



HAL
open science

Accelerated proximal boosting

Erwan Fouillen, Claire Boyer, Maxime Sangnier

► **To cite this version:**

Erwan Fouillen, Claire Boyer, Maxime Sangnier. Accelerated proximal boosting. 2018. hal-01853244v1

HAL Id: hal-01853244

<https://hal.science/hal-01853244v1>

Preprint submitted on 2 Aug 2018 (v1), last revised 29 Nov 2022 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accelerated proximal boosting

Erwan Fouillen¹, Claire Boyer^{1,2}, and Maxime Sangnier^{1,3}

¹Sorbonne Université, CNRS, LPSM, Paris, France

²Ecole Normale Supérieure, CNRS, DMA, Paris, France

³Sorbonne Université, CNRS, LIP6, Paris, France

August 2, 2018

Abstract

Gradient boosting is a prediction method that iteratively combines weak learners to produce a complex and accurate model. From an optimization point of view, the learning procedure of gradient boosting mimics a gradient descent on a functional variable. This paper proposes to build upon the proximal point algorithm when the empirical risk to minimize is not differentiable. In addition, the novel boosting approach, called *accelerated proximal boosting*, benefits from Nesterov's acceleration in the same way as gradient boosting [Biau et al., 2018]. Advantages of leveraging proximal methods for boosting are illustrated by numerical experiments on simulated and real-world data. In particular, we exhibit a favorable comparison over gradient boosting regarding convergence rate and prediction accuracy.

1 Introduction

Boosting is a celebrated machine learning technique, both in statistics and data science. In broad outline, boosting combines simple models (called weak learners) to build a more complex and accurate model. This assembly is performed iteratively, taking into account the performance of the model built at the previous iteration. The way this information is considered leads to several variants of boosting, the most famous of them being Adaboost [Freund and Schapire, 1997] and gradient boosting [Friedman, 2001].

The reason of the success of boosting is twofold: i) from the statistical point of view, boosting is an additive model with an iteratively growing complexity. In this sense, boosting lies between ensemble methods and high capacity models. In practice, it combines the best of both worlds by reducing the variance and the bias of the risk; ii) from the data science perspective, boosting is a model, the fitting of which is computationally cheap. In contrast, it can quickly achieve highly complex models, thus it is able to perform accurately on difficult learning task. As an ultimate feature, the iterative process makes finding the frontier between under and overfitting quite easy. In particular, gradient boosting combined with decision trees (often referred to as gradient tree boosting) is currently regarded as one of the best off-the-shelf learning techniques in data challenges.

As explained by Biau et al. [2018], gradient boosting has its roots in Freund and Schapire's work on combining classifiers, which resulted in the Adaboost algorithm [Schapire, 1990, Freund, 1995, Freund and Schapire, 1996, 1997]. Later, Friedman and its colleagues developed a novel boosting procedure inspired by the numerical optimization literature, and nicknamed gradient boosting [Friedman et al., 2000, Friedman, 2001, 2002]. Such a connection of boosting between statistics and optimization was already stated in several previous analyses by Breiman [Breiman, 1997, 1998, 1999, 2000, 2004] and reviewed as functional optimization [Mason et al., 2000b,a, Meir and Rätsch, 2003, Bühlmann and Hothorn, 2007]: boosting can be seen as an optimization procedure (similar to gradient descent),

aimed at minimizing an empirical risk over the set of linear combinations of weak learners. In this respect, a few theoretical studies prove the convergence, from an optimization point of view, of boosting procedures [Zhang, 2002, 2003, Wang et al., 2015] and particularly of gradient boosting [Temlyakov, 2012, Biau and Cadre, 2017]. Yet, the topic of optimization (or statistical) convergence is out of the scope of this paper and, as a consequence, will not be covered for the proposed method.

It is quite surprising that in gradient boosting (and variants), the number of weak learners controls both the number of optimization steps performed in order to minimize the empirical risk and the statistical complexity of the final predictor. This early stopping feature of gradient boosting can be seen as an iterative regularization mechanism used to prevent overfitting [Lin et al., 2016]. As a consequence, besides the numerical learning procedure of gradient boosting, its statistical performance deeply relies on the algorithm employed. Especially as early stopping operates jointly with another regularization mechanism: the control of the model complexity.

That being said, one may wonder if gradient descent is really a good option. Following this direction, several alternatives have been proposed, such as using the Frank-Wolfe algorithm instead of a gradient descent [Wang et al., 2015], incorporating second order information [Chen and Guestrin, 2016], and applying Nesterov’s acceleration [Biau et al., 2018]. All these variants rely on differentiable loss functions. The contribution of the work described here is to go a step forward *accelerated gradient boosting* [Biau et al., 2018] by proposing a procedure for learning boosted models on non-differentiable loss functions, with a potential acceleration feature.

To go into details, Section 2 reviews boosting with respect to the empirical risk minimization principle and illustrates the flaw of the current learning procedure in a simple non-differentiable case: least absolute deviations. Then, some backgrounds on non-smooth optimization are stated in Section 3 and we explain the main contribution of this paper: adapting the proximal point algorithm [Nesterov, 2004] (and its Nesterov acceleration) to boosting. The proposed method is nicknamed *accelerated proximal boosting*. A second contribution is the derivation of the weights of each weak learners for accelerated descents (including accelerated gradient boosting [Biau et al., 2018]). Finally, the numerical study described in Section 4 shines a light on advantages and limitations of the proposed boosting procedure.

2 Problem and notation

Let \mathcal{X} be an arbitrary input space and $\mathcal{Y} \subseteq \mathbb{R}$ an output space. Given a pair of random variables $(X, Y) \in \mathcal{X} \times \mathcal{Y}$, supervised learning aims at explaining Y given X , thanks to a function $f_0: \mathcal{X} \rightarrow \mathbb{R}$. In this context, $f_0(X)$ may represent several quantities depending on the task at hand, for which the most notable examples are the conditional expectation $x \in \mathcal{X} \mapsto \mathbb{E}(Y|X = x)$ and the conditional quantiles of Y given X for regression, as well as the regression function $x \in \mathcal{X} \mapsto \mathbb{P}(Y = 1|X = x)$ for ± 1 -classification. Often, this target function f_0 is minimizer of the risk $\mathbb{E}(\ell(f(X), Y))$ over all integrable functions f , where $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a suitable convex loss function (respectively the square function and the pinball loss in the regression examples previously mentioned).

Since the distribution of (X, Y) is generally unknown, the minimization of the risk is out of reach. One would rather deal with its empirical version instead. Let $\{(X_i, Y_i)\}_{1 \leq i \leq n} \subseteq \mathcal{X} \times \mathcal{Y}$ be a training sample independent and identically distributed (*iid*) according to the distribution of (X, Y) and $\mathcal{F} \subseteq L^2(\mu_X)$ a class of functions integrable for the distribution μ_X of X . In this work, we consider estimating f_0 thanks to an additive model f^* , that is $f^* = \sum_{t=0}^T w_t g_t$ for an unknown integer T and an unknown sequence $(w_t, g_t)_t \subseteq \mathbb{R} \times \mathcal{F}$, by solving the following optimization problem:

$$\underset{f \in \text{span } \mathcal{F}}{\text{minimize}} \quad C(f), \tag{P1}$$

where

$$C(f) = \mathbb{E}_n(\ell(f(X), Y)) = \frac{1}{n} \sum_{i=1}^n \ell(Y_i, f(X_i))$$

is the empirical risk and $\text{span } \mathcal{F} = \{\sum_{t=1}^m w_t g_t : w \in \mathbb{R}^m, (f_1, \dots, f_m) \in \mathcal{F}^m, m \in \mathbb{N}\}$ is the set of all linear combinations of functions in \mathcal{F} (\mathbb{N} being the set of non-negative integers).




Figure 1: Prediction and training error of a boosting machine trained with a subgradient (top) and a proximal-based method (bottom).

As a simple example, let us consider the regression model $Y = \sin(2\pi X) + \epsilon$, where X is uniformly distributed on $[0, 1]$ and ϵ is normally distributed and independent of X . We aim at solving:

$$\underset{f \in \text{span } \mathcal{F}}{\text{minimize}} \mathbb{E}_n(|Y - f(X)|),$$

with \mathcal{F} being the set of regression trees of depth less than 3.

Two boosting machines $f_T = \sum_{t=0}^T w_t g_t$ are learned (with T fixed to 300): a traditional one with a subgradient-type method, and another with the proposed proximal-based procedure. Figure 1 depicts the prediction of f_T (left) and the training error $C(f_t) = \mathbb{E}_n(|Y - f_t(X)|)$ along the iterations t (right).

In an optimization perspective, it appears clearly that the subgradient method fails to minimize the empirical risk (prediction is far from the data and the training error is stuck above 10^{-2}) while the proximal-based procedure constantly improves the objective. The subgradient method faces a flaw in convergence, in all likelihood due to non-differentiability of the absolute function $|\cdot|$. This simple example illustrates, inside the boosting paradigm, a well-known fact in numerical optimization: proximal-based algorithms prevails over subgradient techniques.

3 Algorithm

There is an ambiguity in (P1), since it is a functional optimization problem but, in practice, we do not necessarily have the mathematical tools to apply standard optimization procedures (in particular concerning differentiation of C). For this reason, C is often regarded as a function from \mathbb{R}^n to \mathbb{R} , considering that it depends on f only through $(f(X_i))_{1 \leq i \leq n}$. To make this remark more precise, let, for all $z \in \mathbb{R}^n$, $D(z) = \frac{1}{n} \sum_{i=1}^n \ell(Y_i, z_i)$. Then, it is enough to remark that for any $f \in L^2(\mu_X)$, $C(f) = D(z_n(f))$, where $z_n(f) = (f(X_1), \dots, f(X_n)) \in \mathbb{R}^n$ is the vector of f computed on the training sample.

Having this remark in mind helps solving (P1), for instance considering that differentiating C with respect to f is roughly equivalent to differentiating C with respect to $f(x)$ (for all observed $x \in \{X_1, \dots, X_n\}$), thus taking in fact the gradient of D . Doing so, the only requirement is to maintain a functional variable while applying standard vectorial optimization procedures. This requires to *approximate* some vectors by a function, which is at the heart of functional optimization methods such that the ones used in boosting [Mason et al., 2000b].

From now on, all necessary computations of C with respect to f , can be forwarded to D . For instance, if ℓ is differentiable with respect to its second argument, we can define, for all $f \in L^2(\mu_X)$, the functional gradient of C as $\nabla_n C(f) = \nabla D(z_n(f))$. On the contrary, if ℓ is not differentiable, we may consider a subgradient of C at f , denoted $\tilde{\nabla}_n C(f)$ and defined as any subgradient of D at $z_n(f)$.

In the forthcoming sections, a common first order optimization algorithm is reviewed. Then, it is explained how to build different procedures for solving (P1), according to the properties of the loss function ℓ .

3.1 Accelerated proximal gradient method

Let us assume for a while that we want to minimize the function $g + h$, where $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and differentiable (with L -Lipschitz continuous gradient, $L > 0$), and $h: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is convex and lower semi-continuous. Then the iterative procedure defined by choosing any $x_0 = v_0 \in \mathbb{R}^d$ and by setting for all $t \in \mathbb{N}$:

$$\begin{cases} x_{t+1} = \text{prox}_{\gamma h}(v_t - \gamma \nabla g(v_t)) \\ v_{t+1} = x_{t+1} + \alpha_{t+1}(x_{t+1} - x_t) \end{cases}$$

where $\gamma \in (0, 1/L]$ and $(\alpha_t)_t$ will be made precise thereafter, is known to converge to a minimizer of $g + h$ [Nesterov, 2004]. The rate of converge depends on the choice of α_t : if $\alpha_t = 0$ for all $t \in \mathbb{N}$, then the previous procedure leads to the well known proximal gradient method, which converges in $O(1/t)$. More formally, assuming that $g + h$ has a minimizer x^* , then $(g + h)(x_t) - (g + h)(x^*) = O(1/t)$. On the other hand, if one chooses the sequence $(\alpha_t)_t$ defined recursively by:

$$\begin{cases} \beta_0 = 0 \\ \beta_{t+1} = \frac{1 + \sqrt{1 + 4\beta_t^2}}{2}, t \in \mathbb{N} \\ \alpha_{t+1} = \frac{\beta_t - 1}{\beta_{t+1}}, t \in \mathbb{N}, \end{cases} \quad (1)$$

then the convergence becomes $O(1/t^2)$. This is in the spirit of the acknowledged acceleration proposed by Nesterov [1983], and generalized to the composite setting by Beck and Teboulle [2009].

Depending on the properties of the objective function to minimize, the procedure described before leads to two simple algorithms (coming with their acceleration):

- the gradient method:

$$x_{t+1} = v_t - \gamma \nabla g(v_t),$$

minimizes a single function g as soon as it is convex and differentiable with Lipschitz-continuous gradient;

- the proximal point algorithm:

$$x_{t+1} = \text{prox}_{\gamma h}(v_t) = v_t - \gamma \left[\frac{1}{\gamma} (v_t - \text{prox}_{\gamma h}(v_t)) \right],$$

minimizes a single function h , which is only required to be convex and lower semi-continuous (in that case, there is no restriction on the step size γ , except being positive).

Without acceleration (*i.e.* with $\alpha_t = 0$, for all $t \in \mathbb{N}$), the proximal gradient method (as well as its two children) has the asset to be a descent method: at each iteration, the objective function monotonically decreases, meaning that $(g + h)(x_{t+1}) \leq (g + h)(x_t)$, with convergence rate at least $O(1/t)$ ($O(1/t^2)$ with Nesterov's acceleration). In particular, this is true when minimizing a single convex and lower semi-continuous function $h: \mathbb{R}^d \rightarrow \mathbb{R}$, even if it is not differentiable.

This has to be put in contrast with the subgradient method: $x_{t+1} = x_t - \gamma_t \tilde{\nabla} h(x_t)$, where $\gamma_t > 0$ and $\tilde{\nabla} h(x_t)$ is any subgradient of h at x_t . This procedure, which is very similar to the gradient descent but replacing the gradient by any subgradient, has a convergence rate $O(1/\sqrt{t})$ in the best case (when γ_t is well chosen) [Nesterov, 2004]. In addition, this rate is optimal: it cannot be improved without extra assumptions on h [Nesterov, 2004, Theorem 3.2.1]. This means that there does not exist an acceleration scheme for this approach. This remark motivates the use of procedures different from the subgradient method when minimizing a non-differentiable function h , such as the proximal point algorithm described above.

3.2 Gradient boosting

Let \mathcal{F}_0 be the set of constant functions on \mathcal{X} and assume that $\mathcal{F}_0 \subseteq \mathcal{F}$. Also, for any $f \in L^2(\mu_X)$, let us denote $\|f\|_{\mu_n} = \sqrt{\mathbb{E}_{\mu_n}(f(X)^2)} = \sqrt{\frac{1}{n} \sum_{i=1}^n f(X_i)^2}$. Then, a simple procedure to approximately solve (P1) is gradient boosting, described in Algorithm 1 [Mason et al., 2000a, Friedman, 2001]. It builds the requested additive model in an iterative fashion, by imitating a gradient (or subgradient if ℓ is not differentiable with respect to its second argument) method. At each iteration t , Algorithm 1 finds a function g_{t+1} that approximates the opposite of a subgradient of C (also called pseudo-residues) and adds it to the model f_t with a weight $w_{t+1} = \nu \gamma_{t+1}$, where $\nu \in (0, 1]$ is a shrinkage coefficient (also called learning rate) and $\gamma_{t+1} \in \mathbb{R}$ is the optimal step given the direction g_{t+1} . At the end of the procedure, the proposed estimator of f_0 is $f_T = \sum_{t=0}^T w_t g_t$, with $w_0 = 1$.

Algorithm 1 Gradient boosting.

Input: $\nu \in (0, 1]$ (shrinkage coefficient).

1: Set $f_0 = g_0 \in \arg \min_{g \in \mathcal{F}_0} C(g)$ (*initialization*).

2: **for** $t = 0$ **to** $T - 1$ **do**

3: Compute

$$g_{t+1} \in \arg \max_{g \in \mathcal{F}, \|g\|_{\mu_n} \leq 1} \left\langle z_n(g), -\tilde{\nabla}_n C(f_t) \right\rangle_{\ell_2} \quad (\text{correlation})$$

or

$$g_{t+1} \in \arg \min_{g \in \mathcal{F}} \left\| z_n(g) + \tilde{\nabla}_n C(f_t) \right\|_{\ell_2}. \quad (\text{least squares})$$

4: Compute

$$\gamma_{t+1} \in \arg \min_{\gamma \in \mathbb{R}} C(f_t + \gamma g_{t+1})$$

5: Set $f_{t+1} = f_t + \nu \gamma_{t+1} g_{t+1}$. (*update*).

6: **end for**

Output: f_T .

Algorithm 1 requires a number of iterations T , which acts on two regularization mechanisms. The first one is statistical (T controls the complexity of the subspace in which f_T lies) and the second one is numerical (T controls the precision to which the empirical risk C is minimized). The shrinkage coefficient ν tunes the balance between these two regularization mechanisms.

Algorithm 1 has two variants according to the way the subgradient of C is approximated (respectively by correlation or by least squares). The first one closely relates to AdaBoost [Mason et al., 2000a] while the second one is officially known as *gradient boosting* [Friedman, 2001].

Let us remark that the line search (Line 4 of Algorithm 1) simply scales the weak learner g_t by a constant factor. However, when the class \mathcal{F} is a set of regression trees, g_t is a piecewise constant function. In this case, it is common to perform a line search sequentially for each leaf of the decision tree [Friedman, 2001] (called a multiple line search). As a consequence, each level of the piecewise constant function g_t is scaled with its own factor.

3.3 Boosting with non-differentiable loss functions

When the function ℓ is not differentiable with respect to its second argument, gradient boosting just uses a subgradient $\tilde{\nabla}_n C(f_t)$ instead of the gradient $\nabla_n C(f_t)$. This is, of course, convenient but as explained previously, far from leading to interesting convergence properties. For this reason, we propose a new procedure for non-differentiable loss functions ℓ , which consists in adapting the proximal point algorithm [Nesterov, 2004] to functional optimization.

For any $f \in L^2(\mu_X)$, let $\psi_\lambda C(f) = \frac{1}{\lambda} (z_n(f) - \text{prox}_{\lambda D}(z_n(f)))$, where $\lambda > 0$ is a parameter. The simple idea underlying the proposed algorithm, nicknamed *accelerated proximal boosting*, is to replace $\tilde{\nabla}_n C(f_t)$ by $\psi_\gamma C(f_t)$, remarking that $z_n(f_t) - \lambda\psi_\lambda C(f_t) = \text{prox}_{\lambda D}(z_n(f_t))$ is exactly the iteration update of the proximal point method. In addition, this iterative procedure can be sped up by applying Nesterov’s acceleration, reviewed in Section 3.1.

The accelerated proximal boosting procedure is described in Algorithm 2. It is very similar to Algorithm 1, except that the pseudo-residues are now given by a proximal operator instead of a subgradient, and that Nesterov’s acceleration is on.

After approximating the direction of optimization $r \in \mathbb{R}^n$ (also called pseudo-residues) by g_{t+1} , the iterate update (Line 11 in Algorithm 2) becomes (see Section 3.1):

$$f_{t+1} = f_t + \alpha_t(f_t - f_{t-1}) + \nu\gamma_{t+1}g_{t+1}. \quad (2)$$

Similarly to regular gradient boosting (Algorithm 1), the estimator returned at the end of Algorithm 2 can be written $f_T = \sum_{t=0}^T w_t g_t$, where the weights (w_0, \dots, w_T) are now slightly more complicated (this is explained in the next section).

Let us remark that the idea of applying Nesterov’s acceleration to boosting originally appeared in [Biau et al., 2018] for a gradient-type procedure. Even though Biau et al. [2018] did not suggest to apply such an acceleration scheme when ℓ is not differentiable, this idea appears natural. However, this is not entirely relevant since it contradicts the optimization theory (see Section 3.1): subgradient method cannot be accelerated. This flaw clearly motivates using proximal-based methods for non-differentiable boosting, as proposed in Algorithm 2.

3.4 Weights with Nesterov’s acceleration

As an additive model, it is of interest to express f_T with respect to the base learners (g_0, \dots, g_T) and their weights: $f_T = \sum_{t=0}^T w_t g_t$. On the first hand, in the non-accelerated case ($\alpha_t = 0$ for all $t \in \mathbb{N}$), the update rule of Algorithm 1 is simply $f_{t+1} = f_t + \nu\gamma_{t+1}g_{t+1}$. Therefore the weights are defined by $w_0 = 1$ and $w_t = \nu\gamma_t$ for all positive integers t .

On the other hand, when Nesterov’s acceleration is on ($(\alpha_t)_{t \in \mathbb{N}}$ is defined by Equation (1)), the update rule becomes (see Line 11 in Algorithm 2):

$$f_{t'+1} = (1 + \alpha_{t'})f_{t'} - \alpha_{t'}f_{t'-1} + \nu\gamma_{t'+1}g_{t'+1},$$

Algorithm 2 Accelerated proximal boosting.

Input: $\nu \in (0, 1]$ (shrinkage coefficient), $\lambda > 0$ (proximal parameter).

1: Set $g_0 \in \arg \min_{g \in \mathcal{F}_0} C(g)$ (initialization).

2: $x_0 \leftarrow z_n(g_0) \in \mathbb{R}^n$ (predictions).

3: $v_0 = x_0$ (interpolated point).

4: $(w_0^{(0)}, \dots, w_T^{(0)}) \leftarrow (1, 0, \dots, 0)$ (weights of weak learners).

5: **for** $t = 0$ **to** $T - 1$ **do**

6: Compute $r \leftarrow \frac{1}{\lambda} (\text{prox}_{\lambda D}(v_t) - v_t)$ (pseudo-residues).

7: Compute

$$g_{t+1} \in \arg \max_{g \in \mathcal{F}, \|g\|_{\mu_n} \leq 1} \langle z_n(g), r \rangle_{\ell_2} \quad (\text{correlation})$$

or

$$g_{t+1} \in \arg \min_{g \in \mathcal{F}} \|z_n(g) - r\|_{\ell_2}. \quad (\text{least squares})$$

8: Compute

$$\gamma_{t+1} \in \arg \min_{\gamma \in \mathbb{R}} C(f_t + \gamma g_{t+1}).$$

9: Set $x_{t+1} \leftarrow v_t + \nu \gamma_{t+1} z_n(g_{t+1})$.

10: Set $v_{t+1} \leftarrow x_{t+1} + \alpha_{t+1}(x_{t+1} - x_t)$.

11: Set $f_{t+1} \leftarrow f_t + \alpha_t(f_t - f_{t-1}) + \nu \gamma_{t+1} g_{t+1}$, which corresponds to updating weights $(w_0^{(t+1)}, \dots, w_{t+1}^{(t+1)})$ according to Equation (3).

12: **end for**

Output: $f_T = \sum_{t=0}^T w_t^{(T)} g_t$.

for all positive integers $t' \leq T-1$. Let us denote, for each iteration $t' \in \{1, \dots, T-1\}$, $f_{t'} = \sum_{t=0}^{t'} w_t^{(t')} g_t$ the expansion of $f_{t'}$. Then

$$f_{t'+1} = \sum_{t=0}^{t'-1} \left((1 + \alpha_{t'}) w_t^{(t')} - \alpha_{t'} w_t^{(t'-1)} \right) g_t + (1 + \alpha_{t'}) w_{t'}^{(t')} g_{t'} + \nu \gamma_{t'+1} g_{t'+1}.$$

First, we see that the weights of $g_{t'}$ and $g_{t'+1}$ in the expansion of $f_{t'+1}$ are respectively:

$$\begin{cases} w_{t'}^{(t'+1)} = (1 + \alpha_{t'}) w_{t'}^{(t')} \\ w_{t'+1}^{(t'+1)} = \nu \gamma_{t'+1}. \end{cases}$$

Second, for each $t \in \{0, \dots, t' - 1\}$, the weight of g_t in the expansion of $f_{t'+1}$ is defined by:

$$w_t^{(t'+1)} = (1 + \alpha_{t'}) w_t^{(t')} - \alpha_{t'} w_t^{(t'-1)}.$$

Therefore, considering that weights take value 0 before being defined, *i.e.* $w_t^{(t-1)} = 0$, we have:

$$w_t^{(t'+1)} - w_t^{(t')} = \alpha_{t'} (w_t^{(t')} - w_t^{(t'-1)}) = \left(\prod_{k=t}^{t'} \alpha_k \right) (w_t^{(t)} - w_t^{(t-1)}) = \left(\prod_{k=t}^{t'} \alpha_k \right) w_t^{(t)}.$$

It follows that:

$$w_t^{(t'+1)} = w_t^{(t')} + \left(\prod_{k=t}^{t'} \alpha_k \right) w_t^{(t)} = w_t^{(t)} + \sum_{j=t}^{t'} \left(\prod_{k=t}^j \alpha_k \right) w_t^{(t)} = \left(1 + \sum_{j=t}^{t'} \prod_{k=t}^j \alpha_k \right) w_t^{(t)}.$$

Then, for $k \leq 1$, one has $\alpha_k = 0$, so $w_0^{(t'+1)} = w_0^{(0)} = 1$ and $w_1^{(t'+1)} = w_1^{(1)} = \nu\gamma_1$. Now, remarking that, for all $t \geq 2$, $w_t^{(t)} = \nu\gamma_t$, we can conclude that the weights of f_T are:

$$\begin{cases} w_0 = 1 \\ w_1 = \nu\gamma_1 \\ w_t = \left(1 + \sum_{j=t}^{T-1} \prod_{k=t}^j \alpha_k\right) \nu\gamma_t, \forall t \in \{2, \dots, T-1\} \\ w_T = \nu\gamma_T. \end{cases}$$

In a computational perspective, it may be efficient to update the weights, at each iteration t , according to the following recursion:

$$\begin{cases} w_0^{(0)} = 1 \\ w_1^{(0)} = \nu\gamma_1 \\ w_1^{(1)} = \nu\gamma_1 \\ w_j^{(t+1)} = (w_j^{(t)} - w_j^{(t-1)})(1 + \alpha_t) + w_j^{(t-1)}, \forall j \in \{1, \dots, t\} \\ w_{t+1}^{(t+1)} = \nu\gamma_{t+1}. \end{cases} \quad (3)$$

4 Numerical analysis

In Section 3, proximal boosting (Algorithm 2) has been introduced in a fairly general way. However, following the success of gradient boosting, the empirical results presented in this section only relate to a least squares approximation of the pseudo-residues, with decision trees of depth at most 3 as base learners (class \mathcal{F}) and a multiple line search.

In the whole section, the four methods involved in the numerical comparison are nicknamed:

Gradient (slow) : gradient boosting [Friedman, 2001];

Gradient (fast) : accelerated gradient boosting [Biau et al., 2018];

Proximal (slow) : Algorithm 2 without acceleration ($\alpha_t = 0$, for all $t \in \mathbb{N}$);

Proximal (fast) : Algorithm 2 with Nesterov's acceleration (α_t defined by Equation (1)).

4.1 Impact of parameters on algorithm behaviors

This section aims at numerically illustrating, based on synthetic data, the performance of our proximal boosting algorithm and at highlighting the benefits of coupling Nesterov's acceleration scheme with proximal boosting. For this purpose, two synthetic models are studied (see description below), both coming from [Biau et al., 2016, 2018]. The other models considered in [Biau et al., 2018] have also been studied but results are not reported because they are very similar to the two models we focus on.

Regression : $n = 800, d = 100, Y = -\sin(2X_1) + X_2^2 + X_3 - \exp(-X_4) + Z_{0,0.5}$.

Classification : $n = 1500, d = 50, Y = 2\mathbf{1}_{X_1 + X_4^3 + X_9 + \sin(X_{12}X_{18}) + Z_{0,0.1} > 0.38} - 1$, where $\mathbf{1}$ is the indicator function.

The first model covers an additive regression problem, while the second covers a binary classification task with covariate interactions. In both cases, we consider an input random variable $X \in \mathbb{R}^d$, the covariate of which, denoted $(x_j)_{1 \leq j \leq d}$, are either uniformly distributed over $(-1, 1)^d$ (uncorrelated design) or normally distributed with zero mean and covariance matrix $\Sigma = (2^{-|i-j|})_{1 \leq i, j \leq d}$ (correlated design). In these synthetic models, an additive and independent noise (normally distributed with mean $\mu \in \mathbb{R}$ and variance σ^2) is embodied by the random variable Z_{μ, σ^2} .

Four different losses are considered (see Table 1 for a brief description): least squares and least absolute deviations for regression; exponential and hinge for classification. Computations for the corresponding (sub)gradients and proximal operators are detailed in Appendix A. Note that we also considered other kind of losses such as the pinball loss for regression and the logistic loss for classification (see Table 1). Nevertheless, since the numerical behaviors are respectively very close to the least absolute deviations and exponential cases, the results are not reported.

LOSS	PARAMETER	$\ell(y, y')$	TYPE
least squares	-	$(y - y')^2/2$	regression
least absolute deviations	-	$ y - y' $	regression
pinball	$\tau \in (0, 1)$	$\max(\tau(y - y'), (\tau - 1)(y - y'))$	regression
exponential	$\beta > 0$	$\exp(-\beta yy')$	classification
logistic	-	$\log_2(1 + \exp(-yy'))$	classification
hinge	-	$\max(0, 1 - yy')$	classification

Table 1: Loss functions.

In the following numerical experiments, the random sample generated based on each model is divided into a training set (50%) to fit the method and a validation set (50%). The performance of the methods are appraised thanks to several curves representing the training and validation losses along the $T = 5000$ iterations with which the algorithms are run.

4.1.1 Learning rate

This subsection tackles the impact of the learning rate parameter $\nu \in (0, 1]$ on the relative performances of both proximal and accelerated proximal methods. Throughout, the proximal step λ is fixed to 1. The convergence rates for $\nu \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ are illustrated (i) in the case of the regression model, for the least squares loss in Figure 2 and for the least absolute deviations loss in Figure 3; (ii) in the case of the classification model, for the exponential loss in Figure 4 and for the hinge loss in Figure 5.

Some general observations can be drawn, which hold true for all losses and both correlated and uncorrelated designs: as one might expect, a higher learning rate leads to a faster convergence rate. In addition, accelerated boosting converges always faster than the vanilla variant but show unpleasant behaviors for high values of learning rate. Those behaviors depends on the loss and can be: divergence (least squares and least absolute deviations), plateau (exponential) and numerous oscillations (hinge).

As far as we can observe, the minimum validation error achieved depends on the shrinkage factor ν and is generally attained for $\nu \in \{10^{-3}, 10^{-2}, 10^{-1}\}$. This value also depends on the acceleration (activated or not).

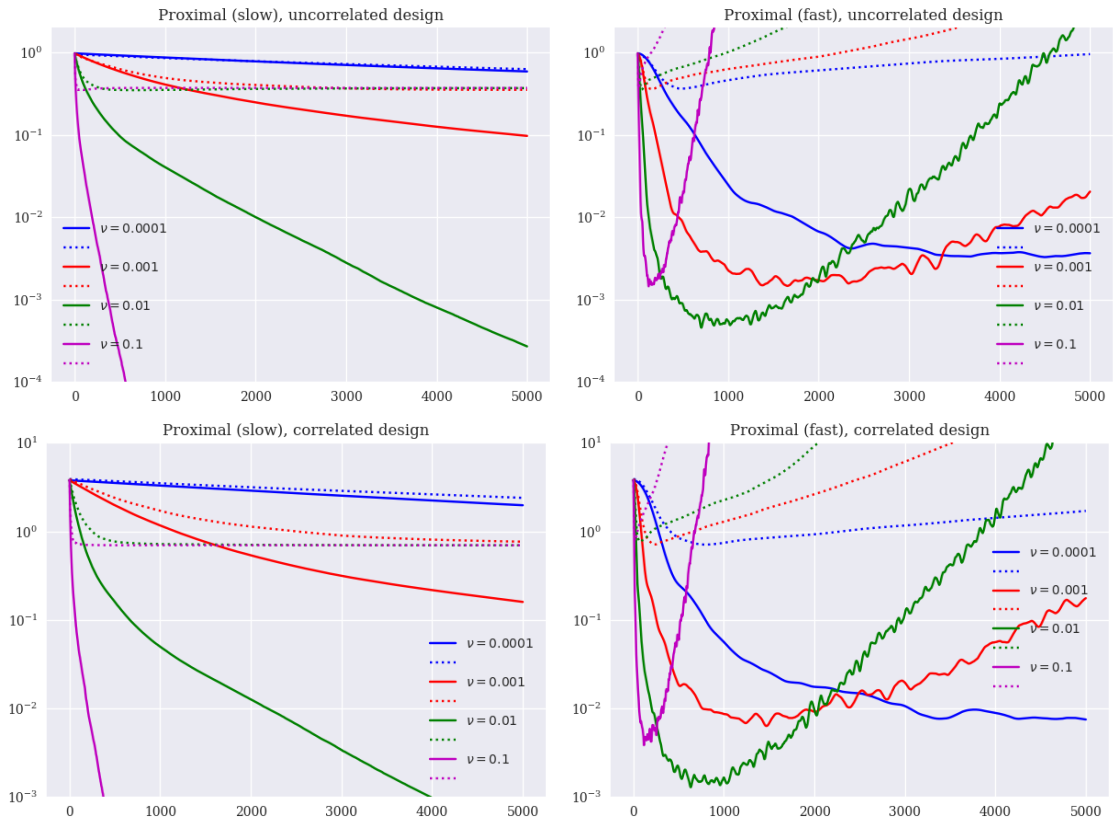


Figure 2: Training (solid lines) and validation (dashed lines) losses for the least squares loss (regression model with λ fixed).




Figure 3: Training (solid lines) and validation (dashed lines) losses for the least absolute deviations loss (regression model with λ fixed).

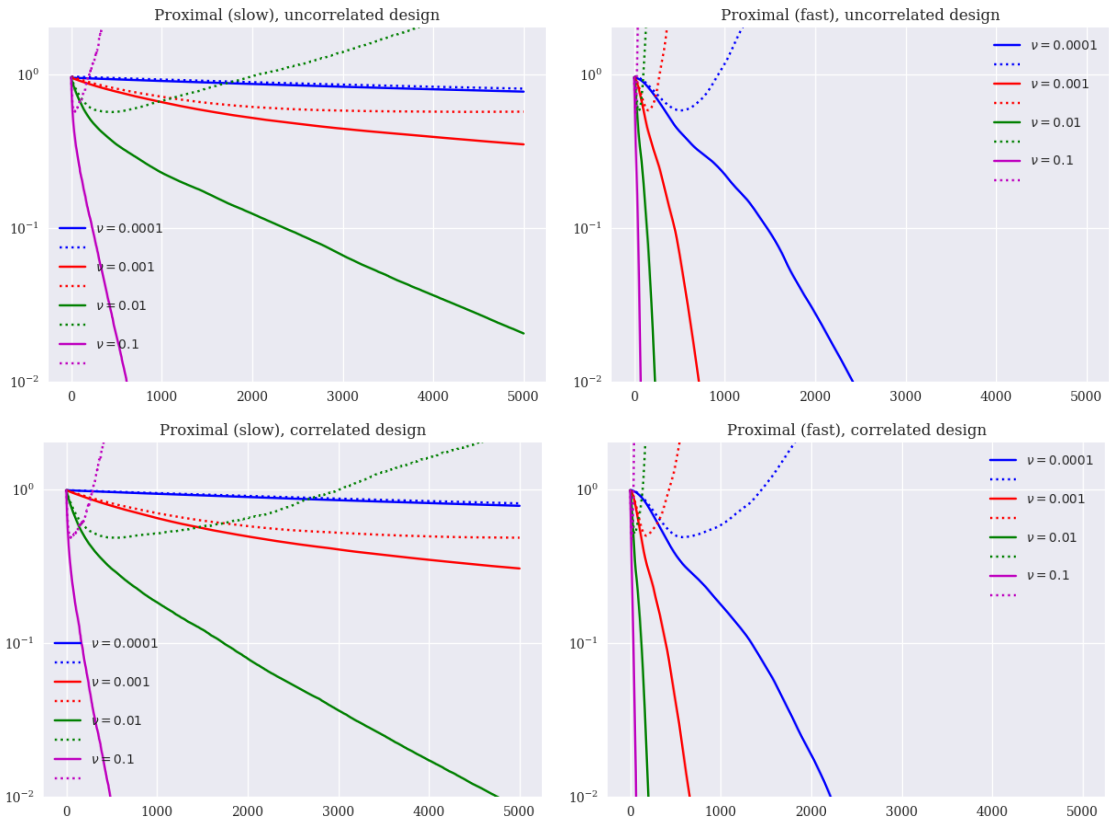


Figure 4: Training (solid lines) and validation (dashed lines) losses for the exponential loss (classification model with λ fixed).

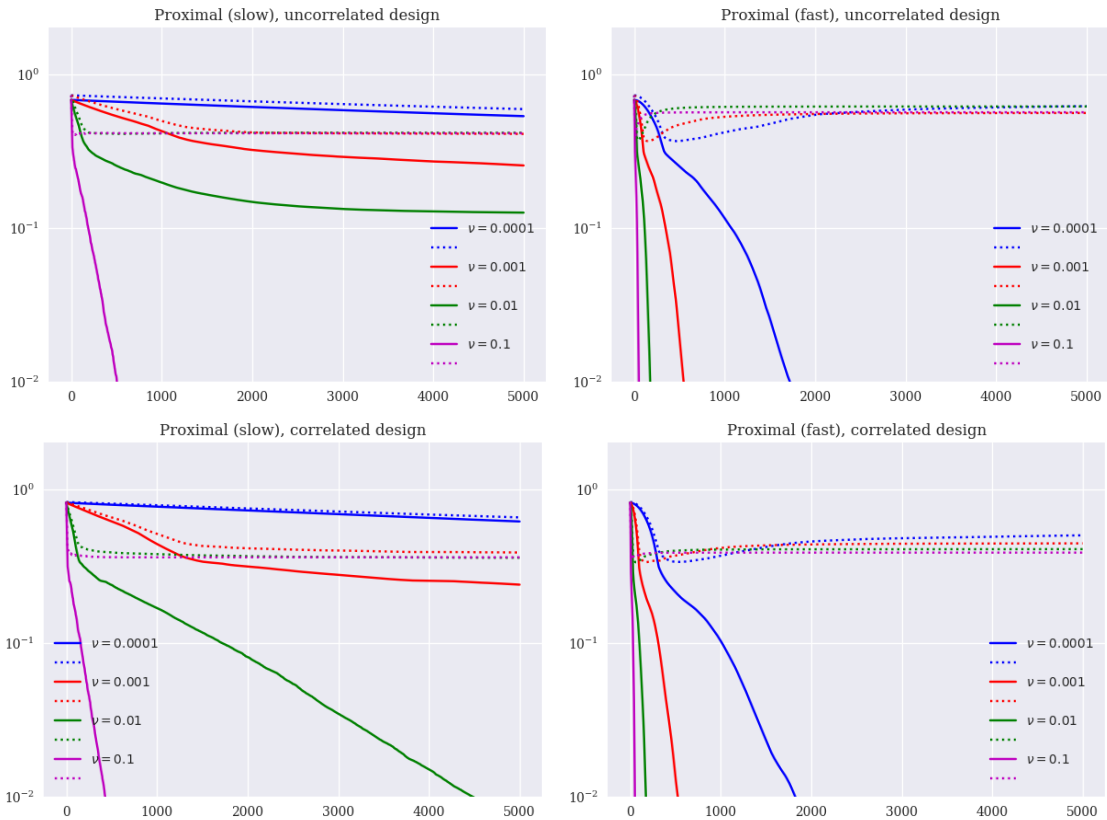


Figure 5: Training (solid lines) and validation (dashed lines) losses for the hinge loss (classification model with λ fixed).

4.1.2 Proximal step

In this subsection, we study the effect of the proximal step parameter $\lambda > 0$ on the performance of both proximal and accelerated proximal boosting methods. Throughout, the learning rate ν is fixed to 10^{-2} . The convergence rates for $\lambda \in \{10^{-2}, 10^{-1}, \dots, 10^2\}$ are illustrated, in the case of the regression model, for the least squares loss in Figure 6 and for the least absolute deviations in Figure 7. In the case of the classification model, numerical results for the exponential and hinge losses, are depicted respectively in Figure 8 and in Figure 9. In all cases, we can observe that non-accelerated method do not converge after 5000 iterations.

For the least squares loss, λ do not seem to have an impact on proximal boosting, which yields similar performances to gradient boosting. For the accelerated variants, there is only a few differences when λ varies and accelerated proximal boosting seems to achieve a lower training error than accelerated gradient boosting for $\lambda = 100$.

Concerning least absolute deviations, λ induces large differences in the training curves, both for regular and accelerated proximal boosting. Besides, (accelerated) proximal boosting achieves a lower training error than (accelerated) gradient boosting with large values of λ ($\lambda > 1$). These observations are similar regarding the exponential loss, except that improvements of proximal boosting occur for small values of λ ($\lambda < 1$) and are less important.

Finally, proximal boosting with the hinge loss depends a lot on the parameter λ but always decreases the training error faster than gradient boosting. For the accelerated version of proximal boosting, unpleasant behaviors may occur and λ should be chosen adequately.

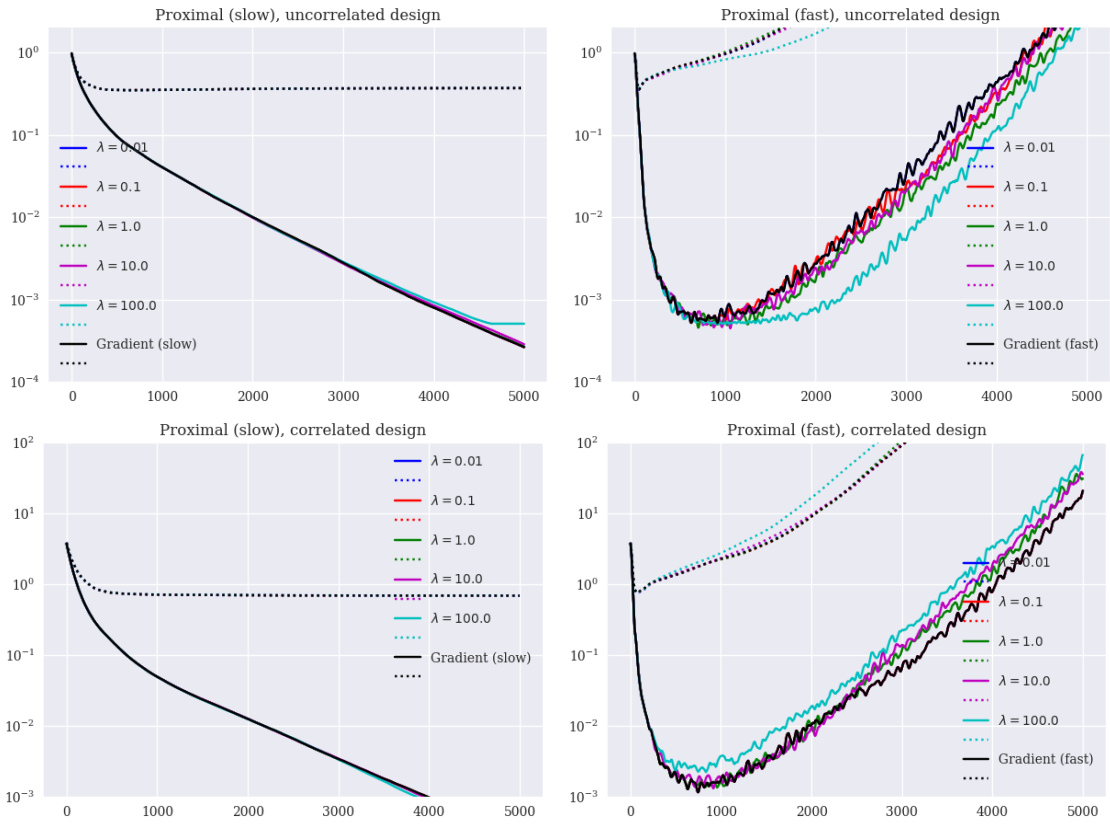


Figure 6: Training (solid lines) and validation (dashed lines) losses for the least squares loss (regression model with ν fixed).




Figure 7: Training (solid lines) and validation (dashed lines) losses for the least absolute deviations loss (regression model with ν fixed).

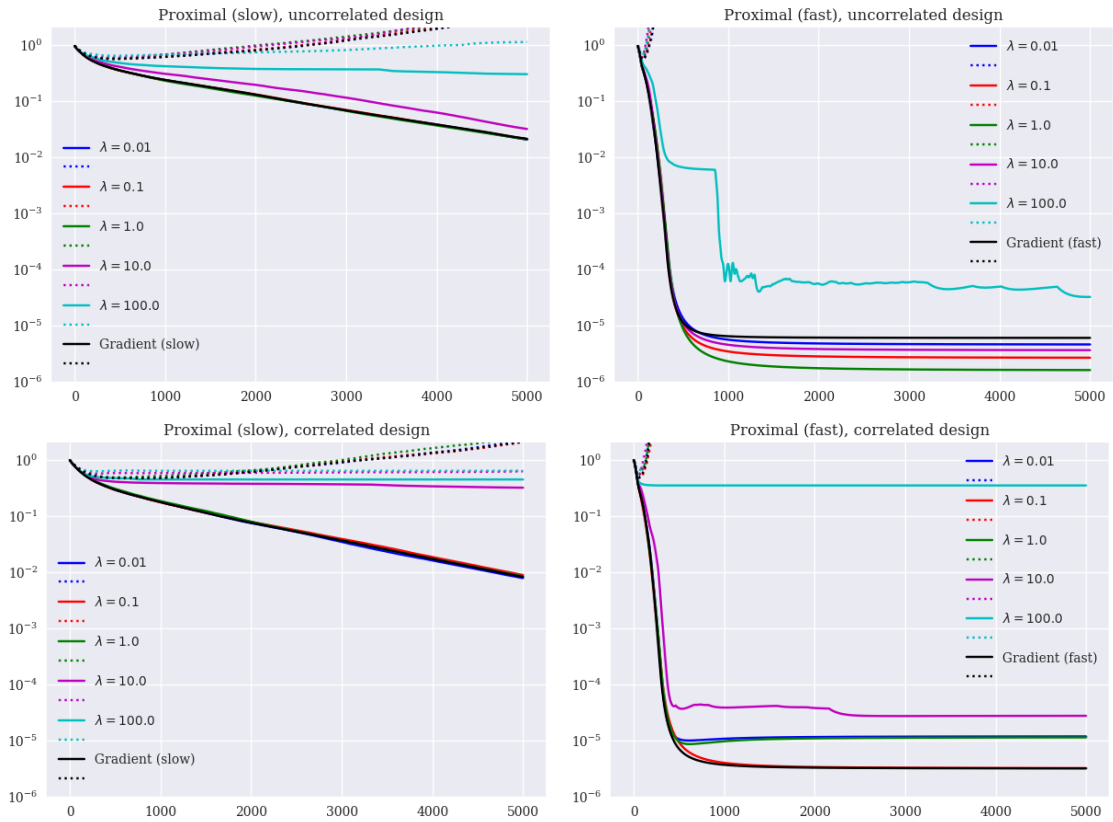


Figure 8: Training (solid lines) and validation (dashed lines) losses for the exponential loss (classification model with ν fixed).

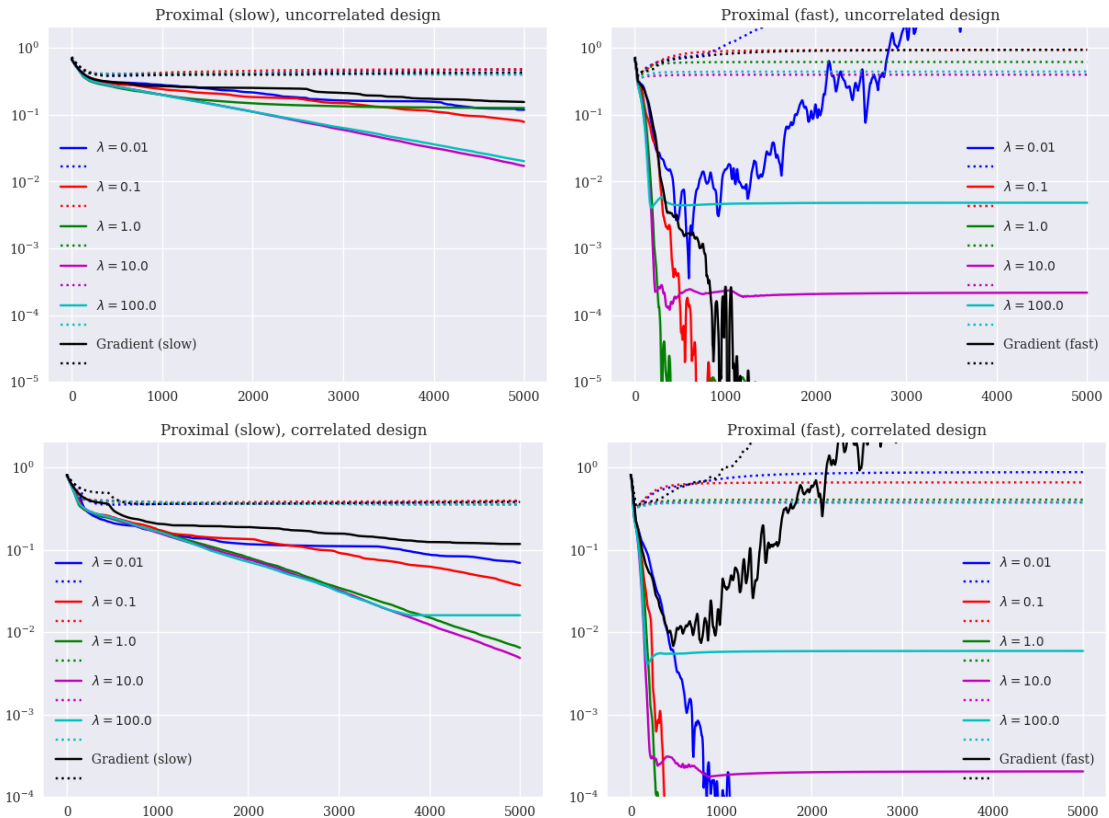


Figure 9: Training (solid lines) and validation (dashed lines) losses for the hinge loss (classification model with ν fixed).

To conclude, the convergence rate as well as the generalization ability of proximal boosting depends on the shrinkage coefficient ν . In addition, the parameter λ (which only exists for proximal boosting) has a noticeable impact on the decrease of the training error. There is no general rule to choose appropriate values for these parameters, but with such values, (accelerated) proximal boosting demonstrates favorable performances, compared to (accelerated) gradient boosting.

4.2 Comparison in real-world cases

This section aims at assessing the generalization ability and the size of the final model for the proposed approach (see Section 3.3) in comparison to known variants. Intuitively, Algorithm 2 is expected to behave better than gradient-type boosting when the loss function ℓ is not differentiable. Benefits are expected on the generalization ability (proximal methods are able to minimize $C(f)$, with $f \in \text{span } \mathcal{F}$, as much as one likes) and on the number of iterations (or weak learners) necessary for producing accurate predictions (weak learners $(g_t)_t$ are more likely to minimize C if they are based on proximal methods).

Comparison is based on four datasets (available on the UCI Machine Learning repository), for which the characteristics are described in Table 2. The first two are univariate regression datasets, while the three others relate to binary classification problems. In both situations, the sample is split into a training set (50%), a validation set (25%) and a test set (25%). The parameters of the methods (number of weak classifiers $T \in [1, 2000]$, learning rate $\nu \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ and proximal step $\lambda \in \{10^{-3}, 10^{-2}, \dots, 10^2\}$) are selected as minimizers of the loss computed on the validation set for models fitted on the training set. Then, models are refitted on the training and the validation sets with

selected parameters. Finally, the generalization capability of the methods are estimated by computing the loss (and the misclassification rate for classification models) on the test set. These quantities are reported thanks to statistics computed on 20 random splits of the datasets.

DATASET	n	d	TYPE
crime	1994	101	regression
wine	1599	11	regression
adult	30162	13	classification
advertisements	2359	1558	classification
spam	4601	57	classification

Table 2: Real-world datasets.

The losses considered in these experiments are least squares, least absolute deviations and pinball ($\tau = 0.9$) for the regression problems, as well as exponential ($\beta = 1$) and hinge for the classification tasks (see Table 1 for a quick definition and Appendix A for the details). For usual regression (least square loss), the four methods described above are also confronted to random forests [Breiman, 2001] with a number of trees selected in accordance with the evaluation procedure described above.

4.2.1 Regression problems

Test losses for the least squares (left), least absolute deviations (middle) and pinball (right) losses are described in Figure 10, along with number of weak learners selected in Figure 11. One can observe that, with respect to both test losses and size of the final predictor, all the boosting methods yield comparable results on the least squares loss (which is differentiable), and better performances than random forests. However, regarding least absolute deviations and pinball losses, proximal approaches achieve similar or better accuracy with equally sized or smaller final models.




Figure 10: Losses on the test datasets for the least squares (left), least absolute deviations (middle) and pinball (right) losses.

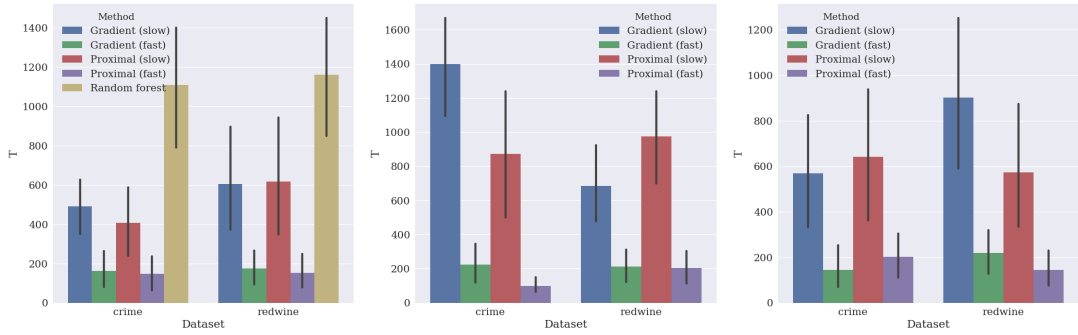


Figure 11: Selected number of weak learners for the least squares (left), least absolute deviations (middle) and pinball (right) losses.

In addition, these results confirm an observation stated in [Biau et al. \[2018\]](#): accelerated boosting (with gradient and proximal-based directions) produces models roughly as accurate as usual boosting but with much less weak learners.

For the two non-differentiable loss functions considered here (least absolute deviations and pinball losses), the good performances of subgradient-based boosting could be explained by the following reason: in this numerical experiment, we are interested in the test loss, which does not require minimizing the training loss entirely. Thus, even though subgradient techniques may not converge, the decrease of the empirical loss may be sufficient for achieving good generalization performances.

4.2.2 Classification problems

Losses (left) and misclassification rates (right) computed on the test dataset are depicted in Figure 12 for exponential (top) and hinge (bottom) losses. One can observe that for the exponential loss, which is a differentiable function, all methods achieve comparable errors and accuracies, up to a slight waste for accelerated variants. This may be explained by the limited values offered to the learning rate in our experimental setting.

However, for the non-differentiable hinge loss, proximal-based boosting clearly outperforms gradient-based techniques, concerning both the loss value and the misclassification rate. This is in agreement with what was expected.

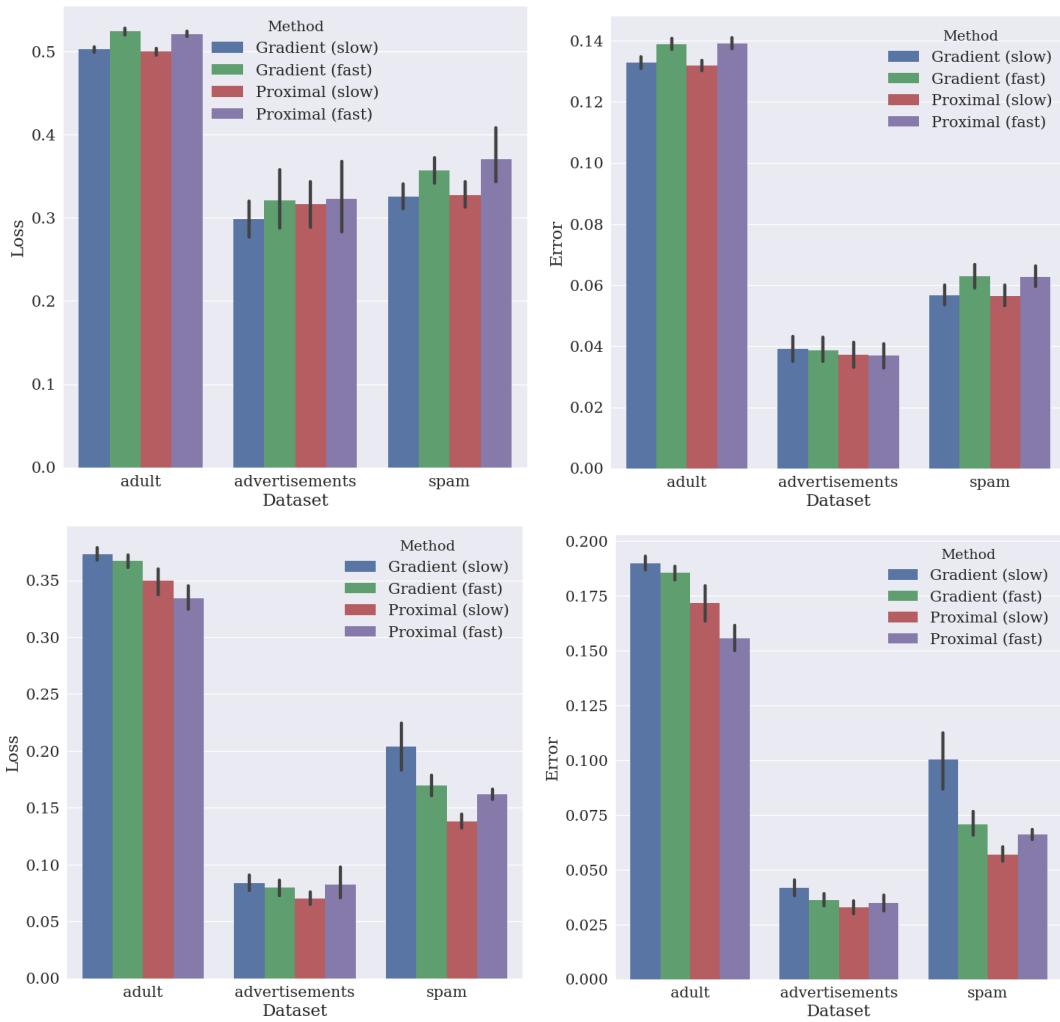


Figure 12: Losses (left) and misclassification rates (right) on the test datasets for the exponential (top) and hinge (bottom) losses.

In addition, Figure 13 shows that the selected number of weak learners are generally greater for proximal boosting (but decreases dramatically when Nesterov’s acceleration is activated). This unfavorable comparison for proximal boosting can be explained by the selected values of shrinkage coefficient ν , which are different for all methods.




Figure 13: Selected number of weak learners for the exponential (left) and the hinge (right) losses.

5 Conclusion

This paper has introduced a novel boosting algorithm, along with an accelerated variant, which have appeal for non-differentiable loss functions ℓ . The main idea is to use a proximal-based direction of optimization, coupled with Nesterov’s acceleration (as already introduced to boosting by [Biau et al. \[2018\]](#)).

Numerical experiments on synthetic data show a significant impact of the newly introduced parameter λ , but also improvements on regular gradient boosting for adequate values of λ . Moreover, in real-world situations, the proposed proximal boosting algorithm achieves comparable or better accuracies than gradient boosting and random forests, depending on the loss employed and the dataset. Finally, it does not appear that proximal boosting needs noticeably less weak learners than gradient boosting but in any case, the size of the final model can be dramatically reduced by activating Nesterov’s acceleration.

We believe that the connection between boosting and functional optimization can be much more investigated. In particular, advances in optimization theory can spread to boosting, like the recently revisited Frank-Wolfe algorithm impacted boosting [[Jaggi, 2013](#), [Wang et al., 2015](#)]. This may also hold true for non-differentiable and non-convex optimization (see for instance [[Ochs et al., 2014](#)]).

Acknowledgement

The authors are thankful to Gérard Biau and Jalal Fadili for enlightening discussions.

References

- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- G. Biau and B. Cadre. Optimization by gradient boosting. *arXiv:1707.05023 [cs, math, stat]*, 2017.
- G. Biau, A. Fischer, B. Guedj, and J.D. Malley. COBRA: A combined regression strategy. *Journal of Multivariate Analysis*, 146:18–28, 2016.

- G. Biau, B. Cadre, and L. Rouvière. Accelerated Gradient Boosting. *arXiv:1803.02042 [cs, stat]*, 2018.
- L. Breiman. Arcing the Edge. Technical Report 486, Statistics Department, University of California, Berkeley, 1997.
- L. Breiman. Arcing classifier (with discussion and a rejoinder by the author). *The Annals of Statistics*, 26(3):801–849, 1998.
- L. Breiman. Prediction Games and Arcing Algorithms. *Neural Computation*, 11(7):1493–1517, 1999.
- L. Breiman. Some Infinite Theory for Predictor Ensembles. Technical Report 577, Statistics Department, University of California, Berkeley, 2000.
- L. Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 2001.
- L. Breiman. Population theory for boosting ensembles. *The Annals of Statistics*, 32(1):1–11, 2004.
- P. Bühlmann and T. Hothorn. Boosting Algorithms: Regularization, Prediction and Model Fitting. *Statistical Science*, 22(4):477–505, 2007.
- T. Chen and C. Guestrin. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, New York, NY, USA, 2016. ACM.
- Y. Freund. Boosting a Weak Learning Algorithm by Majority. *Information and Computation*, 121(2):256–285, 1995.
- Y. Freund and R.E. Schapire. Experiments with a New Boosting Algorithm. In *Proceedings of the Thirteenth International Conference on International Conference on Machine Learning*, San Francisco, CA, USA, 1996.
- Y. Freund and R.E. Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- J. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.
- J. Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, February 2002.
- J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The Annals of Statistics*, 28(2):337–407, 2000.
- M. Jaggi. Revisiting {Frank-Wolfe}: Projection-Free Sparse Convex Optimization. In *Proceedings of The 30th International Conference on Machine Learning*, pages 427–435, 2013.
- J. Lin, L. Rosasco, and D.-X. Zhou. Iterative Regularization for Learning with Convex Loss Functions. *Journal of Machine Learning Research*, 17(77):1–38, 2016.
- L. Mason, J. Baxter, P.L. Bartlett, and M. Frean. Boosting Algorithms as Gradient Descent. In S.A. Solla, T.K. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, pages 512–518. MIT Press, 2000a.
- L. Mason, J. Baxter, P.L. Bartlett, and M. Frean. Functional gradient techniques for combining hypotheses. In A.J. Smola, P.L. Bartlett, B. Shölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 221–246. The MIT Press, 2000b.
- R. Meir and G. Rätsch. An Introduction to Boosting and Leveraging. In *Advanced Lectures on Machine Learning*, Lecture Notes in Computer Science, pages 118–183. Springer, Berlin, Heidelberg, 2003.

- Y. Nesterov. A method of solving a convex programming problem with convergence rate $0(1/k^2)$. *Soviet Mathematics Doklady*, 27, 1983.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, 2004.
- P. Ochs, Y. Chen, T. Brox, and T. Pock. iPiano: Inertial Proximal Algorithm for Nonconvex Optimization. *SIAM Journal on Imaging Sciences*, 2014.
- R.E. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.
- V.N. Temlyakov. Greedy expansions in convex optimization. In *Proceedings of the Steklov Institute of Mathematics*, 2012.
- C. Wang, Y. Wang, W. E, and R. Schapire. Functional Frank-Wolfe Boosting for General Loss Functions. *arXiv:1510.02558 [cs, stat]*, 2015.
- T. Zhang. A General Greedy Approximation Algorithm with Applications. In T.G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 1065–1072. MIT Press, 2002.
- T. Zhang. Sequential greedy approximation for certain convex optimization problems. *IEEE Transactions on Information Theory*, 49(3):682–691, March 2003.

A Implementation details

As explained previously, given a loss function $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, gradient and proximal boosting aim at minimizing the risk functional

$$C(f) = \frac{1}{n} \sum_{i=1}^n \ell(Y_i, F(X_i)) = D(z_n(f))$$

for $f \in \text{span } \mathcal{F}$ (where \mathcal{F} is a class of weak learners $f: \mathbb{R}^d \rightarrow \mathbb{R}$), thereby measuring the cost incurred by predicting $f(X_i)$ when the answer is Y_i .

In the forthcoming subsections, implementation details are given for six popular losses: least squares, least absolute deviations and pinball losses (regression), as well as exponential, logistic and hinge losses (binary classification). In this latter case, the predicted label of a point $x \in \mathbb{R}^d$ is given by +1 if $f(x) \geq 0$ and -1 otherwise.

For each loss, we lay out the following information:

Definition: the mapping of the loss function $\ell: (y, y') \in \mathbb{R}^2 \mapsto \ell(y, y')$.

Initial estimator: the constant function $f_0 \in \arg \min_{f \in \mathcal{F}_0} C(f)$.

Line search: the optimal step size $\gamma_{t+1} \in \arg \min_{\gamma \in \mathbb{R}} C(f_t + \gamma g_t)$.

(Sub)gradient: the direction of optimization to follow at the iterate f_t .

Proximal operator: for all $z \in \mathbb{R}^n$, $\text{prox}_{\lambda n D}(z) = \arg \min_{u \in \mathbb{R}^d} \lambda n D(u) + \frac{1}{2} \|u - z\|_{\ell_2}^2$.

First, for the exponential and the logistic loss, the line search and the proximal operator have no closed-form solution, but are known to be roots of some equations. In that case, we perform one or several steps of the Newton-Raphson method to obtain an approximation of the desired quantity.

Second, when using decision trees as base learners, its common to perform a line search for each leaf of the tree g_t . In that case, the line search may take a simpler form than the one given below.

A.1 Least squares

Definition: $\ell(y, y') = (y - y')^2/2$.

Initial estimator: $f_0 = \frac{1}{n} \sum_{i=1}^n Y_i$.

Line search: $\gamma_{t+1} = \begin{cases} \frac{\sum_{i=1}^n (Y_i - f_t(X_i)) g_{t+1}(X_i)}{\sum_{i=1}^n g_{t+1}(X_i)^2} & \text{if } \sum_{i=1}^n g_{t+1}(X_i)^2 > 0 \\ 0 & \text{otherwise.} \end{cases}$

Gradient: $\nabla_n C(f_t) = ((f_t(X_i) - Y_i)/n)_{1 \leq i \leq n}$.

Proximal operator: $\text{prox}_{\lambda n D}(z) = ((\lambda Y_i + z_i)/(1 + \lambda))_{1 \leq i \leq n}$.

A.2 Least absolute deviations

Definition: $\ell(y, y') = |y - y'|$.

Initial estimator: f_0 is the empirical median of the sample $\{y_1, \dots, y_n\}$.

Line search: $\gamma_{t+1} = \arg \min_{\gamma \in \{0\} \cup \{\frac{Y_i - f_t(X_i)}{g_{t+1}(X_i)} : g_{t+1}(X_i) \neq 0\}} C(f_t + \gamma g_t)$.

Subradient: $\tilde{\nabla}_n C(f_t) = ((\text{sign}(f_t(X_i) - Y_i))/n)_{1 \leq i \leq n}$, where for all $x \in \mathbb{R}$, $\text{sign}(x) = \begin{cases} -1 & \text{if } x < 0 \\ 1 & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases}$

Proximal operator: $\text{prox}_{\lambda n D}(z) = \left(\max\left(0, 1 - \frac{\lambda}{|z_i - Y_i|}\right) (z_i - Y_i) + Y_i \right)_{1 \leq i \leq n}$.

A.3 Pinball

Definition: $\ell(y, y') = \max(\tau(y - y'), (\tau - 1)(y - y'))$, $\tau \in (0, 1)$.

Initial estimator: f_0 is the τ -quantile of the sample $\{y_1, \dots, y_n\}$.

Line search: $\gamma_{t+1} = \arg \min_{\gamma \in \{0\} \cup \{\frac{Y_i - f_t(X_i)}{g_{t+1}(X_i)} : g_{t+1}(X_i) \neq 0\}} C(f_t + \gamma g_t)$.

Subradient: $\tilde{\nabla}_n C(f_t) = \left(\begin{cases} (\tau - 1)/n & \text{if } Y_i - f_t(X_i) < 0 \\ \tau/n & \text{if } Y_i - f_t(X_i) > 0 \\ 0 & \text{otherwise} \end{cases} \right)_{1 \leq i \leq n}$.

Proximal operator: $\text{prox}_{\lambda n D}(z) = \left(\begin{cases} z_i + \lambda\tau & \text{if } Y_i - z_i > \lambda\tau \\ z_i + \lambda(\tau - 1) & \text{if } Y_i - z_i < \lambda(\tau - 1) \\ Y_i & \text{otherwise} \end{cases} \right)_{1 \leq i \leq n}$.

A.4 Exponential loss

Definition: $\ell(y, y') = \exp(-\beta y y')$, $\beta > 0$.

Initial estimator: $f_0 = \frac{\log(\frac{p}{n-p})}{2\beta}$, where $p = \sum_{Y_i=1}^{1 \leq i \leq n} 1$.

Line search: no closed-form solution.

Gradient: $\nabla_n C(f_t) = \left(\frac{-\beta Y_i e^{-Y_i f_t(X_i)}}{n} \right)_{1 \leq i \leq n}$.

Proximal operator: no closed-form solution.

A.5 Logistic loss

Definition: $\ell(y, y') = \log_2(1 + \exp(-yy'))$.

Initial estimator: $f_0 = \log\left(\frac{p}{n-p}\right)$, where $p = \sum_{\substack{1 \leq i \leq n \\ Y_i=1}} 1$.

Line search: no closed-form solution.

Gradient: $\nabla_n C(f_t) = \left(\frac{-Y_i e^{-Y_i f_t(X_i)}}{n \log 2 (1 + e^{-Y_i f_t(X_i)})} \right)_{1 \leq i \leq n}$.

Proximal operator: no closed-form solution.

A.6 Hinge loss

Definition: $\ell(y, y') = \max(0, 1 - yy')$.

Initial estimator: $f_0 = \text{sign}\left(\sum_{i=1}^n Y_i\right)$.

Line search: $\gamma_{t+1} = \arg \min_{\gamma \in \{0\} \cup \left\{ \frac{1 - Y_i f_t(X_i)}{Y_i g_{t+1}(X_i)} : g_{t+1}(X_i) \neq 0 \right\}} C(f_t + \gamma g_t)$.

Subgradient: $\tilde{\nabla}_n C(f_t) = \left(\begin{cases} -Y_i/n & \text{if } Y_i f_t(X_i) < 1 \\ 0 & \text{otherwise} \end{cases} \right)_{1 \leq i \leq n}$.

Proximal operator: $\text{prox}_{\lambda n D}(z) = \left(\begin{cases} z_i + \lambda Y_i & \text{if } Y_i z_i < 1 - \lambda \\ z_i & \text{if } Y_i z_i > 1 \\ Y_i & \text{otherwise} \end{cases} \right)_{1 \leq i \leq n}$.