4)      |                  |
| msr_r                    | 2.42 (0.0281)     | 2.36 (0.0236)    | <b>2.34</b> (0.0202)    |                  |
| malnutrition             | 6.53 (0.00648)    | 6.41 (0.0077)    | <b>6.41</b> (0.00798)   |                  |
| Av. rank                 | 3.57              | 2.29             | 1.29                    | 2.86             |
| p-val Friedman test      | 0.000886          |                  |                         |                  |
| Adj. p-val Wilcoxon test | 0.129             | 0.129            |                         | 0.0469           |

Table 6: Results when the maximal tree depth is three: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

| Data                     | Grad              | Hybrid            | Newton                   | XGBoost           |
|--------------------------|-------------------|-------------------|--------------------------|-------------------|
| bin_classif              | 0.0464 (0.0122)   | 0.0408 (0.0104)   | <b>0.0379</b> (0.0101)   | 0.0434 (0.0111)   |
| multi_classif_fht        | 0.401 (0.00676)   | 0.366 (0.00593)   | <b>0.337</b> (0.00748)   | 0.389 (0.0118)    |
| digits                   | 0.0351 (0.00589)  | 0.0297 (0.00465)  | <b>0.0234</b> (0.00445)  | 0.0384 (0.00427)  |
| satimage                 | 0.105 (0.00579)   | 0.0981 (0.00716)  | <b>0.0958</b> (0.00604)  | 0.102 (0.00613)   |
| ijcnn                    | 0.0143 (0.000823) | 0.0131 (0.000699) | <b>0.0116</b> (0.000586) | 0.0133 (0.000672) |
| poisson_r                | 7080 (82.5)       | 7020 (75.6)       | <b>6990</b> (77.1)       | 7010 (72)         |
| gamma_r                  | 4610 (158)        | 4620 (156)        | <b>4610</b> (158)        | 4620 (159)        |
| tobit_r                  | 4060 (47.8)       | 4050 (48.2)       | <b>4050</b> (47.4)       |                   |
| msr_r                    | 2.41 (0.0575)     | 2.19 (0.0434)     | <b>2.17</b> (0.0273)     |                   |
| malnutrition             | 6.53 (0.00658)    | 6.42 (0.00823)    | <b>6.42</b> (0.00737)    |                   |
| Av. rank                 | 3.57              | 2.43              | 1                        | 3                 |
| p-val Friedman test      | 2.73e-05          |                   |                          |                   |
| Adj. p-val Wilcoxon test | 0.00586           | 0.00586           |                          | 0.0156            |

Table 7: Results when the maximal tree depth is eight: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

| Data                     | Grad              | Hybrid             | Newton                  | XGBoost          |
|--------------------------|-------------------|--------------------|-------------------------|------------------|
| bin_classif              | 0.0457 (0.0122)   | 0.0415 (0.0112)    | <b>0.0385</b> (0.00998) | 0.0436 (0.0103)  |
| multi_classif_fht        | 0.41 (0.00486)    | 0.388 (0.00579)    | <b>0.35</b> (0.00678)   | 0.394 (0.0138)   |
| digits                   | 0.0351 (0.00555)  | 0.0285 (0.00427)   | <b>0.0232</b> (0.0041)  | 0.0386 (0.00477) |
| satimage                 | 0.104 (0.00604)   | 0.1 (0.00746)      | <b>0.0951</b> (0.00552) | 0.103 (0.00542)  |
| ijcnn                    | 0.0147 (0.000878) | 0.0137 (0.000396)  | <b>0.0118</b> (0.00064) | 0.014 (0.000548) |
| poisson_r                | 7100 (75.6)       | <b>7030</b> (77.3) | 7050 (81.2)             | 7060 (75.3)      |
| gamma_r                  | 4610 (158)        | 4640 (159)         | <b>4610</b> (158)       | 4630 (159)       |
| tobit_r                  | 4070 (47.1)       | <b>4050</b> (46.8) | 4060 (46.3)             |                  |
| msr_r                    | 2.42 (0.0457)     | 2.19 (0.0434)      | <b>2.18</b> (0.0316)    |                  |
| malnutrition             | 6.53 (0.00688)    | 6.42 (0.00783)     | <b>6.42</b> (0.00743)   |                  |
| Av. rank                 | 3.57              | 2.14               | 1.14                    | 3.14             |
| p-val Friedman test      | 0.000298          |                    |                         |                  |
| Adj. p-val Wilcoxon test | 0.00586           | 0.322              |                         | 0.0313           |

Table 8: Results when the maximal tree depth is twenty: Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.

## D Results for the simulated data with a smaller sample size

| Data                     | Grad           | Hybrid               | Newton                 | XGBoost           |
|--------------------------|----------------|----------------------|------------------------|-------------------|
| bin_classif              | 0.111 (0.026)  | 0.103 (0.0248)       | <b>0.0899</b> (0.0233) | 0.106 (0.0247)    |
| bin_classif_fht          | 0.248 (0.0227) | 0.246 (0.0214)       | <b>0.237</b> (0.0216)  | 0.242 (0.0244)    |
| multi_classif            | 0.356 (0.0386) | 0.346 (0.035)        | <b>0.329</b> (0.0349)  | 0.346 (0.0374)    |
| multi_classif_fht        | 0.591 (0.0245) | 0.583 (0.0242)       | <b>0.558</b> (0.025)   | 0.587 (0.0231)    |
| poisson_r                | 778 (40.2)     | 769 (40.5)           | 773 (37.3)             | <b>764</b> (36.9) |
| poisson_f3               | 1210 (17.5)    | 1210 (17)            | <b>1200</b> (18.2)     | 1220 (20.2)       |
| gamma_r                  | 471 (43.5)     | 476 (42.9)           | <b>468</b> (43.3)      | 472 (42.9)        |
| gamma_f3                 | 1440 (11.3)    | 1440 (11.4)          | <b>1440</b> (11.5)     | 1440 (11.6)       |
| tobit_r                  | 445 (18.5)     | <b>442</b> (18.7)    | 443 (18.8)             |                   |
| tobit_f3                 | 480 (16.1)     | <b>478</b> (16)      | 478 (15.6)             |                   |
| msr_f3                   | 3.38 (0.0354)  | <b>3.37</b> (0.0381) | 3.38 (0.0419)          |                   |
| msr_r                    | 3.23 (0.199)   | 3.19 (0.254)         | <b>3.07</b> (0.236)    |                   |
| Av. rank                 | 3.5            | 2.38                 | 1.25                   | 2.88              |
| p-val Friedman test      | 0.000627       |                      |                        |                   |
| Adj. p-val Wilcoxon test | 0.00146        | 0.47                 |                        | 0.297             |

Table 9: Results for simulated data with a sample size of  $n=500$ : Average test error rates for classification and test negative log-likelihoods for regression. In parentheses are approximate standard deviations. Below are average ranks of the methods over the different datasets (only considering datasets for which all four methods are run). Further, a p-value of a Friedman test with an Iman and Davenport correction for comparing the different algorithms is reported. The last row shows Holm-Bonferroni corrected p-values of Wilcoxon signed-rank tests for pairwise comparison of Newton boosting with the novel number of weighted samples parameter and the three alternative methods.