



HAL
open science

Infinitesimal gradient boosting

Clément Dombry, Jean-Jil Duchamps

► **To cite this version:**

| Clément Dombry, Jean-Jil Duchamps. Infinitesimal gradient boosting. 2021. hal-03209503

HAL Id: hal-03209503

<https://hal.science/hal-03209503>

Preprint submitted on 27 Apr 2021

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.



Distributed under a Creative Commons Attribution 4.0 International License

Infinitesimal gradient boosting

Clément Dombry* and Jean-Jil Duchamps*

April 26, 2021

Abstract

We define infinitesimal gradient boosting as a limit of the popular tree-based gradient boosting algorithm from machine learning. The limit is considered in the vanishing-learning-rate asymptotic, that is when the learning rate tends to zero and the number of gradient trees is rescaled accordingly. For this purpose, we introduce a new class of randomized regression trees bridging totally randomized trees and Extra Trees and using a softmax distribution for binary splitting. Our main result is the convergence of the associated stochastic algorithm and the characterization of the limiting procedure as the unique solution of a nonlinear ordinary differential equation in a infinite dimensional function space. Infinitesimal gradient boosting defines a smooth path in the space of continuous functions along which the training error decreases, the residuals remain centered and the total variation is well controlled.

Keywords: gradient boosting, softmax regression tree, vanishing-learning-rate asymptotic, convergence of Markov processes.

MSC 2020 subject classifications: primary 60F17; secondary 60J20, 62G05.

*Université Bourgogne Franche-Comté, Laboratoire de Mathématiques de Besançon, UMR CNRS 6623, F-25000 Besançon, France. Email: clement.dombry@univ-fcomte.fr, jean-jil.duchamps@univ-fcomte.fr

Contents

1	Introduction and main results	2
1.1	Background	2
1.2	Regression trees and gradient boosting	5
1.3	Main results	9
2	Regression trees	11
2.1	Notation and background	11
2.2	Softmax regression trees	11
3	The space of tree functions	15
3.1	Tree function space and total variation norm	15
3.2	Decomposition of tree functions and L^2 -norm	16
4	Infinitesimal gradient boosting	18
4.1	Existence and uniqueness of the ODE solution	18
4.2	Convergence of gradient boosting	19
4.3	Properties of infinitesimal gradient boosting	21
5	Proofs	24
5.1	Proofs related to Section 2	24
5.2	Proofs related to Section 3	26
5.3	Proofs related to Section 4	28
5.3.1	Proofs related to Section 4.1 and proof of Theorem 1.11	28
5.3.2	Proofs related to Section 4.2 and proof of Theorem 1.9	31
5.3.3	Proofs related to Section 4.3	37
	References	42

1 Introduction and main results

1.1 Background

Trying to understand how a set of covariates $X = (X^1, \dots, X^p)$ impacts a quantity of interest Y is a major task of statistical learning and machine learning that is crucial for predicting the response Y when only X can be observed. Among many modern numeric methods such as random forests or neural networks, gradient boosting is a prominent one, as demonstrated by the incredible success of XGBoost ([Chen and Guestrin, 2016](#)) in machine learning contests. This paper proposes a first in-depth mathematical analysis of the dynamics underneath gradient boosting.

In a prediction framework, statistical learning aims at approaching the performance of the Bayes predictor which minimizes the expected prediction loss and is defined by

$$f^* = \arg \min_f \mathbb{E}[L(Y, f(X))]. \quad (1)$$

Here $f : [0, 1]^p \rightarrow \mathbb{R}$ is a measurable function called predictor and the function L , called loss function, measures the discrepancy between observation Y and prediction $f(X)$. For instance, the fundamental task of regression uses the squared error loss $L(y, z) = (y - z)^2$ and the Bayes predictor is the regression function $f^*(x) = \mathbb{E}[Y \mid X = x]$. Because the distribution of (X, Y) is unknown and accessible only through a sample of observations $(x_i, y_i)_{1 \leq i \leq n}$, the optimization problem (1) is replaced by its empirical counterpart

$$\arg \min_f \frac{1}{n} \sum_{i=1}^n L(y_i, f(x_i)). \quad (2)$$

Here optimization is not performed over the entire function space because it would yield an overfit. In regression, interpolating functions may achieve zero loss but with poor generalization capacity to new observations. Common strategies are to restrict the problem to smaller parametric classes (e.g. linear regression, neural network) and/or to add penalties that impose more regularity of the solution (e.g. smoothing splines, generalized additive models). The choices of a suitable parametric class and/or of suitable penalty terms for regularization are critical and related to the generalization capacity of the predictor in relation with under/over-fitting. When using a parametric model, algorithms from numerical analysis, such as gradient descent and their variants, can be used efficiently to solve the associated finite-dimensional optimization problem. For more general background on statistical learning, the reader should refer to [Hastie et al. \(2009\)](#).

With a different strategy, the gradient boosting method, as proposed by [Friedman \(2001\)](#), is an original approach to tackle the optimization problem (2) directly in the infinite dimensional function space. Akin to gradient descent, it is a recursive procedure that tries to improve the current predictor by performing small steps in suitable directions. A “suitable direction” is obtained by fitting a predictor, called base learner, to the residuals of the current model (i.e. the negative gradient of the empirical loss at the observation points). In practice, the most successful base learners are regression trees — a precise description of regression trees and gradient boosting are provided in the next subsection. Efficient implementation in the R package GBM or Python library XGBoost makes tree-based gradient boosting one of the most useful techniques from modern machine learning.

The development of boosting started with the algorithm AdaBoost for classification by [Freund and Schapire \(1999\)](#), where the idea to combine many weak learners trained sequentially in order to improve the prediction proved successful. [Friedman et al. \(2000\)](#) were able to see the wider statistical framework of stagewise additive modeling which lead to gradient boosting ([Friedman, 2001](#)) and its stochastic version ([Friedman, 2002](#)). In the mathematical analysis of boosting, the main issue discussed in the statistical literature is consistency, i.e. the ability of the procedure to achieve the optimal Bayes error rate when the sample size tends to infinity. Such consistency results were proved for AdaBoost ([Jiang, 2004](#)), more general boosting procedures ([Lugosi and Vayatis, 2004](#); [Blanchard et al., 2004](#); [Zhang and Yu, 2005](#)) or an infinite-population version of boosting ([Breiman, 2004](#)). These papers mostly use the fact that the Bayes predictor can be approximated by linear combinations

taken over a class of simple functions (such as trees) together with some measure of the complexity of the class (such as VC-dimension).

The precise dynamics of the boosting procedure are seldom considered, an exception being the study of linear L2-boosting by [Bühlmann and Yu \(2003\)](#), where the authors consider a linear base learner and provide an explicit expression of the associated boosting procedure relying on linear algebra. Then, the precise knowledge of the behavior of the base learner eigenvalues in the large sample limit allows them to derive consistency. In this linear framework, [Dombry and Esstafa \(2020\)](#) recently introduced the vanishing-learning-rate asymptotic for linear L2-boosting. They proved that, as the learning rate converges to zero and the number of iterations is rescaled accordingly, the boosting procedure converges to the solution of a linear differential equation in a function space. The motivation comes from the fact that small learning rates are known to provide better performances and are commonly used in practice ([Ridgeway, 2007](#)). This vanishing-learning-rate asymptotic sheds some new light on linear L2-boosting, putting the emphasis on the dynamics of the procedure with finite sample size, rather than on its consistency as the sample size tends to infinity.

In the present paper, we extend the vanishing-learning-rate asymptotic beyond linear boosting and obtain the existence of a limit for gradient boosting with a general convex loss function and non linear base learner given by regression trees. Dealing with trees implies new technical issues because the state space of the boosting sequence is now truly infinite dimensional, while the proofs in [Dombry and Esstafa \(2020\)](#) rely on assumptions ensuring that the boosting sequence remains in a finite dimensional space. To tackle this issue, we work in a suitable space that we call the *space of tree functions*. The construction relies on a new encoding of regression trees by discrete signed measures, see [Section 3](#). Another issue with regression trees is their non-linearity and discontinuity with respect to the training sample. Both are due to the splitting procedure that makes use of the response variable in a greedy way, where the arg max functional underneath the best split selection is not continuous. This has led us to design a new class of randomized regression trees, that we call *softmax regression trees*, where the classical arg max selection is replaced by a softmax selection. An important feature of this model is that the expected tree is Lipschitz continuous with respect to the training sample, which makes it possible to develop the differential equation approach from [Dombry and Esstafa \(2020\)](#). Now the dynamics are driven by a nonlinear differential equation and the Lipschitz property ensures the existence and uniqueness of solutions.

We shortly mention the limitation and perspectives of the present work. Strong regularity properties of the expected base learner are assumed that are tailored for the two main important tasks of statistical learning that are least squares regression and binary classification. Further statistical tasks such as quantile regression or robust regression are not covered by the theory we develop here and should be the subject of further research. We establish here a probabilistic theory for the vanishing-learning-rate asymptotic of gradient boosting based on a finite and fixed sample. The variability with respect to the sample distribution and the large sample behavior should be considered in a future work putting the emphasis on statistical

issues. Partial result for linear boosting are provided in [Dombry and Esstafa \(2020\)](#) with a bias/variance decomposition of the training and test errors.

The remainder of this section is devoted to a description of our main results with limited technical details. We present our framework for tree-based gradient boosting and our main results concerning the existence of the vanishing-learning-rate asymptotic (Theorem 1.9) and the characterization of its dynamics in terms of a differential equation (Theorem 1.11). The technical material required to state and prove our results is developed in the next three sections. Section 2 focuses on the base learner and develops the theory of softmax regression trees. Section 3 is devoted to the construction and study of a new function space tailored to the analysis of tree-based gradient boosting. Section 4 proposes a detailed analysis of the vanishing-learning-rate asymptotic for gradient boosting and the associated dynamics. Finally, all the technical proofs are postponed and gathered in Section 5.

1.2 Regression trees and gradient boosting

Regression trees

Gradient boosting is usually implemented with regression trees as base learners. Contrary to random forest where fully grown trees are used, gradient boosting usually makes use of shallow trees, typically with depth between 1 and 5. We refer to [Ridgeway \(2007\)](#) for practical guidelines on gradient boosting. We describe below the classical regression tree models and also introduce a new model that is crucial in our theory.

The original regression tree by [Breiman et al. \(1984\)](#) is a recursive top-down procedure that uses binary splitting to partition the feature space into hypercubes called leaves. Each split is greedy in the sense that it minimizes the mean squared error. More precisely, define the mean squared error on a region A by

$$\text{mse}(A) = \frac{1}{n} \sum_{i=1}^n (y_i - \bar{y}(A))^2 \mathbf{1}_A(x_i), \quad (3)$$

where $\bar{y}(A)$ denotes the mean of $\{y_i : x_i \in A\}$. A split of region A into two subregions A_0 and A_1 gives rise to a decrease of the mean squared error equal to

$$\text{mse}(A) - \text{mse}(A_0) - \text{mse}(A_1) = \frac{1}{n} \left(n(A_0) [\bar{y}(A_0) - \bar{y}(A)]^2 + n(A_1) [\bar{y}(A_1) - \bar{y}(A)]^2 \right),$$

where $n(A_k)$ denotes the number of observations in A_k . The algorithm maximizes this decrease over all admissible splits, where an admissible splits consists in dividing the hypercube A into two hypercubes A_0 and A_1 according to whether some variable X^j is below or above threshold. The search for the best threshold requires mostly to sort the observations according to the value of X^j and to compute cumulative sums of the corresponding values of Y . This is repeated p times to determine both the optimal split variable X^j and the optimal threshold. For large sample and/or high dimensional covariate space, this can still be computationally expensive and the computational burden can be alleviated by the use of Extra-Trees ([Geurts et al., 2006](#)). In this algorithm the splits are randomized and optimization is not performed

over all admissible splits, but only among K random proposals, with K a parameter of the Extra-Trees algorithm. For both algorithms, the regression tree of depth d is grown recursively by applying the greedy binary splitting procedure to the 2^{d-1} leaves of the regression tree of depth $(d-1)$ — starting from a 0-depth tree consisting of a unique leaf defined as the entire hypercube $[0, 1]^p$. This results in a partition of $[0, 1]^p$ into 2^d regions, or leaves. Note that the Extra-Trees algorithm with $K = 1$ involves no optimization and reduces to the so-called *totally randomized tree* where the splits do not depend on the input. On the opposite, as $K \rightarrow \infty$, Extra-Trees approaches Breiman’s regression tree.

Our analysis requires a regularity property of the (expected) tree output with respect to its input $(y_i)_{1 \leq i \leq n}$ that is not satisfied by the two aforementioned models. This lack of regularity is due to the discontinuity of the arg max operator. It is indeed well-known that regression trees tend to be unstable since small modifications in the input may modify the first split and thereby the whole tree structure. Our main idea for regularization is simply to modify the Extra-Trees construction by replacing the hard arg max selection by a smoother softmax selection, where the softmax function is defined by

$$\text{softmax}_\beta(z) = \left(\frac{e^{\beta z_k}}{\sum_{l=1}^K e^{\beta z_l}} \right)_{1 \leq k \leq K}, \quad z \in \mathbb{R}^K, \quad \beta \geq 0, \quad (4)$$

and is interpreted as a probability distribution on $\llbracket 1, K \rrbracket$.

We introduce a new class of randomized regression trees called *softmax regression trees* and whose construction relies on softmax binary splitting defined as follows. The splits are randomized and involve two parameters: $\beta > 0$ and $K \geq 1$. As for Extra-Trees, $K \geq 1$ random proposals are considered for a split. But instead of taking the one maximizing the mse decrease, the split is randomly selected among the K proposals with distribution $\text{softmax}_\beta(\text{scores})$, where the vector of scores records the mse decrease of the K proposed splits. Finally, binary splitting is recursively performed up to depth d so as to obtain 2^d leaves.

Definition 1.1. *We call softmax regression tree with parameter (β, K, d) the tree function*

$$T(x; (x_i, y_i)_{1 \leq i \leq n}, \zeta) = \sum_{1 \leq k \leq 2^d} \bar{y}(A_k) \mathbf{1}_{A_k}(x), \quad x \in [0, 1]^p,$$

where $(x_i, y_i)_{1 \leq i \leq n}$ denotes the input sample, ζ the auxiliary randomness used to perform the splits and $(A_k)_{1 \leq k \leq 2^d}$ the resulting partition of $[0, 1]^p$.

We describe formally in Section 2.2 below what is the structure of the auxiliary randomness ζ and how it determines the partition $(A_k)_{1 \leq k \leq 2^d}$.

Interestingly, the model provides a bridge between the totally randomized tree (as $\beta \rightarrow 0$) and Extra-Trees (as $\beta \rightarrow +\infty$). Furthermore we retrieve Breiman’s regression tree in the limit $K \rightarrow +\infty$ and $\beta \rightarrow +\infty$. See Remarks 2.3 and 2.4 for a formal justification. A property crucial to our theory is that the mean softmax regression tree is Lipschitz continuous in its input $(y_i)_{1 \leq i \leq n}$, see Proposition 2.5.

Tree-based gradient boosting

We focus in this paper on gradient boosting as introduced by [Friedman \(2001\)](#). Recall that the task (2) is to minimize the empirical loss. We assume here that the loss function L is convex and twice differentiable in its second variable. The boosting procedure is an ensemble method that combines many instances of the base learner that are fitted sequentially so as to gradually improve the current model. We use here our softmax regression tree with parameter (β, K, d) as base learner. Gradient boosting with learning rate $\lambda > 0$ then produces the sequence of predictors $(\hat{F}_m^\lambda)_{m \geq 0}$ defined as follows:

1. (Initialization). Set \hat{F}_0^λ equal to the constant predictor:

$$\hat{F}_0^\lambda(x) \equiv \arg \min_{z \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n L(y_i, z), \quad x \in [0, 1]^p. \quad (5)$$

2. (Recursion). At step $m \geq 0$,

- (i) Compute the residuals (negative loss gradients):

$$r_{m,i}^\lambda = -\frac{\partial L}{\partial z}(y_i, \hat{F}_m^\lambda(x_i)), \quad 1 \leq i \leq n; \quad (6)$$

- (ii) Fit a softmax regression tree with parameter (β, K, d) to the residuals:

$$T_{m+1}(x; (x_i, r_{m,i}^\lambda)_{1 \leq i \leq n}, \zeta_{m+1}) = \sum_{1 \leq k \leq 2^d} \bar{r}_m^\lambda(A_k) \mathbf{1}_{A_k}(x), \quad (7)$$

where ζ_{m+1} denotes the auxiliary randomness (independent of the past), $(A_k)_{1 \leq k \leq 2^d}$ the resulting partition of $[0, 1]^p$ into 2^d leaves and $\bar{r}_m^\lambda(A_k)$ the mean residual in leaf A_k ;

- (iii) Modify the leaf values according to a line search one-step approximation:

$$\tilde{T}_{m+1}(x) = \sum_{1 \leq k \leq 2^d} \tilde{r}_m^\lambda(A_k) \mathbf{1}_{A_k}(x), \quad (8)$$

with

$$\tilde{r}_m^\lambda(A_k) = -\frac{\sum_{i=1}^n \frac{\partial L}{\partial z}(y_i, \hat{F}_m^\lambda(x_i)) \mathbf{1}_{A_k}(x_i)}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, \hat{F}_m^\lambda(x_i)) \mathbf{1}_{A_k}(x_i)} \quad (9)$$

- (iv) Update the model by adding a shrunken version of the tree:

$$\hat{F}_{m+1}^\lambda(x) = \hat{F}_m^\lambda(x) + \lambda \tilde{T}_{m+1}(x). \quad (10)$$

We emphasize that at each step, the tree \tilde{T}_{m+1} is obtained by fitting a softmax regression tree to the residuals in order to compute the partition (step *iii*) and then modifying the tree values according to the line search one-step approximation (step *iv*). This motivates the following definition of *softmax gradient trees*.

Definition 1.2. Let $(x_i)_{1 \leq i \leq n}$ be fixed. For a bounded function $F : [0, 1]^p \rightarrow \mathbb{R}$, the residuals at F are given by

$$r_i = -\frac{\partial L}{\partial z}(y_i, F(x_i)), \quad 1 \leq i \leq n.$$

The softmax gradient tree with parameter (β, K, d) is defined as the tree function

$$\tilde{T}(x; F, \zeta) = \sum_{1 \leq k \leq 2^d} \tilde{r}(A_k) \mathbf{1}_{A_k}(x), \quad x \in [0, 1]^p,$$

where ζ denotes the auxiliary randomness, $(A_k)_{1 \leq k \leq 2^d}$ the partition associated with the softmax regression tree $T(\cdot; (x_i, r_i)_{1 \leq i \leq n}, \zeta)$ and

$$\tilde{r}(A_k) = -\frac{\sum_{i=1}^n \frac{\partial L}{\partial z}(y_i, F(x_i)) \mathbf{1}_{A_k}(x_i)}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbf{1}_{A_k}(x_i)}, \quad 1 \leq k \leq 2^d, \quad (11)$$

the leaf values.

With this definition, Equations (6)-(10) take the simple form

$$\hat{F}_{m+1}^\lambda = \hat{F}_m^\lambda + \lambda \tilde{T}(\cdot; \hat{F}_m^\lambda, \zeta_{m+1}), \quad m \geq 0, \quad (12)$$

clearly evidencing the Markov structure of the sequence $(\hat{F}_m^\lambda)_{m \geq 0}$.

Remark 1.3. In the original definition of gradient boosting by Friedman (2001), stagewise additive modeling is considered in a greedy way and the modification of the leaf values in step (iii) is given by

$$\tilde{r}(A_k) = \arg \min_{z \in \mathbb{R}} \sum_{i=1}^n L(r_i, F(x_i) + z) \mathbf{1}_{A_k}(x_i) \quad (13)$$

This corresponds to a line search for optimally updating the current model in an additive way on leaf A_k . Quite often, the line search problem (13) has no explicit solution and numerical optimization has to be used. To alleviate the computational burden, it is usually replaced by its one-step approximation

$$\tilde{r}(A_k) = \arg \min_{z \in \mathbb{R}} \sum_{i=1}^n \left(L(r_i, F(x_i)) + \frac{\partial L}{\partial z} L(r_i, F(x_i)) z + \frac{1}{2} \frac{\partial^2 L}{\partial z^2} L(r_i, F(x_i)) z^2 \right) \mathbf{1}_{A_k}(x_i),$$

where the function to optimize is replaced by its second order Taylor approximation. Solving for this quadratic problem we retrieve exactly Equation (11). Because of its constant use in modern implementation of gradient boosting such as XGBoost (Chen and Guestrin, 2016), we directly use this one-step approximation in our definition of gradient boosting and softmax gradient trees.

Example 1.4. (regression with square loss). In the case of square loss $L(y, z) = \frac{1}{2}(y - z)^2$, the residuals are $r_i = y_i - F(x_i)$. We recover the usual notion of residual, that is the difference between observation and predicted value. In this important case, the line search and its one-step approximation are both equal to the mean residual. The gradient tree is the same as the regression tree, i.e. $T_{m+1} = \tilde{T}_{m+1}$ in Equations (7)-(8). Gradient boosting for L^2 -regression thus consists in sequentially updating the model in an additive way with a shrunken version of the regression tree fitted to the current residuals.

Example 1.5. (binary classification with cross-entropy). Binary classification aims at predicting a binary response variable $Y \in \{0, 1\}$ given $X = x$, the values 0 and 1 being often interpreted as failure and success respectively. The goal is to predict the success probability $p^*(x) = \mathbb{P}(Y = 1 \mid X = x)$. The binary cross entropy corresponds to the negative log-likelihood $L(y, z) = -yz + \log(1 + e^z)$. The Bayes predictor is then equal to $f^*(x) = \log(p^*(x)/(1 - p^*(x)))$, which is the logit of the success probability. The loss derivatives are given by

$$\frac{\partial L}{\partial z}(y, z) = y - p \quad \text{and} \quad \frac{\partial^2 L}{\partial z^2}(y, z) = p(1 - p) \quad \text{with} \quad p = \frac{e^z}{1 + e^z}.$$

Example 1.6. (AdaBoost). The algorithm AdaBoost by [Freund and Schapire \(1999\)](#) was at the origin of the success of boosting. [Friedman \(2001\)](#) showed that it fits in the framework of gradient boosting when the loss is exponential. Here $Y \in \{-1, 1\}$ and the loss function is $L(y, z) = \exp(-yz)$. The Bayes predictor is $f^*(x) = \frac{1}{2} \log(p^*(x)/(1 - p^*(x)))$ with $p^*(x) = \mathbb{P}(Y = 1 \mid X = x)$. The loss derivatives are simply

$$\frac{\partial L}{\partial z}(y, z) = ye^{-yz} \quad \text{and} \quad \frac{\partial^2 L}{\partial z^2}(y, z) = e^{-yz}.$$

Remark 1.7. Gradient boosting is a versatile procedure that can handle many different statistical tasks such as quantile regression or robust regression thanks to suitable choices of the loss function. For instance, the least absolute deviation $L(y, z) = |y - z|$ leads to the median regression and the corresponding Bayes predictor is the conditional median of Y given $X = x$. The theory we develop in this paper uses a strong regularity assumption on the loss function and does not cover this example. The same remark applies for the τ th quantile loss used in quantile regression or for the Huber loss used in robust regression. Such examples will be considered in further research. The regularity assumption we assume in the present paper are tailored for least squares regression and binary classification.

1.3 Main results

We present our main results on the existence of the vanishing-learning-rate asymptotic for gradient boosting based on softmax regression trees and the characterization of the corresponding dynamics. For the sake of simplicity, we state first our results in the Banach space $\mathbb{B} = \mathbb{B}([0, 1]^p, \mathbb{R})$ of bounded functions endowed with the sup norm. However, because of separability issues, we will develop our theory in the space $\mathbb{T} = \mathbb{T}([0, 1]^p, \mathbb{R})$ of tree functions introduced in the next section. We state our results under the following assumptions.

Assumption 1.8. *Let the input $(x_i, y_i)_{1 \leq i \leq n}$ be fixed and consider gradient boosting with loss function L , softmax regression tree with parameter (β, K, d) as base learner and learning rate $\lambda > 0$. Denote by $(\hat{F}_m^\lambda)_{m \geq 0}$ the corresponding sequence of predictors defined by Equations (5)-(10).*

Furthermore, assume the loss function satisfies the following conditions:

(A1) $L : \mathbb{R}^2 \rightarrow [0, \infty)$ is twice differentiable with respect to the second variable and, for all $y \in \mathbb{R}$, $z \mapsto \frac{\partial^2 L}{\partial z^2}(y, z)$ is positive and locally Lipschitz.

(A2) For each $C > 0$, $\sup_{(y,z):L(y,z) \leq C} \left| \frac{\partial L}{\partial z}(y, z) / \frac{\partial^2 L}{\partial z^2}(y, z) \right| < \infty$.

Note that these assumptions are satisfied in the case of regression and binary classification considered in Examples 1.4, 1.5 and 1.6 above.

Equation (12) implies that the sequence $(\hat{F}_m^\lambda)_{m \geq 0}$ is a \mathbb{B} -valued Markov chain. We consider its asymptotic as the learning rate $\lambda \rightarrow 0$ and the number of iterations m is rescaled accordingly. Our first result states that the limit does exist and is deterministic. We call this regime the vanishing-learning-rate asymptotic.

Theorem 1.9. *Under Assumption 1.8, there exists a \mathbb{B} -valued process $(\hat{F}_t)_{t \geq 0}$ such that, for all $T > 0$,*

$$\sup_{t \in [0, T]} \sup_{x \in [0, 1]^p} |\hat{F}_{[t/\lambda]}^\lambda(x) - \hat{F}_t(x)| \longrightarrow 0 \quad \text{in probability as } \lambda \rightarrow 0. \quad (14)$$

Furthermore, the limit process is continuous as a function of t and deterministic.

Note that, due to some path regularity, the processes in Equation (14) are separable and we can freely restrict the supremum to rational values so as to avoid measurability issues.

For $\lambda > 0$, the rescaled sequence $(\hat{F}_{[t/\lambda]}^\lambda)_{t \geq 0}$ defines a càdlàg stochastic process. At times λm , $m \geq 1$, jumps occur as randomized trees are added to the model. Theorem 1.9 states that both jumps and randomness disappear in the vanishing-learning-rate asymptotic. We call the limit process $(\hat{F}_t)_{t \geq 0}$ the *infinitesimal gradient boosting process*. Our second main result is the characterization of this process as the solution of a differential equation in the Banach space \mathbb{B} . We first define the infinitesimal boosting operator that drives the dynamics.

Definition 1.10. *Let $(x_i)_{1 \leq i \leq n}$ be fixed. The infinitesimal boosting operator $\mathcal{T} : \mathbb{B} \rightarrow \mathbb{B}$ is defined by*

$$\mathcal{T}(F)(x) = \mathbb{E}_\zeta[\tilde{T}(x; F, \zeta)], \quad x \in [0, 1]^p, F \in \mathbb{B},$$

where $\tilde{T}(\cdot; F, \zeta)$ denotes the softmax gradient tree from Definition 1.2.

In words, the infinitesimal boosting operator at F is the expectation of the softmax gradient tree used in gradient boosting when updating the predictor F . In general, \mathcal{T} is a non linear operator. It implicitly depends on the sample $(x_i, y_i)_{1 \leq i \leq n}$, on the loss function L and on the parameter (β, K, d) used for the softmax regression trees.

Theorem 1.11. *Suppose Assumption 1.8 is satisfied. In the Banach space \mathbb{B} , consider the differential equation*

$$F'(t) = \mathcal{T}(F(t)), \quad t \geq 0. \quad (15)$$

The following properties are satisfied:

- (i) For all $F_0 \in \mathbb{B}$, Equation (15) admits a unique solution defined on $[0, \infty)$ and started at F_0 ;
- (ii) The infinitesimal gradient boosting process $(\hat{F}_t)_{t \geq 0}$ is the solution of (15) started at $F_0 \equiv \arg \min_{z \in \mathbb{R}} \frac{1}{n} L(y_i, z)$.

Theorem 1.11 reveals the dynamics associated with infinitesimal gradient boosting. In Section 4, we will state a similar result in the space of tree functions and study more precisely the properties of the infinitesimal gradient boosting process in terms of regularity and long time behavior.

2 Regression trees

We give a detailed construction of the softmax regression trees of Definition 1.1 after introducing some notation and recalling background on the usual regression trees (Breiman et al., 1984) and Extra-Trees (Geurts et al., 2006).

2.1 Notation and background

In order to introduce softmax regression trees, let us fix some notation. In this section, we consider a fixed sample $(x_i, r_i)_{1 \leq i \leq n} \in ([0, 1]^p \times \mathbb{R})^n$. Recall from Section 1.2 that regression trees are grown with a top-down recursive procedure based on greedy binary splitting. Here we let couples $(j, u) \in \llbracket 1, p \rrbracket \times [0, 1]$ encode specific splits of a region $A \subset [0, 1]^p$, in the following manner. The split $A = A_0 \cup A_1$ corresponding to (j, u) is defined by

$$A_0 = \{x \in A : x^j < a^j + u(b^j - a^j)\} \quad \text{and} \quad A_1 = \{x \in A : x^j \geq a^j + u(b^j - a^j)\},$$

where x^j denotes the j th coordinate of x , and $a^j = \inf_{x \in A} x^j$ and $b^j = \sup_{x \in A} x^j$.

With this encoding in mind, we define the mse decrease — or simply score — associated with the (j, u) -split of region A by

$$\Delta(j, u; A) = \frac{n(A_0)}{n} (\bar{r}(A_0) - \bar{r}(A))^2 + \frac{n(A_1)}{n} (\bar{r}(A_1) - \bar{r}(A))^2. \quad (16)$$

Using this notation, Breiman's regression trees are grown by recursively choosing a (j, u) -split that is optimal, in the sense of maximizing $\Delta(j, u; A)$, for each region A . In the Extra-Trees algorithm of Geurts et al. (2006), in order to split a region A , we first draw a fixed number $K \geq 1$ of uniformly chosen candidate splits $(j_1, u_1), \dots, (j_K, u_K)$ and then select (j, u) among them so as to maximize the score $\Delta(j, u; A)$. Note that in the case $K = 1$ where only one candidate is considered for each split, we obtain the *totally randomized trees* where the splits do not depend on the input sample $(x_i, r_i)_{1 \leq i \leq n}$.

2.2 Softmax regression trees

We propose a new model of randomized regression trees closely related to Extra-Trees that we call softmax regression trees. The main idea is to replace the strong

arg max selection in the Extra-Trees algorithm by a softmax selection with parameter $\beta \geq 0$. Recall that the softmax function defined in Equation (4) provides a distribution $\text{softmax}_\beta(z)$ over $\llbracket 1, K \rrbracket$ that depends on $z = (z_k)_{1 \leq k \leq K}$ and on $\beta \geq 0$. When dividing the region A , the split (j, u) is randomly chosen among the candidate splits $(j_1, u_1), \dots, (j_K, u_K)$ with distribution

$$\text{softmax}_\beta\left(\left(\Delta(j_k, u_k; A)\right)_{1 \leq k \leq K}\right).$$

Interestingly, when $\beta = 0$, the distribution $\text{softmax}_\beta(z)$ is the uniform distribution on $\llbracket 1, K \rrbracket$ and, as $\beta \rightarrow \infty$, it concentrates on the subset $I_z \subset \llbracket 1, K \rrbracket$ where the components of z are maximal. This implies that the class of softmax regression trees provides a bridge between totally randomized trees (when $\beta = 0$) and Extra-Trees (as $\beta \rightarrow +\infty$). Also, as the number of candidates $K \rightarrow +\infty$ and $\beta \rightarrow +\infty$, we recover Breiman's regression trees. See Remarks 2.3 and 2.4 below for more details.

We propose now a formal definition of the softmax regression tree with parameter (β, K, d) and first set up some notation. The binary rooted tree with depth $d \geq 1$ (from graph theory) is defined on the vertex set $\mathcal{T}_d = \cup_{l=0}^d \{0, 1\}^l$. A vertex $v \in \{0, 1\}^l$ is seen as a word of size l in the letters 0 and 1. The empty word $v = \emptyset$ corresponds to the tree root. The vertex set is divided into the internal nodes $v \in \mathcal{T}_{d-1}$ and the terminal nodes $v \in \{0, 1\}^d$, also called leaves. Each internal node v has two children denoted $v0$ and $v1$ (concatenation of words) while the terminal nodes have no offspring.

A regression tree is encoded by a *splitting scheme*

$$\xi = (j_v, u_v)_{v \in \mathcal{T}_{d-1}} \in (\llbracket 1, p \rrbracket \times (0, 1))^{\mathcal{T}_{d-1}}$$

giving the splits at each internal node, and its leaf values $(\bar{r}_v)_{v \in \{0, 1\}^d}$. The splitting scheme ξ allows to associate to each vertex $v \in \mathcal{T}_d$ a region $A_v = A_v(\xi)$ defined recursively by $A_\emptyset = [0, 1]^p$ and, for $v \in \mathcal{T}_{d-1}$,

$$\begin{aligned} A_{v0} &= A_v \cap \{x \in [0, 1]^p : x^{j_v} < a_v + u_v(b_v - a_v)\}, \\ A_{v1} &= A_v \cap \{x \in [0, 1]^p : x^{j_v} \geq a_v + u_v(b_v - a_v)\}, \end{aligned}$$

with $a_v = \inf_{x \in A_v} x^{j_v}$ and $b_v = \sup_{x \in A_v} x^{j_v}$. Note that A_v depends only on the splits attached to the ancestors of v . For each level $l = 0, \dots, d$, $(A_v)_{v \in \{0, 1\}^l}$ is a partition of $[0, 1]^p$ into 2^l hypercubes. The leaf values are then

$$\bar{r}_v = \bar{r}(A_v) = \frac{1}{n(A_v)} \sum_{i: x_i \in A_v} r_i.$$

Finally, the tree associated with the sample $(x_i, r_i)_{1 \leq i \leq n}$ and splitting scheme $\xi = (j_v, u_v)_{v \in \mathcal{T}_{d-1}}$ is the piecewise constant function

$$T(x; (x_i, r_i)_{1 \leq i \leq n}, \xi) = \sum_{v \in \{0, 1\}^d} \bar{r}(A_v) \mathbf{1}_{A_v}(x). \quad (17)$$

Softmax regression trees are randomized trees and we next describe their distribution. In view of the previous discussion, it is enough to give the distribution of

the associated splitting scheme. Let P_0 be the distribution of the splitting scheme of a totally randomized tree, that is the splits $(j_v, u_v)_{v \in \mathcal{T}_{d-1}}$ are i.i.d. uniformly distributed on $\llbracket 1, p \rrbracket \times (0, 1)$ under P_0 . The candidate splits that appear during the procedure can be gathered into candidate splitting schemes ξ^1, \dots, ξ^K assumed i.i.d. with distribution P_0 . A splitting scheme ξ is constructed by selection among the candidate splits, so that $\xi = (j_v^{\varphi(v)}, u_v^{\varphi(v)})_{v \in \mathcal{T}_{d-1}} =: \xi^\varphi$ for some selection map $\varphi : \mathcal{T}_{d-1} \rightarrow \llbracket 1, K \rrbracket$ specifying which split is selected at each internal node. According to the softmax selection rule, the distribution of ξ given ξ^1, \dots, ξ^K is given by

$$\begin{aligned} \mathbb{P}(\xi = (j_v^{\varphi(v)}, u_v^{\varphi(v)})_{v \in \mathcal{T}_{d-1}} \mid \xi^1, \dots, \xi^K) \\ = \prod_{v \in \mathcal{T}_{d-1}} \left(\text{softmax}_\beta(\Delta(s_v^k, u_v^k; A_v(\xi^\varphi))_{1 \leq k \leq K}) \right)_{\varphi(v)}. \end{aligned} \quad (18)$$

In this equation, the factor indexed by the internal node v corresponds to the probability that the region $A_v(\xi)$ is split according to $(j_v^{\varphi(v)}, u_v^{\varphi(v)})$ among the candidates (j_v^k, u_v^k) , $1 \leq k \leq K$. It should be noted that $A_v(\xi)$ only depends on the splits $(j_{v'}, u_{v'})$ attached to the ancestors v' of v which implies that $(A_v(\xi))_{v \in \mathcal{T}_{d-1}}$ enjoys a branching Markov property which naturally corresponds to the recursive procedure described above.

We next deduce a characterization of the distribution of the splitting scheme ξ , which, in view of Equation (17), characterizes the distribution of the softmax regression tree. The distribution depends on the parameter (β, K, d) and also on the input $(x_i, r_i)_{1 \leq i \leq n}$ but we only write $P_{\beta, K}$ for the sake of readability.

Proposition 2.1. *Let $(x_i, r_i)_{1 \leq i \leq n}$ and (β, K, d) be fixed. We denote by $P_{\beta, K}$ the distribution of the splitting scheme ξ associated with the softmax regression tree with parameter (β, K, d) grown on the sample $(x_i, r_i)_{1 \leq i \leq n}$. Then $P_{\beta, K}$ is absolutely continuous with respect to P_0 with Radon-Nykodym derivative*

$$\frac{dP_{\beta, K}}{dP_0}(\xi) = \int \prod_{v \in \mathcal{T}_{d-1}} \frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi)))}{K^{-1} \sum_{k=1}^K \exp(\beta \Delta(j_v^k, u_v^k; A_v(\xi)))} \delta_\xi(d\xi^1) \otimes_{k=2}^K P_0(d\xi^k), \quad (19)$$

with $\xi^k = (j_v^k, u_v^k)_{v \in \mathcal{T}_{d-1}}$, $1 \leq k \leq K$, and δ_ξ the Dirac mass at ξ .

Note that the Radon-Nykodym derivative (19) is bounded from above by K^{2^d-1} .

Remark 2.2. In Definition 1.1, the softmax regression tree was introduced in terms of an external randomness ζ . In view of the preceding discussion, we can take $\zeta = (\xi^1, \dots, \xi^K, \gamma)$ where ξ^1, \dots, ξ^K denote K independent splitting schemes with distribution P_0 giving the candidate splits and $\gamma = (\gamma_v)_{v \in \mathcal{T}_{d-1}}$ be independent uniform random variables on $[0, 1]$ used to perform the softmax selection at each internal node (using e.g. the probability integral transform). In the following, it will be convenient to write

$$T(\cdot, (x_i, r_i)_{1 \leq i \leq n}, \zeta) = T(\cdot, (x_i, r_i)_{1 \leq i \leq n}, \xi)$$

where the splitting scheme ξ has distribution $P_{\beta, K}$ and can be seen as a (deterministic) function of $\zeta = (\xi^1, \dots, \xi^K, \gamma)$ and $(x_i, r_i)_{1 \leq i \leq n}$.

Remark 2.3. The softmax distribution converges to the uniform distribution as $\beta \rightarrow 0$ and concentrates on the arg max set as $\beta \rightarrow +\infty$. More precisely, for $z \in \mathbb{R}^K$ and $I_z = \arg \max_{1 \leq k \leq K} z_k$,

$$\frac{\exp(\beta z_1)}{\sum_{k=1}^K \exp(\beta z^k)} \longrightarrow \begin{cases} 1/K & \text{as } \beta \rightarrow 0 \\ \mathbb{1}_{I_z}(1)/\text{card}(I_z) & \text{as } \beta \rightarrow +\infty \end{cases} .$$

We deduce readily that $dP_{\beta,K}/dP_0$ has limits when $\beta \rightarrow 0$ or $\beta \rightarrow +\infty$. By Scheffé's theorem, this implies convergence in total variation of the splitting scheme distribution. These limits correspond respectively to totally randomized trees and Extra-Trees.

Remark 2.4. The law of large numbers implies that, for $(\xi^k)_{k \geq 2}$ i.i.d. with distribution P_0 ,

$$\frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi)))}{K^{-1} \sum_{k=1}^K \exp(\beta \Delta(j_v^k, u_v^k; A_v(\xi)))} \longrightarrow \frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi)))}{\mathbb{E}[\exp(\beta \Delta(j, u; A_v(\xi)))]} \quad \text{a.s.},$$

as $K \rightarrow \infty$, with expectation taken with respect to (j, u) uniformly distributed on $\llbracket 1, p \rrbracket \times (0, 1)$. This limit is interpreted as an exponentially tilted distribution for the law of the split (j_v^1, u_v^1) of the region $A_v(\xi)$. Proposition 2.1 and dominated convergence then implies, as $K \rightarrow \infty$,

$$\frac{dP_{\beta,K}}{dP_0}(\xi) \longrightarrow \frac{dP_{\beta,\infty}}{dP_0}(\xi) := \prod_{v \in \mathcal{T}_{d-1}} \frac{\exp(\beta \Delta(j_v, u_v; A_v(\xi)))}{\mathbb{E}[\exp(\beta \Delta(j, u; A_v(\xi)))]},$$

where $\xi = ((j_v, u_v))_{v \in \mathcal{T}_{d-1}}$. This limit corresponds to the distribution of the splitting scheme of a randomized regression tree grown thanks to a binary splitting rule based on the exponentially tilted distribution. Furthermore, letting $\beta \rightarrow +\infty$, the exponentially tilted distribution concentrates on the arg max set so that Breiman's regression trees are recovered.

Our motivation for softmax regression trees is the regularity property of the mean tree defined by

$$\begin{aligned} \bar{T}_{\beta,K,d}(x; (x_i, r_i)_{1 \leq i \leq n}) &= \mathbb{E}_{\zeta}[T(x; (x_i, r_i)_{1 \leq i \leq n}, \zeta)] \\ &= \int T(x; (x_i, r_i)_{1 \leq i \leq n}, \xi) P_{\beta,K}(d\xi). \end{aligned}$$

See Remark 2.2 for the relationship between ζ and ξ and recall that $P_{\beta,K}$ depends on $(x_i, r_i)_{1 \leq i \leq n}$. This will be handled with care in the proof with more explicit notation.

Proposition 2.5. *Let $(x_i)_{1 \leq i \leq n} \in [0, 1]^p$ and (β, K, d) be fixed. The mean softmax regression tree is locally Lipschitz in its input $(r_i)_{1 \leq i \leq n}$, i.e*

$$\begin{cases} (\mathbb{R}^n, \|\cdot\|_{\infty}) & \longrightarrow (\mathbb{B}, \|\cdot\|_{\infty}) \\ (r_i)_{1 \leq i \leq n} & \mapsto \bar{T}_{\beta,K,d}(\cdot; (x_i, r_i)_{1 \leq i \leq n}) \end{cases}$$

is locally Lipschitz.

The same result does not hold when $\beta = \infty$, i.e. for Extra-Trees and Breiman's regression trees, because of the lack of continuity of the arg max selection.

3 The space of tree functions

3.1 Tree function space and total variation norm

An original and powerful point of view to study the sequence of functions $(\hat{F}_m^\lambda)_{m \geq 0}$ produced by gradient boosting is to see it as a Markov chain in a suitable function space that we call the space of tree functions $\mathbb{T} = \mathbb{T}([0, 1]^p, \mathbb{R})$. The main idea is that a regression tree is naturally associated with a signed measure. This is especially useful because the space of non-negative measures enjoys nice compactness properties that will ease tightness considerations. We introduce here the space of tree functions and discuss some useful properties.

We first introduce some notation. Let $\mathcal{M} = \mathcal{M}([0, 1]^p)$ (resp. $\mathcal{M}^+ = \mathcal{M}^+([0, 1]^p)$) denote the space of Borel finite signed measures on $[0, 1]^p$ (resp. Borel finite non-negative measures). The total variation of a signed measure $\mu \in \mathcal{M}$ is

$$\|\mu\|_{\text{TV}} = \sup \left\{ \mu(f) : f \in \mathbb{B}, \|f\|_\infty \leq 1 \right\},$$

where we let $\mu(f) = \int f d\mu$ denote the integral. We note $\mu = \mu^+ - \mu^-$ the Jordan decomposition of μ and $|\mu| = \mu^+ + \mu^-$ the variation of μ — see Billingsley (1995, Section 32) for background on signed measures. We also consider the subset $\mathcal{M}_0 \subset \mathcal{M}$ of signed measures μ satisfying $|\mu|([0, 1]^p \setminus [0, 1]^p) = 0$.

Definition 3.1. *The space of tree functions $\mathbb{T} = \mathbb{T}([0, 1]^p, \mathbb{R})$ is defined as*

$$\mathbb{T} = \left\{ T : [0, 1]^p \rightarrow \mathbb{R} : T(x) = \mu([0, x]) \text{ for some } \mu \in \mathcal{M}_0 \text{ and all } x \in [0, 1]^p \right\}.$$

For all $T \in \mathbb{T}$, the measure $\mu \in \mathcal{M}_0$ such that $T(x) = \mu([0, x])$ is unique and denoted by μ_T . The total variation of T is defined by $\|T\|_{\text{TV}} = \|\mu_T\|_{\text{TV}}$.

Clearly, $\|\cdot\|_{\text{TV}}$ defines a norm on \mathbb{T} , and the mapping $\mu \in \mathcal{M}_0 \mapsto T \in \mathbb{T}$ is an isomorphism so that the space of tree functions $(\mathbb{T}, \|\cdot\|_{\text{TV}})$ is a Banach space. Furthermore, the inclusion $\mathbb{T} \subset \mathbb{B}$ holds with inequality $\|T\|_\infty \leq \|T\|_{\text{TV}}$ so that the injection is continuous.

Remark 3.2. In dimension $p = 1$, $\mathbb{T}([0, 1], \mathbb{R})$ is exactly the set of càdlàg functions with finite total variation that are continuous at 1. In dimension $p \geq 2$, tree functions $T \in \mathbb{T}$ can be seen as the “cumulative distribution function” of a signed measure and therefore have quartant limits and are right-continuous. For instance, when $p = 2$, the four limits $\lim_{y \rightarrow x, y \in Q_x^l} T(y)$ exist for all $x \in [0, 1]^2$ and $l = 1, \dots, 4$, where

$$\begin{aligned} Q_x^1 &= \{y : y^1 \geq x^1, y^2 \geq x^2\}, & Q_x^2 &= \{y : y^1 \geq x^1, y^2 < x^2\}, \\ Q_x^3 &= \{y : y^1 < x^1, y^2 \geq x^2\}, & Q_x^4 &= \{y : y^1 < x^1, y^2 < x^2\}. \end{aligned}$$

Since $x \in Q_x^1$, $T(x)$ is equal to the limit on Q_x^1 . More generally, in dimension $p \geq 1$ there are 2^p such quartant. We refer to Neuhaus (1971) for the definition and properties of the multiparameter Skorokod space $\mathbb{D}([0, 1]^p, \mathbb{R})$. The inclusion $\mathbb{T}([0, 1]^p, \mathbb{R}) \subset \mathbb{D}([0, 1]^p, \mathbb{R})$ holds. For each $T \in \mathbb{T}$, the condition $\mu_T \in \mathcal{M}_0$ implies that T is continuous on $[0, 1]^p \setminus [0, 1]^p$, i.e. at each point having at least one component equal to 1.

We now explain the relationship between regression trees and the space \mathbb{T} .

Proposition 3.3. *Let $\xi \in ([1, p] \times (0, 1))^{\mathcal{T}_{d-1}}$ be a splitting scheme with depth d and $(\tilde{r}_v)_{v \in \{0,1\}^d} \in \mathbb{R}^{\{0,1\}^d}$ be leaf values. Then the function*

$$T(x) = \sum_{v \in \{0,1\}^d} \tilde{r}_v \mathbf{1}_{A_v(\xi)}(x), \quad x \in [0, 1]^p,$$

belongs to \mathbb{T} and satisfies

$$\|T\|_{\text{TV}} \leq 2^{d+\min(d,p)} \|T\|_{\infty} \leq 4^d \|T\|_{\infty}.$$

Furthermore, if $d < p$, the measure μ_T is supported by the set \mathcal{F}_d of points $x \in [0, 1]^p$ with at most d positive components.

Figure 1 provides an illustration of the fact that regression trees belong to \mathbb{T} .




Figure 1: A tree function T with $d = p = 2$. The splitting scheme is represented on the left and corresponds to $\xi = (j_v, u_v)_{v \in \mathcal{T}_{d-1}}$ given by $(j_{\emptyset}, u_{\emptyset}) = (1, 0.4)$, $(j_0, u_0) = (2, 0.5)$ and $(j_1, u_1) = (1, 0.5)$. The leaf values a, b, c, d are below the corresponding leaves. The induced partition $(A_v)_{v \in \{0,1\}^d}$ of $[0, 1]^2$ is represented on the right; the colored dots represent the atoms of the measure μ_T and their labels the associated masses. We can see that $\|T\|_{\text{TV}} \leq 4|a| + 2|b| + 2|c| + |d| \leq 4^2 \|T\|_{\infty}$.

3.2 Decomposition of tree functions and L^2 -norm

We now consider a specific decomposition of tree functions that will be useful in the sequel. For $J \subset [1, p]$ and $\varepsilon \in \{0, 1\}^{J^c}$, define the $|J|$ -dimensional face $\mathcal{F}_{J,\varepsilon}$ of $[0, 1]^p$ by

$$\mathcal{F}_{J,\varepsilon} = \{x \in [0, 1]^p : \forall j \in J, x^j \in (0, 1); \forall j \in J^c, x^j = \varepsilon_j\}.$$

When ε is identically null on J^c , we simply write \mathcal{F}_J . Each measure $\mu \in \mathcal{M}_0$ can be decomposed as

$$\mu = \sum_{J \subset [1,p]} \mu^J \quad \text{with } \mu^J(\cdot) = \mu(\cdot \cap \mathcal{F}_J).$$

Correspondingly, each tree function $T \in \mathbb{T}$ can be decomposed as

$$T = \sum_{J \subset [1,p]} T^J \quad \text{with } T^J(x) = \mu_T^J([0, x]),$$

where $T^J(x)$ depends on x only through the components x^j for $j \in J$. Alternatively the decomposition is characterized recursively by $T^\emptyset = T(0)$ and

$$T^J(x) = T(x^J) - \sum_{J' \subsetneq J} T^{J'}(x), \quad J \neq \emptyset,$$

where $x^J = (x^j \mathbf{1}_{j \in J})_j$ is the vector with components x^j if $j \in J$ and 0 if $j \notin J$.

Remark 3.4.

1. From a statistical perspective, $T^J(x)$ is interpreted as the interaction effect of the covariates x^j , $j \in J$, when the value 0 is taken as reference. This is more easily explained in the case $p = 2$: $T^\emptyset(x) = T(0, 0)$ corresponds to the value of T with both covariates at reference; $T^{\{1\}}(x^1, x^2) = T(x^1, 0) - T(0, 0)$ corresponds to the additional effect of x^1 with x^2 at reference and similarly for $T^{\{2\}}(x^1, x^2)$; finally, $T^{\{1,2\}}(x^1, x^2) = T(x^1, x^2) - T(x^1, 0) - T(0, x^2) + T(0, 0)$ corresponds to the additional interaction effect between x^1 and x^2 .
2. According to Proposition 3.3, the measure μ_T of a regression tree T with depth $d < p$ is supported by $\mathcal{F}_d = \cup_{|J| \leq d} \mathcal{F}_J$. This implies that $T^J = 0$ whenever $|J| > d$, i.e. there are no interaction effect involving more than d variables. In particular, when $d = 1$, $T(x) = T(0) + \sum_{j=1}^p T^{\{j\}}(x)$ where $T^{\{j\}}(x)$ depends on x^j only, which corresponds to an additive model in the covariates x^1, \dots, x^p .

Next we consider \mathbb{T} as a subset of a well-chosen L^2 space. This is useful because the total variation norm induces a topology that is too strong for our purpose, making the Banach space $(\mathbb{T}, \|\cdot\|_{\text{TV}})$ non-separable. For $J \subset \llbracket 1, p \rrbracket$ and $\varepsilon \in \{0, 1\}^{J^c}$, we define the measure $\text{Leb}_{J,\varepsilon}$ on $[0, 1]^p$ supported by the $|J|$ -dimensional face $\mathcal{F}_{J,\varepsilon}$ by

$$\text{Leb}_{J,\varepsilon}(dx) = \prod_{j \in J^c} \delta_{\varepsilon_j}(dx^j) \prod_{j \in J} dx^j.$$

We consider the probability measure ν on $[0, 1]^p$ defined by

$$\nu = \frac{1}{3^p} \sum_{J,\varepsilon} \text{Leb}_{J,\varepsilon}$$

and the Hilbert space $L^2 = L^2([0, 1]^p, \nu)$ with associated norm $\|\cdot\|_2$. Note that if the tree functions $T_1, T_2 \in \mathbb{T}$ are equal $\nu(dx)$ -almost everywhere, then they are equal everywhere because they are right-continuous in the sense of Remark 3.2. This implies that $\|\cdot\|_2$ is a norm on \mathbb{T} . Therefore we can write $\mathbb{T} \subset L^2$, and we have $\|T\|_2 \leq \|T\|_\infty \leq \|T\|_{\text{TV}}$ for all $T \in \mathbb{T}$. Note that \mathbb{T} is not closed in L^2 since for instance $\mathbf{1}_{[1/n, 1]^p} \in \mathbb{T}$ converges in L^2 to $\mathbf{1}_{(0, 1]^p} \notin \mathbb{T}$ as $n \rightarrow +\infty$.

Proposition 3.5.

- (i) \mathbb{T} is dense in L^2 .
- (ii) Any subset $K \subset \mathbb{T}$ satisfying $\sup_{T \in K} \|T\|_{\text{TV}} < \infty$ is relatively compact in L^2 .
- (iii) Let $T, T_1, T_2, \dots \in \mathbb{T}$ be such that $(\|T_n\|_{\text{TV}})_{n \geq 1}$ is bounded. Then, $\|T_n - T\|_2 \rightarrow 0$ if and only if $\mu_{T_n}^J \xrightarrow{w} \mu_T^J$ for all $J \subset \llbracket 1, p \rrbracket$, where \xrightarrow{w} denotes weak convergence of measures.

To deal with tightness and uniform convergence in the proof of Theorem 1.9, it will be convenient to consider the subset $\mathbb{T}^+ = \{T \in \mathbb{T} : \mu_T \in \mathcal{M}^+\}$ of tree functions with positive measure. We denote by $\overline{\mathbb{T}^+}$ the adherence of \mathbb{T}^+ in L^2 and equip it with the induced metric.

Proposition 3.6.

- (i) $\overline{\mathbb{T}^+}$ is a proper metric space, i.e. bounded sets are relatively compact.
- (ii) Let $T, T_1, T_2, \dots \in \mathbb{T}^+$ and assume T is continuous. Then, $\|T_n - T\|_2 \rightarrow 0$ implies $\|T_n - T\|_\infty \rightarrow 0$.

When dealing with gradient boosting, we need the function F to be well-defined at $(x_i)_{1 \leq i \leq n}$ because the algorithm involves the values $F(x_i)$. The following remark provides a slight modification of the previous result in order to deal with this issue and introduces the function space $L_{\mathbf{x}}^2$ that will be useful in Section 4.

Remark 3.7. With the sample $\mathbf{x} = (x_i)_{i=1}^n$ being fixed, we consider the space $L_{\mathbf{x}}^2 = L^2([0, 1]^p, \nu_{\mathbf{x}})$, where $\nu_{\mathbf{x}}$ is the probability measure

$$\nu_{\mathbf{x}} = \frac{\nu}{2} + \frac{1}{2n} \sum_{i=1}^n \delta_{x_i}.$$

For $F \in L_{\mathbf{x}}^2$, the values $(F(x_i))_{1 \leq i \leq n}$ are well defined and we can consider the gradient tree $\tilde{T}(\cdot; F, \xi)$ for $F \in L_{\mathbf{x}}^2$.

It is readily checked that $\mathbb{T} \subset L_{\mathbf{x}}^2$, with $\|T\|_{L_{\mathbf{x}}^2} \leq \|T\|_{\text{TV}}$ and that the statements and proofs of Propositions 3.5 and 3.6 still hold with L^2 replaced by $L_{\mathbf{x}}^2$, the only modification being that (iii) in Proposition 3.5 has to be replaced by

- (iii') Let $T, T_1, T_2, \dots \in \mathbb{T}$ be such that $(\|T_n\|_{\text{TV}})_{n \geq 1}$ is bounded. Then, $\|T_n - T\|_{L_{\mathbf{x}}^2} \rightarrow 0$ if and only if $\mu_{T_n}^J \xrightarrow{w} \mu_T^J$ for all $J \subset \llbracket 1, p \rrbracket$ and $T_n(x_i) \rightarrow T(x_i)$ for all $i \in \llbracket 1, n \rrbracket$.

The proof is easily adapted thanks to the fact that $L_{\mathbf{x}}^2$ is isomorphic to $L^2 \times \mathbb{R}^n$ via the isomorphism $F \mapsto (F, (F(x_i))_{1 \leq i \leq n})$.

4 Infinitesimal gradient boosting

4.1 Existence and uniqueness of the ODE solution

For the sake of simplicity, our results were first stated in Section 1.3 in the Banach space $\mathbb{B} = \mathbb{B}([0, 1]^p, \mathbb{R})$ of bounded function. We now state and prove a version of Theorem 1.11 in the Banach space \mathbb{T} of tree functions endowed with the total variation norm. Existence and uniqueness of the ODE solution follow essentially from the fact that the infinitesimal boosting operator introduced in Definition 1.10 is locally Lipschitz. The proof is similar to the proof of Proposition 2.5 but slightly more involved because we have to consider softmax gradient trees (Definition 1.2), that have a more complex structure than softmax regression trees (Definition 1.1).

Proposition 4.1. *Suppose Assumption 1.8 is satisfied. Then the infinitesimal boosting operator \mathcal{T} defined in Definition 1.10 satisfies $\mathcal{T}(F) \in \mathbb{T}$, for all $F \in \mathbb{B}$, and the map $\mathcal{T} : (\mathbb{B}, \|\cdot\|_\infty) \rightarrow (\mathbb{T}, \|\cdot\|_{\text{TV}})$ is locally Lipschitz.*

Remark 4.2. Because $\mathbb{T} \subset \mathbb{B}$ and $\|T\|_\infty \leq \|T\|_{\text{TV}}$ for all $T \in \mathbb{T}$, the previous proposition implies that \mathcal{T} can be seen as a locally Lipschitz operator on $(\mathbb{B}, \|\cdot\|_\infty)$ and $(\mathbb{T}, \|\cdot\|_{\text{TV}})$. This is useful when considering the ODE (15) on both spaces, in Theorem 1.11 or Theorem 4.3 below respectively.

Theorem 4.3. *Consider the differential equation (15) in the space $(\mathbb{T}, \|\cdot\|_{\text{TV}})$. Under Assumption 1.8, for any initial condition $F_0 \in \mathbb{T}$, it has a unique solution defined on $[0, \infty)$ and satisfying $F(0) = F_0$.*

Since the differential equation (15) is driven by the infinitesimal boosting operator which is locally Lipschitz according to Proposition 4.1, the Cauchy–Lipschitz theorem ensures the existence and uniqueness of local solutions. The proof of Theorem 4.3 consists in verifying that the maximal solution corresponding to the local solution started at F_0 at time $t = 0$ is defined for all time $t \geq 0$.

4.2 Convergence of gradient boosting

In this section, we consider convergence of the gradient boosting sequence $(\hat{F}_m^\lambda)_{m \geq 0}$ in the vanishing learning rate asymptotic $\lambda \rightarrow 0$. Theorem 1.9 states, for all $T > 0$, the convergence in probability

$$(\hat{F}_{\lfloor t/\lambda \rfloor}^\lambda)_{t \in [0, T]} \xrightarrow{\mathbb{P}} (\hat{F}_t)_{t \in [0, T]}$$

in the Skohorod space $\mathbb{D}([0, T], \mathbb{B})$ endowed with the norm of uniform convergence. The non-separability of the Banach space \mathbb{B} makes it not suitable for the analysis and we first develop our results in the separable space $L_{\mathbf{x}}^2$ introduced in Section 3.2. More precisely, we prove the convergence of the positive and negative parts of the gradient boosting sequence seen in the space $\overline{\mathbb{T}}^+ \subset L_{\mathbf{x}}^2$. The fact that $\overline{\mathbb{T}}^+$ is a proper metric space (see Proposition 3.6 *i*) makes tightness issues easily tractable. Furthermore, the fact that L^2 -convergence in $\overline{\mathbb{T}}^+$ to a continuous limit implies uniform convergence (see Proposition 3.6 *ii*) will ultimately be used to derive the uniform convergence stated in Theorem 1.9.

We now consider the decomposition of gradient boosting into positive and negative parts. The Jordan decomposition $\mu = \mu^+ - \mu^-$ of the signed measure $\mu \in \mathcal{M}_0$ naturally induces the following decomposition of tree functions: for $T \in \mathbb{T}$, we have

$$T = T^+ - T^- \quad \text{with } \mu_{T^+} = (\mu_T)^+ \text{ and } \mu_{T^-} = (\mu_T)^-.$$

The tree functions T^+, T^- are respectively associated with the positive and negative part of μ_T and hence belong to $\overline{\mathbb{T}}^+$. This decomposition can be applied to the softmax gradient tree $\tilde{T}(\cdot; F, \zeta)$ and we note

$$\tilde{T}(\cdot; F, \zeta) = \tilde{T}^+(\cdot; F, \zeta) - \tilde{T}^-(\cdot; F, \zeta).$$

The gradient boosting sequence $(\hat{F}_m^\lambda)_{m \geq 0}$ is then decomposed recursively as $\hat{F}_0^\lambda = \hat{F}_0^{\lambda+} - \hat{F}_0^{\lambda-}$ and

$$\begin{cases} \hat{F}_{m+1}^{\lambda+} &= \hat{F}_m^{\lambda+} + \lambda \tilde{T}^+(\cdot; \hat{F}_m^{\lambda+}, \zeta_{m+1}) \\ \hat{F}_{m+1}^{\lambda-} &= \hat{F}_m^{\lambda-} + \lambda \tilde{T}^-(\cdot; \hat{F}_m^{\lambda+}, \zeta_{m+1}) \end{cases}, \quad m \geq 0.$$

Clearly, we have $\hat{F}_m^\lambda = \hat{F}_m^{\lambda+} - \hat{F}_m^{\lambda-}$, and $\hat{F}_m^{\lambda+}, \hat{F}_m^{\lambda-} \in \mathbb{T}^+$ but it is not true in general that $\hat{F}_m^{\lambda+}$ and $\hat{F}_m^{\lambda-}$ are the positive and negative parts of \hat{F}_m^λ because the measures associated to the two components may share some common mass (at point 0) that vanishes when taking the difference. This slight abuse of notation should however cause no confusion. Finally, the infinitesimal gradient boosting operator can also be decomposed into a positive and negative part as $\mathcal{T} = \mathcal{T}^+ - \mathcal{T}^-$ with

$$\begin{cases} \mathcal{T}^+(F)(x) &= \mathbb{E}_\zeta[\tilde{T}^+(x; F, \zeta)] \\ \mathcal{T}^-(F)(x) &= \mathbb{E}_\zeta[\tilde{T}^-(x; F, \zeta)] \end{cases}, \quad F \in L_{\mathbf{x}}^2, x \in [0, 1]^p.$$

Note here that the operator are defined on $L_{\mathbf{x}}^2$ and not L^2 — this is required because we need the values $(F(x_i))_{1 \leq i \leq n}$ to be well-defined in order to define the softmax gradient tree $\tilde{T}(\cdot; F, \zeta)$. The measures associated with $\mathcal{T}^+(F)$ and $\mathcal{T}^-(F)$ may share some common mass so that the decomposition does not coincide with the Jordan decomposition.

The following theorem characterizes the convergence of $(\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-})_{t \geq 0}$ with deterministic limit given as the solution of a differential equation.

Theorem 4.4.

(i) *Given an initial condition $(F_0^+, F_0^-) \in L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2$, there is a unique solution $(F^+(t), F^-(t))_{t \geq 0}$ to the ODE system*

$$\begin{cases} \frac{d}{dt} F^+(t) = \mathcal{T}^+(F(t)) \\ \frac{d}{dt} F^-(t) = \mathcal{T}^-(F(t)) \end{cases} \quad \text{with } F(t) = F^+(t) - F^-(t), \quad t \geq 0, \quad (20)$$

started from (F_0^+, F_0^-) .

Furthermore, if $(F_0^+, F_0^-) \in \mathbb{T}^+ \times \mathbb{T}^+$, then $(F^+(t), F^-(t)) \in \mathbb{T}^+ \times \mathbb{T}^+$ for all $t \geq 0$ and $(F(t))_{t \geq 0}$ is equal to the solution of (15) started from $F_0 = F_0^+ - F_0^- \in \mathbb{T}$.

(ii) *The convergence in distribution*

$$(\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-})_{t \geq 0} \xrightarrow{d} (\hat{F}_t^+, \hat{F}_t^-)_{t \geq 0}, \quad \text{as } \lambda \rightarrow 0,$$

holds in the Skorokod space $\mathbb{D}([0, \infty), \overline{\mathbb{T}^+} \times \overline{\mathbb{T}^+})$ endowed with the J_1 -topology where $(\hat{F}_t^+, \hat{F}_t^-)_{t \geq 0}$ is equal to the solution of (20) started from $(\hat{F}_0^+, \hat{F}_0^-)$.

An important tool in the proof of Theorem 4.4 is a control of the training error

$$L_n(F) = \sum_{i=1}^n L(y_i, F(x_i)), \quad F \in \mathbb{B}.$$

The next proposition states that for small learning rates, the training error is almost surely non-increasing along the gradient boosting sequence $(\hat{F}_m^\lambda)_{m \geq 0}$. This in turn yields a control of the increments of the sequence because — as we will show in the course of the proof of Theorem 4.4 — the total variation of the gradient tree $\|\tilde{T}(\cdot; F, \zeta)\|_{\text{TV}}$ is uniformly bounded on the level set $\{F \in \mathbb{B} : L_n(F) \leq C\}$.

Proposition 4.5. *For all $T > 0$, there exists $\lambda_0 > 0$ such that $\lambda \in (0, \lambda_0]$ implies that*

$$t \mapsto L_n(\hat{F}_{\lfloor t/\lambda \rfloor}^\lambda) \quad \text{is a.s. non-increasing on } [0, T].$$

Note that in the case of regression $L(y, z) = \frac{1}{2}(y - z)^2$, Proposition 4.5 trivially holds for all $T > 0$ and $\lambda > 0$ because the tree values are obtained by line search so that the training error cannot increase, see Example 1.4 and Equation (13). In the general case however, the one-step approximation (9) does not ensure such a property but the proposition states that monotonicity still holds for small learning rates.

4.3 Properties of infinitesimal gradient boosting

We call infinitesimal gradient boosting the solution $(\hat{F}_t)_{t \geq 0}$ of the differential Equation (15) with initial condition (5). In the following, we consider some properties of infinitesimal gradient boosting, including the behavior of training error and residuals, the asymptotic behavior as $t \rightarrow +\infty$ and also the space-time regularity.

Proposition 4.6. *Infinitesimal gradient boosting satisfies the following properties:*

- (i) *the training error $L_n(\hat{F}_t) = \sum_{i=1}^n L(y_i, \hat{F}_t(x_i))$, $t \geq 0$, is non-increasing;*
- (ii) *the residuals $r_{t,i} = \frac{\partial L}{\partial z}(y_i, \hat{F}_t(x_i))$, $1 \leq i \leq n$, are centered, i.e.*

$$\bar{r}_t := \frac{1}{n} \sum_{i=1}^n r_{t,i} \equiv 0, \quad t \geq 0.$$

The first point is very natural since gradient boosting aims at minimizing the training error. The second point can be interpreted as follows: it is not possible to reduce the training error by adding a constant term to the model; indeed, a simple computation shows that $c \mapsto L_n(\hat{F}_t + c)$ has derivative $n\bar{r}_t = 0$ at $c = 0$ and, the function being convex, this corresponds to a minimum. As will be clear from the proof, this property is due to the initialization (5) and to the line search approximation (9).

We next consider the long time behavior of infinitesimal gradient boosting. A few more assumptions are required for our analysis.

Assumption 4.7.

(A3) *for all $y \in \mathbb{R}$, $\inf_z L(y, z) = 0$;*

(A4) *for all $y \in \mathbb{R}$ and all $R > 0$, $\sup_{z: |\frac{\partial L}{\partial z}(y, z)| \leq R} \frac{\partial^2 L}{\partial z^2}(y, z) < \infty$;*

(A5) there is $J \subset \llbracket 1, p \rrbracket$ with $|J| \leq d$ such that $(x_i^J)_{1 \leq i \leq n}$ are pairwise distinct, where $x_i^J = (x_i^j)_{j \in J} \in [0, 1]^J$.

Assumptions (A3) and (A4) state conditions on the loss function L . It is easily checked that the classical loss functions for regression and classification satisfy those conditions, see Examples 1.4–1.6. Furthermore, (A3) can always be assumed without loss of generality because shifting the loss function by its infimum does not affect the gradient boosting algorithm. On the other hand, Assumption (A5) states a condition on the training sample $(x_i)_{1 \leq i \leq n}$ and is discussed in Remark 4.9 below. It is used in the following proposition characterizing the critical points of the ODE (15).

Proposition 4.8. *Under Assumption (A5), $\mathcal{T}(F) = 0$ if and only if*

$$\frac{\partial L}{\partial z}(y_i, F(x_i)) = 0, \quad \text{for all } 1 \leq i \leq n.$$

This means that critical points of the ODE (15) are functions for which the residuals are all null.

Remark 4.9. We briefly comment upon Assumption (A5). It is trivially satisfied in the following two cases:

- the $(x_i)_{1 \leq i \leq n}$ are pairwise distinct and $d \geq p$ (consider $J = \llbracket 1, p \rrbracket$);
- there exists $j \in \llbracket 1, p \rrbracket$ such that the $(x_i^j)_i$ are pairwise distinct (consider $J = \{j\}$).

Next, we discuss a simple example where (A5) does not hold. Consider the problem of regression with $p = 2, d = 1$ and $(x_i, y_i)_{1 \leq i \leq 4}$ given by

$$\begin{cases} x_1 = (\frac{1}{3}, \frac{1}{3}), y_1 = 1, & x_2 = (\frac{2}{3}, \frac{1}{3}), y_2 = -1, \\ x_3 = (\frac{1}{3}, \frac{2}{3}), y_3 = -1, & x_4 = (\frac{2}{3}, \frac{2}{3}), y_4 = 1. \end{cases}$$

For this configuration, starting from $F \equiv 0$, any randomized tree with depth $d = 1$ is null because, for any vertical or horizontal split, the sum of residuals compensate in each region. Therefore, $\mathcal{T}(F) = 0$ but the residuals are nonzero. Similar examples can be built in higher dimension as long as $d < p$.

The analysis of critical point suggests that, in the long-time asymptotic, infinitesimal gradient boosting should converge to a critical point with null residuals. This is the subject of our next result.

Proposition 4.10. *Under Assumption 4.7, we have*

$$\begin{aligned} \frac{\partial L}{\partial z}(y_i, \hat{F}_t(x_i)) &\longrightarrow 0, \quad 1 \leq i \leq n, \\ \text{and } L_n(\hat{F}_t) &\longrightarrow 0 \end{aligned}$$

as $t \rightarrow \infty$.

In the specific case of regression or classification, we get under Assumption 4.7:

- $\hat{F}_t(x_i) \rightarrow y_i$ in the case of regression (Example 1.4);
- $\hat{F}_t(x_i) \rightarrow +\infty$ if $y_i = 1$ and $\hat{F}_t(x_i) \rightarrow -\infty$ otherwise in the case of classification (Examples 1.5 and 1.6).

This shows that infinitesimal gradient boosting is prone to overfitting as $t \rightarrow +\infty$ as it tries to match the training sample perfectly. This is a well-known feature of gradient boosting and early stopping is usually used to avoid overfitting and obtain good generalization properties, see Zhang and Yu (2005).

Finally, we study the regularity of infinitesimal gradient boosting with respect to space and time. To this aim, we define a function space \mathbb{W} and first a reference measure $\pi_0 \in \mathcal{M}_0$ as follows. For a splitting scheme ξ , we consider $\pi_\xi = \sum \delta_x \in \mathcal{M}_0$ the point measure with a Dirac mass at each vertex in $[0, 1]^p$ of the partition $(A_v(\xi))_{v \in \{0,1\}^d}$ into hypercubes; the measure π_0 is then defined as the intensity measure of the point process π_ξ under $P_0(d\xi)$, that is

$$\pi_0 = \int \pi_\xi P_0(d\xi) \in \mathcal{M}_0.$$

The function space $\mathbb{W} = \mathbb{W}([0, 1]^p, \mathbb{R})$ is defined by

$$\mathbb{W} = \left\{ T \in \mathbb{T} : \mu_T \ll \pi_0 \text{ and } \frac{d\mu_T}{d\pi_0} \in L^\infty([0, 1]^p, \pi_0) \right\}$$

and is endowed with the norm

$$\|T\|_{\mathbb{W}} = \left\| \frac{d\mu_T}{d\pi_0} \right\|_\infty, \quad T \in \mathbb{W}.$$

Clearly, $\mathbb{W} \subset \mathbb{T}$ with $\|T\|_{TV} \leq \|T\|_{\mathbb{W}}$, for all $T \in \mathbb{W}$. Furthermore, $(\mathbb{W}, \|\cdot\|_{\mathbb{W}})$ is a Banach space isomorphic to $L^\infty([0, 1]^p, \pi_0)$ via the map $T \mapsto d\mu_T/d\pi_0$.

Proposition 4.11. *Infinitesimal gradient boosting $(\hat{F}_t)_{t \geq 0}$ defines a smooth space in \mathbb{W} , i.e. $\hat{F}_t \in \mathbb{W}$ for all $t \geq 0$ and the mapping $t \mapsto \hat{F}_t$ is continuously differentiable. Furthermore, the same result holds for the positive and negative parts $(\hat{F}_t^+)_{t \geq 0}$ and $(\hat{F}_t^-)_{t \geq 0}$ defined jointly as the solution of (20).*

As a consequence, there exists a continuously differentiable mapping $t \mapsto \varphi_t \in L^\infty([0, 1]^p, \pi_0)$ such that

$$\hat{F}_t(x) = \int_{[0,x]} \varphi_t(y) \pi_0(dy), \quad x \in [0, 1]^p, t \geq 0.$$

Note that π_0 is absolutely continuous with respect to ν , because all possible vertices of a partition corresponding to a splitting scheme ξ with distribution P_0 have coordinates that are either null, or of the form $\prod_{i=1}^N U_i$, where N is random and the U_i are i.i.d. uniform in $[0, 1]$. As a consequence, Proposition 4.11 implies that $(t, x) \mapsto \hat{F}_t(x)$ is jointly continuous on $[0, \infty) \times [0, 1]^p$ and the same holds for $\hat{F}_t^+(x)$ and $\hat{F}_t^-(x)$.

Surprisingly, infinitesimal gradient boosting is regular not only in time but also in space. Recall that $(\hat{F}_t)_{t \geq 0}$ appears as the limit of gradient boosting $(\hat{F}_{[t/\lambda]}^\lambda)_{t \geq 0}$ which is highly discontinuous. The limit $\lambda \rightarrow 0$ as a regularizing effect with respect to time and space. The temporal regularizing effect is quite natural since the jump size λ vanishes in the limit. The spatial regularizing effect is due to the randomization of softmax gradient trees: due to split randomization, the discontinuities in the different gradient trees occur at different places and an averaging effect provides a regular function in the limit. This heuristic reasoning is made rigorous in Lemma 5.4 where we state that $\mathcal{T}(F) \in \mathbb{W}$ for all $F \in \mathbb{B}$, evidencing the spatial regularization effect of the infinitesimal gradient boosting operator.

5 Proofs

5.1 Proofs related to Section 2

Proof of Proposition 2.1. Let $A \subset ([1, p] \times (0, 1))^{\mathcal{T}_{d-1}}$ be measurable and consider the event $\{\xi \in A\}$. Modulo null sets, it can be decomposed into the disjoint union

$$\{\xi \in A\} = \bigcup_{\varphi} \{\xi \in A, \xi = (j_v^{\varphi(v)}, u_v^{\varphi(v)})_{v \in \mathcal{T}_{d-1}}\}$$

with the union taken over all selection maps $\varphi : \mathcal{T}_{d-1} \rightarrow [1, K]$. We deduce

$$\mathbb{P}(\xi \in A) = \sum_{\varphi} \mathbb{P}(\xi \in A, \xi = (j_v^{\varphi(v)}, u_v^{\varphi(v)})_{v \in \mathcal{T}_{d-1}} \mid \xi^1, \dots, \xi^K) \otimes_{k=1}^K P_0(d\xi^k).$$

By invariance of the product measure, all the terms in the sum are equal to the term corresponding to $\varphi \equiv 1$ and there are K^{2^d-1} such terms, with $2^d - 1$ the cardinal of \mathcal{T}_{d-1} . We deduce

$$\begin{aligned} \mathbb{P}(\xi \in A) &= K^{2^d-1} \int \mathbb{P}(\xi \in A, \xi = \xi^1 \mid \xi^1, \dots, \xi^K) \otimes_{k=1}^K P_0(d\xi^k) \\ &= K^{2^d-1} \int \mathbb{1}_{\{\xi^1 \in A\}} \prod_{v \in \mathcal{T}_{d-1}} \left(\text{softmax}_{\beta}(\Delta(j_v^k, u_v^k; A_v(\xi^1))_{1 \leq k \leq K}) \right)_1 \otimes_{k=1}^K P_0(d\xi^k) \\ &= \int_A \left(\int \prod_{v \in \mathcal{T}_{d-1}} \frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi^1)))}{K^{-1} \sum_{k=1}^K \exp(\beta \Delta(j_v^k, u_v^k; A_v(\xi^1)))} \otimes_{k=2}^K P_0(d\xi^k) \right) P_0(d\xi^1), \end{aligned}$$

where the second equality relies on Equation (18) and the third on the definition of the softmax function. This characterizes the distribution $P_{\beta, K}$ of ξ and proves that the Radon-Nykodym derivative is given by Equation (19). \square

Proof of Proposition 2.5. Consider input $\mathbf{r} = (r_i)_{1 \leq i \leq n}$ and $\mathbf{r}' = (r'_i)_{1 \leq i \leq n}$ and let $(x_i)_{1 \leq i \leq n}$ be fixed. We may replace without risk of confusion $(x_i, r_i)_{1 \leq i \leq n}$ and $(x_i, r'_i)_{1 \leq i \leq n}$ by \mathbf{r} and \mathbf{r}' respectively in the notation below. For a fixed (deterministic) splitting scheme $\xi = (j_v, u_v)_{v \in \mathcal{T}_{d-1}}$, we compare $T(\cdot; \mathbf{r}, \xi)$ and $T(\cdot; \mathbf{r}', \xi)$. Denoting by $(A_v)_{v \in \{0,1\}^d}$ the partition associated to ξ , we have

$$T(\cdot; \mathbf{r}, \xi) = \sum_{v \in \{0,1\}^d} r(A_v) \mathbb{1}_{A_v} \quad \text{and} \quad T(\cdot; \mathbf{r}', \xi) = \sum_{v \in \{0,1\}^d} r'(A_v) \mathbb{1}_{A_v},$$

with $r(A_v)$ (resp. $r'(A_v)$) the mean of the values r_i (resp. r'_i) with $x_i \in A_v$. Since

$$|r(A_v) - r'(A_v)| \leq \|\mathbf{r} - \mathbf{r}'\|_\infty,$$

we deduce

$$\|T(\cdot; \mathbf{r}, \xi) - T(\cdot; \mathbf{r}', \xi)\|_\infty \leq \|\mathbf{r} - \mathbf{r}'\|_\infty. \quad (21)$$

Next, we consider a random splitting scheme ξ . It should be stressed that the splitting scheme distribution depends on the input (respectively \mathbf{r} and \mathbf{r}') and we denote by $P_{\beta,K}$ and $P'_{\beta,K}$ the respective splitting scheme distributions. Proposition 2.1 implies

$$\bar{T}_{\beta,K,d}(x; \mathbf{r}) = \int T(x; \mathbf{r}, \xi) \frac{dP_{\beta,K}}{dP_0}(\xi) P_0(d\xi)$$

and similarly for $\bar{T}_{\beta,K,d}(x; \mathbf{r}')$ with \mathbf{r} and $dP_{\beta,K}$ replaced by \mathbf{r}' and $dP'_{\beta,K}$ respectively. We deduce

$$\begin{aligned} \left| \bar{T}_{\beta,K,d}(x; \mathbf{r}) - \bar{T}_{\beta,K,d}(x; \mathbf{r}') \right| &\leq \int |T(x; \mathbf{r}, \xi) - T(x; \mathbf{r}', \xi)| \frac{dP_{\beta,K}}{dP_0}(\xi) P_0(d\xi) \\ &\quad + \int |T(x; \mathbf{r}', \xi)| \left| \frac{dP_{\beta,K}}{dP_0}(\xi) - \frac{dP'_{\beta,K}}{dP_0}(\xi) \right| P_0(d\xi). \end{aligned} \quad (22)$$

The first term is bounded from above by $\|\mathbf{r} - \mathbf{r}'\|_\infty$ thanks to Equation (21). For the second term, we use $|T(x; \mathbf{r}', \xi)| \leq \|\mathbf{r}'\|_\infty$ and the following Lipschitz property of the Radon-Nykodym derivative, justified below,

$$\sup_\xi \left| \frac{dP_{\beta,K}}{dP_0}(\xi) - \frac{dP'_{\beta,K}}{dP_0}(\xi) \right| \leq CR \|\mathbf{r} - \mathbf{r}'\|_\infty \quad (23)$$

for $\max(\|\mathbf{r}\|_\infty, \|\mathbf{r}'\|_\infty) \leq R$ and $C = \beta(2^d - 1)K^{2^d - 1}$. With these bound, Equation (22) implies

$$\left\| \bar{T}_{\beta,K,d}(\cdot; \mathbf{r}) - \bar{T}_{\beta,K,d}(\cdot; \mathbf{r}') \right\|_\infty \leq (1 + CR^2) \|\mathbf{r} - \mathbf{r}'\|_\infty,$$

proving that $\mathbf{r} \mapsto \bar{T}_{\beta,K,d}(\cdot; \mathbf{r})$ is locally Lipschitz.

We finally prove Equation (23) using the explicit formula (19) for the Radon-Nykodym derivative. In definition (16), the score $\Delta(j, u; A)$ implicitly depends on the input \mathbf{r} so for clarity we write below $\Delta(j, u; A, \mathbf{r})$. Simple computations yield

$$|\Delta(j, u; A, \mathbf{r}) - \Delta(j, u; A, \mathbf{r}')| \leq 2R \|\mathbf{r} - \mathbf{r}'\|_\infty$$

for $\max(\|\mathbf{r}\|_\infty, \|\mathbf{r}'\|_\infty) \leq R$ and all (j, u) and A . Furthermore, it is elementary to see that the softmax function $z \mapsto e^{z_1} / \sum_{k=1}^K e^{z_k}$ is 1/2-Lipschitz for the uniform norm on \mathbb{R}^K . Therefore, for fixed $\xi, \xi_1, \xi_2, \dots, \xi_K$ and $v \in \mathcal{T}_{d-1}$, the different factors in Equation (19) satisfy

$$\left| \frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi), \mathbf{r}))}{\sum_{k=1}^K \exp(\beta \Delta(j_v^k, u_v^k; A_v(\xi), \mathbf{r}))} - \frac{\exp(\beta \Delta(j_v^1, u_v^1; A_v(\xi), \mathbf{r}'))}{\sum_{k=1}^K \exp(\beta \Delta(j_v^k, u_v^k; A_v(\xi), \mathbf{r}'))} \right| \leq \beta R \|\mathbf{r} - \mathbf{r}'\|_\infty.$$

Using the inequality $|\Pi_{i \in I} a_i - \Pi_{i \in I} b_i| \leq \sum_{i \in I} |a_i - b_i|$ for finite families $(a_i), (b_i) \in [-1, 1]^I$ with $I = \mathcal{T}_{d-1}$ of cardinal $2^d - 1$ and integrating with respect to $\delta_\xi(d\xi^1) \otimes_{k=2}^K P_0(d\xi^k)$, we deduce Equation (23). \square

5.2 Proofs related to Section 3

Proof of Proposition 3.3. Consider $a, b \in [0, 1]^p$ with $0 \leq a^i < b^i \leq 1$ for all $1 \leq i \leq p$. Using the notation of Neuhaus (1971), we consider the map $S = \mathbb{1}_{[a,b]}$, where

$$[a, b] := I_1 \times I_2 \times \cdots \times I_p, \quad \text{with} \quad I_i = \begin{cases} [a^i, b^i] & \text{if } b^i < 1 \\ [a^i, b^i] & \text{if } b^i = 1. \end{cases}$$

Let us show that $S \in \mathbb{T}$ and characterize the measure μ_S . For each vertex $\varepsilon = (\varepsilon_1, \dots, \varepsilon_p) \in \{0, 1\}^p$, we denote by $|\varepsilon| = \sum_{i=1}^p \varepsilon_i$ the number of positive components. The vertices of the hypercube $[a, b]$ are $a + \varepsilon \cdot (b - a)$ where \cdot denotes coordinate-wise multiplication. Then the signed measure μ_S

$$\mu_S = \sum_{\varepsilon \in \{0,1\}^p} (-1)^{|\varepsilon|} \mathbb{1}_{[0,1]^p}(a + \varepsilon \cdot (b - a)) \delta_{a + \varepsilon \cdot (b - a)} \quad (24)$$

belongs to \mathcal{M}_0 and is such that $\mu_S([0, x]) = S(x)$ for all $x \in [0, 1]^p$. This proves that $\mathbb{1}_{[a,b]} = S \in \mathbb{T}$. Furthermore, we see from (24) that $\|\mathbb{1}_{[a,b]}\|_{\text{TV}} = 2^q \leq 2^p$, with q the number of coordinates i such that $b^i < 1$.

Now recall the recursive construction of the regions $(A_v(\xi))_{v \in \mathcal{F}_d}$ from the splitting scheme ξ , described in Section 2.2. For each leaf $v \in \{0, 1\}^d$, the region $A_v(\xi)$ is defined as the intersection of d subsets of $[0, 1]^p$ of the form $[a, b]$, where there is a unique $i \in \llbracket 1, p \rrbracket$ such that $a^i > 0$ or $b^i < 1$. Therefore A_v is necessarily of the form $[a, b]$, where there are $q \leq \min(p, d)$ coordinates $i \in \llbracket 1, p \rrbracket$ such that $a^i > 0$ or $b^i < 1$. From the discussion above, we deduce the bound

$$\|T\|_{\text{TV}} \leq \sum_{v \in \{0,1\}^d} |\tilde{r}_v| \|\mathbb{1}_{A_v(\xi)}\|_{\text{TV}} \leq 2^{d+\min(p,d)} \|T\|_{\infty}.$$

In the case $d < p$, each region $A_v(\xi)$ is of the form $[a, b]$ where the set J of coordinates $j \in \llbracket 1, p \rrbracket$ such that $a^j > 0$ or $b^j < 1$ has cardinal $q \leq d < p$. If $\varepsilon \in \{0, 1\}^p$ is such that $a + \varepsilon \cdot (b - a) \in [0, 1]^p$ and $j \notin J$, then $(a + \varepsilon \cdot (b - a))^j = a^j = 0$. This shows that the signed measure μ_S associated with $S = \mathbb{1}_{[a,b]}$ and defined in (24) has its support in

$$\mathcal{F}_J = \{x \in [0, 1]^p : x^j = 0 \text{ for all } j \notin J\} \subset \mathcal{F}_d.$$

Since T is a linear combination of such functions, the proof is complete. \square

Proof of Proposition 3.5. (i) The density results from standard approximation arguments, for instance, of continuous functions, which are dense in L^2 , by step functions.

(ii) Since L^2 is a metric space, it is sufficient to consider a sequence $T_1, T_2, \dots \in \mathbb{T}$ that is bounded in total variation, and show there exists a convergent subsequence. Because $(\mu_{T_n})_{n \geq 1}$ is a sequence of signed measures that is supported on the compact space $[0, 1]^p$ and bounded in total variation, we can extract a weakly convergent subsequence. In fact, we can apply this argument to the sequences $(\mu_{T_n}^J)_{n \geq 1}$ and assume that $\mu_{T_n}^J \xrightarrow{w} \mu^J$ along some subsequence, jointly for all $J \subset \llbracket 1, p \rrbracket$, where the μ^J are finite signed measures supported on $\overline{\mathcal{F}}_J \subset [0, 1]^p$ — note that considering closures is necessary.

We aim at proving ν -a.e. convergence of the sequence of functions $(T_n)_{n \geq 1}$ — this implies L^2 convergence by the dominated convergence theorem because ν is finite and the T_n are bounded. More precisely, we prove that

$$T_n(x) \xrightarrow{n \rightarrow \infty} T(x) := \sum_{J \subset \llbracket 1, p \rrbracket} \mathbb{1}_{\mathcal{F}_{\geq J}}(x) \mu^J([0, x]) \quad \nu\text{-a.e.}, \quad (25)$$

where $\mathcal{F}_{\geq J} = \{x \in [0, 1]^p : \forall j \in J, x^j > 0\}$. Note that of course the limit is not necessarily in \mathbb{T} .

First consider a point $a \in [0, 1]^p$ and some $J \subset \llbracket 1, p \rrbracket$. Viewing $\mu_{T_n}^J$ and μ^J as measures on the space $\overline{\mathcal{F}}_J$, by the Portmanteau theorem, the weak convergence $\mu_{T_n}^J \xrightarrow{w} \mu^J$ implies

$$\mu_{T_n}^J([0, a]) \xrightarrow{n \rightarrow \infty} \mu^J([0, a])$$

if $\mu^J(\partial([0, a^J])) = 0$, with $a^J = (a^j \mathbb{1}_{j \in J})_j$ and where the boundary in the previous expression is the topological boundary within the space $\overline{\mathcal{F}}_J$ — *not within* $[0, 1]^p$. Notice that

- if $a \notin \mathcal{F}_{\geq J}$, that is if there exists $j \in J$ such that $a^j = 0$, then $[0, a] \cap \mathcal{F}_J = \emptyset$, so that $\mu_{T_n}^J([0, a]) = 0$ for all $n \geq 1$;
- we have the inclusion

$$\partial([0, a^J]) \subset \bigcup_{\substack{j \in J \\ 0 \leq a^j < 1}} \{x \in [0, 1]^p : x^j = a^j\}.$$

Now let us define

$$A = \{0, 1\} \cup \left\{ t \in [0, 1] : \sum_{J \subset \llbracket 1, p \rrbracket} \sum_{j=1}^p \mu^J(\{x \in [0, 1]^p : x^j = t\}) = 0 \right\}.$$

The set A has a countable complement in $[0, 1]$ and therefore has Lebesgue measure 1. Since $0, 1 \in A$, we have $\nu(A^p) = 1$ as well. Finally consider $a \in A^p$ and let us show that the sequence $(T_n(a))_{n \geq 1}$ converges to $T(a)$ defined in (25). It is sufficient to show that $\mu_{T_n}^J([0, a])$ converges to $\mathbb{1}_{\mathcal{F}_{\geq J}}(a) \mu^J([0, a])$ for any $J \subset \llbracket 1, p \rrbracket$. From the discussion above, this is obvious whenever $a \notin \mathcal{F}_{\geq J}$. If this is not the case, we have

$$\partial([0, a^J]) \subset \bigcup_{\substack{j \in J \\ 0 < a^j < 1}} \{x \in [0, 1]^p : x^j = a^j\},$$

which is a null μ^J -measure set by construction of A , implying

$$\mu_{T_n}^J([0, a]) \longrightarrow \mu^J([0, a]).$$

As this is true for all $J \subset \llbracket 1, p \rrbracket$, we conclude that

$$T_n(a) \xrightarrow{n \rightarrow \infty} T(a), \quad \text{for all } a \in A^p,$$

therefore $T_n \rightarrow T$ in L^2 , completing the proof.

(iii) The sequence $(T_n)_{n \geq 1}$ being bounded in total variation, it is relatively compact in L^2 by (ii). Using the identification of the possible adherence points (25) of this sequence in the proof of the previous point, (iii) is easily deduced. \square

Proof of Proposition 3.6. (i) To show that $\overline{\mathbb{T}}^+$ is a proper metric space, it is sufficient to show that any bounded sequence $T_1, T_2, \dots \in \mathbb{T}^+$ is relatively compact in L^2 .

The fact that the sequence $(T_n)_{n \geq 1}$ is bounded in L^2 implies that the real sequence $(T_n(1))_{n \geq 1}$ is bounded because the reference measure ν includes a Dirac mass at 1. Since $T_n \in \mathbb{T}^+$, the equality $T_n(1) = \|T_n\|_{TV}$ holds. Therefore the sequence (T_n) is bounded in total variation so that relative compactness is deduced from point (ii) in Proposition 3.5.

(ii) This is a straightforward consequence of the multidimensional version of Pólya's uniform convergence theorem for distribution functions — for a general form of this convergence theorem, see Billingsley and Topsøe (1967, Theorem 2). \square

5.3 Proofs related to Section 4

5.3.1 Proofs related to Section 4.1 and proof of Theorem 1.11

Proof of Proposition 4.1. From Definition 1.8, the function $\mathcal{T}(F)$ is defined point-wise by $\mathcal{T}(F)(x) = \mathbb{E}_\zeta[\tilde{T}(x; F, \zeta)]$ for all $x \in [0, 1]^p$. By Proposition 3.3, $\tilde{T}(\cdot; F, \zeta) \in \mathbb{T}$ for all ζ . Note that it is not straightforward to define $\mathcal{T}(F)$ as an expectation in the Banach space $(\mathbb{T}, \|\cdot\|_{TV})$ because of separability and measurability issues. Alternatively, we can take the expectation of the random measure $\mu_{\tilde{T}(\cdot; F, \zeta)}$ and define $\mu = \mathbb{E}_\zeta[\mu_{\tilde{T}(\cdot; F, \zeta)}] \in \mathcal{M}_0$. It is then straightforward to see that $\mathcal{T}(F)(x) = \mu([0, x])$ for all $x \in [0, 1]^p$ so that $\mathcal{T}(F) \in \mathbb{T}$ with

$$\mu_{\mathcal{T}(F)} = \mathbb{E}_\zeta[\mu_{\tilde{T}(\cdot; F, \zeta)}]. \quad (26)$$

We next prove that $\mathcal{T} : \mathbb{B} \rightarrow \mathbb{T}$ is locally Lipschitz. Assumption (A1) plays here a crucial role because it implies the following property: writing $\mathbf{z} = (F(x_i))_{1 \leq i \leq n} \in \mathbb{R}^n$, $\mathbf{r} = (r_i)_{1 \leq i \leq n}$ with $r_i = \frac{\partial L}{\partial z}(x_i, F(x_i))$ and

$$\tilde{r}(A) = - \frac{\sum_{i=1}^n \frac{\partial L}{\partial z}(y_i, F(x_i)) \mathbb{1}_{\{x_i \in A\}}}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{\{x_i \in A\}}}, \quad A \subset [0, 1]^p,$$

the maps $\mathbf{z} \mapsto \mathbf{r}$ and $\mathbf{z} \mapsto \tilde{r}(A)$ are locally Lipschitz for the uniform norm on \mathbb{R}^n . They are indeed continuously differentiable in \mathbf{z} so that their derivatives are locally bounded. Furthermore, because only finitely many different functions $\mathbf{z} \mapsto \tilde{r}(A)$ can be generated for different regions A , the Lipschitz constant can be assumed independent of A . Finally, since $F \mapsto \mathbf{z}$ is linear with operator norm 1, the maps $F \mapsto \mathbf{r}$ and $F \mapsto \tilde{r}(A)$ are locally Lipschitz functions on \mathbb{B} .

As in Remark 2.2, we shall consider that the splitting scheme ξ depends on the auxiliary randomness ζ and let $\tilde{T}(\cdot; F, \xi)$ denote the gradient tree with splitting scheme ξ . Let us recall that the distribution $P_{\beta, K}$ of ξ depends (implicitly) on the vector of residuals \mathbf{r} and that the splitting scheme ξ induces a partition $(A_v)_{v \in \{0, 1\}^d}$. Then the leaf values of the gradient trees are given by $(\tilde{r}(A_v))_{v \in \{-1, 1\}^d}$ — not by $(r(A_v))_{v \in \{-1, 1\}^d}$ as in the case of regression trees. The gradient tree can thus be written

$$\tilde{T}(\cdot; F, \xi) = \sum_{v \in \{0, 1\}^d} \tilde{r}(A_v) \mathbb{1}_{A_v}.$$

As in the proof of Proposition 2.5, we first analyze the case of a fixed splitting scheme ξ inducing a fixed partition $(A_v)_{v \in \{0,1\}^d}$. Since the maps $F \mapsto \tilde{r}(A_v)$, $v \in \{0,1\}^d$, are locally Lipschitz, there exists, for any $R \geq 0$, a constant $C \geq 0$ that does not depend on the splitting scheme ξ and such that

$$\|\tilde{T}(\cdot; F, \xi) - \tilde{T}(\cdot; F', \xi)\|_\infty \leq C\|F - F'\|_\infty,$$

for all $F, F' \in \mathbb{B}$ with $\max(\|F\|_\infty, \|F'\|_\infty) \leq R$. Since ξ is induced by a d -depth splitting scheme, Proposition 3.3 implies

$$\begin{aligned} \|\mu_{\tilde{T}(\cdot; F, \xi)} - \mu_{\tilde{T}(\cdot; F', \xi)}\|_{\text{TV}} &\leq 4^d \|\tilde{T}(\cdot; F, \xi) - \tilde{T}(\cdot; F', \xi)\|_\infty \\ &\leq 4^d C \|F - F'\|_\infty. \end{aligned} \quad (27)$$

Next, we compare $\mathcal{T}(F)$ and $\mathcal{T}(F')$ for $F, F' \in \mathbb{B}$. With similar notation as in the proof of Proposition 2.5, we denote by $\mathbf{r} = (r_i)_{1 \leq i \leq n}$ (resp. $\mathbf{r}' = (r'_i)_{1 \leq i \leq n}$) the residuals and by $P_{\beta, K}$ (resp. $P'_{\beta, K}$) the splitting scheme distribution associated with F (resp. F'). Equation (26) can be rewritten as

$$\mu_{\mathcal{T}(F)} = \int \mu_{\tilde{T}(\cdot; F, \xi)} \frac{dP_{\beta, K}}{dP_0} dP_0.$$

The same equation holds for $\mu_{\mathcal{T}(F')}$ with F replaced by F' and $P_{\beta, K}$ by $P'_{\beta, K}$. We deduce, as in Equation (22),

$$\begin{aligned} \|\mathcal{T}(F) - \mathcal{T}(F')\|_{\text{TV}} &= \|\mu_{\mathcal{T}(F)} - \mu_{\mathcal{T}(F')}\|_{\text{TV}} \\ &\leq \int \|\mu_{\tilde{T}(\cdot; F, \xi)} - \mu_{\tilde{T}(\cdot; F', \xi)}\|_{\text{TV}} \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) \\ &\quad + \int \|\mu_{\tilde{T}(\cdot; F', \xi)}\|_{\text{TV}} \left| \frac{dP_{\beta, K}}{dP_0}(\xi) - \frac{dP'_{\beta, K}}{dP_0}(\xi) \right| P_0(d\xi). \end{aligned}$$

The end of the proof is then similar to the end of the proof of Proposition 2.5: the first term is bounded from above by Equation (27) and the second one by Equation (23) together with the fact that the vector of residuals \mathbf{r} is locally Lipschitz in F . \square

The following lemma is a simple consequence of Assumption (A2) and is crucial for controlling the total variation of gradient trees, which is needed in the proof of Theorem 4.3.

Lemma 5.1. *Under Assumption 1.8, for $C > 0$, define*

$$M(C) = \sup_{(y, z): L(y, z) \leq C} \left| \frac{\partial L}{\partial z}(y, z) / \frac{\partial^2 L}{\partial z^2}(y, z) \right| < \infty.$$

Let $F \in \mathbb{T}$ be such that $L_n(F) = \sum_{i=1}^n L(y_i, F(x_i)) \leq C$. Then

$$\|\tilde{T}(\cdot; F, \zeta)\|_\infty \leq M(C), \quad \|\tilde{T}(\cdot; F, \zeta)\|_{\text{TV}} \leq 4^d M(C), \quad (28)$$

and consequently $\|\mathcal{T}(F)\|_{\text{TV}} \leq 4^d M(C)$.

Proof. Because the loss function is non negative, the inequality $L_n(F) \leq C$ yields $L(y_i, F(x_i)) \leq C$ for all $i = 1, \dots, n$ and Assumption (A2) implies

$$\left| \frac{\partial L(y_i, F(x_i))}{\partial^2 L(y_i, F(x_i))} \right| \leq M(C), \quad i = 1, \dots, n,$$

where we use the short notation ∂L and $\partial^2 L$ for the first and second order partial derivatives of L with respect to the second variable. Let $(A_v)_{v \in \{0,1\}^d}$ be the leaves of $\tilde{T}(\cdot; F, \xi)$ so that

$$\tilde{T}(\cdot; F, \xi) = \sum_{v \in \{0,1\}^d} \tilde{r}(A_v) \mathbb{1}_{A_v}$$

with

$$\tilde{r}(A_j) = - \frac{\sum_{i: x_i \in A_j} \partial L(y_i, F(x_i))}{\sum_{i: x_i \in A_j} \partial^2 L(y_i, F(x_i))}.$$

We finally observe that

$$|\tilde{r}(A_v)| \leq \frac{\sum_{i: x_i \in A_v} \partial^2 L(y_i, F(x_i)) \left| \frac{\partial L(y_i, F(x_i))}{\partial^2 L(y_i, F(x_i))} \right|}{\sum_{i: x_i \in A_v} \partial^2 L(y_i, F(x_i))} \leq M(C)$$

because we recognize a weighted average of terms bounded from above by $M(C)$. This implies the bound for the uniform norm in Equation (28), and the bound in total variation follows from Proposition 3.3. The last bound is deduced from (28) and the fact that we have $\|\mathcal{T}(F)\|_{\text{TV}} \leq \mathbb{E}[\|\tilde{T}(\cdot; F, \zeta)\|_{\text{TV}}]$. \square

Proof of Theorem 4.3. Because the operator \mathcal{T} is Lipschitz as proven in Proposition 4.1, the Cauchy–Lipschitz theorem implies the *local* existence and uniqueness of solutions for Equation (15). It follows that, with the initial condition $F(0) = F_0$, there exists a unique local solution of (15) and that this solution can be extended uniquely into a maximal solution, still noted F , defined on a maximal interval $[0, t^*)$. We need to prove that $t^* = +\infty$. For the sake of clarity, we write F_t for the solution at time t and $F_t(x)$ for its evaluation at x . We first prove that the function

$$t \mapsto L_n(F_t) = \sum_{i=1}^n L(y_i, F_t(x_i))$$

is non-increasing on $[0, t^*)$. Indeed, its derivative is equal to

$$\begin{aligned} \frac{d}{dt} L_n(F_t) &= \sum_{i=1}^n \left(\frac{d}{dt} F_t(x_i) \right) \frac{\partial L}{\partial z}(y_i, F_t(x_i)) \\ &= \sum_{i=1}^n \mathcal{T}(F_t)(x_i) \frac{\partial L}{\partial z}(y_i, F_t(x_i)) \\ &= -\mathbb{E}_{\zeta} \left[\sum_{v \in \{0,1\}^d} \frac{\left(\sum_{i=1}^n \frac{\partial L}{\partial z}(y_i, F_t(x_i)) \mathbb{1}_{A_v}(x_i) \right)^2}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F_t(x_i)) \mathbb{1}_{A_v}(x_i)} \right] \\ &\leq 0. \end{aligned} \tag{29}$$

The second equality uses the fact that $(F_t)_{t \in [0, t^]}$ is a solution of (15). The third equality uses Definitions 1.2 and 1.10 where the partition $(A_v)_{v \in \{0,1\}^d}$ depends on the external randomness ζ . The derivative being non-positive, we have proved that $t \mapsto L_n(F_t)$ is non-increasing on $[0, t^]$.

Consider the level set

$$\Lambda := \{G \in \mathbb{T} : L_n(G) \leq L_n(F_0)\}.$$

We have just proved that $t \mapsto L_n(F_t)$ is non-increasing so that $F_t \in \Lambda$ for all $t \in [0, t^]$. Therefore, by Lemma 5.1, there exists a constant $C > 0$ such that $\|\mathcal{T}(F_t)\|_{\text{TV}} \leq C$ for all $t \in [0, t^]$. We deduce that for all $0 \leq t \leq s < t^$,

$$\|F_s - F_t\|_{\text{TV}} = \left\| \int_t^s \mathcal{T}(F_u) \, du \right\|_{\text{TV}} \leq C(s - t).$$

If $t^$ were finite, this would imply that (F_t) is Cauchy as $t \rightarrow t^$, contradicting the maximality assumption of the solution $(F(t))_{t \in [0, t^]}$. We deduce that $t^ = +\infty$ and that the total variation norm of F_t increases at most linearly. \square

Proof of Theorem 1.11 (i). Using the fact that \mathcal{T} is a locally Lipschitz operator of $(\mathbb{B}, \|\cdot\|_\infty)$, as mentioned in Remark 4.2, the proof is the same as above. \square

5.3.2 Proofs related to Section 4.2 and proof of Theorem 1.9

The following Lemma is crucial for the proof of Proposition 4.5 below.

Lemma 5.2. *Let L satisfy Assumption 1.8 and consider K_1 (resp. K_2) a finite (resp. compact) subset of \mathbb{R} . Then there exists $\lambda_0 > 0$ such that for all $\lambda \in (0, \lambda_0]$, $k \geq 1$, $y_1, \dots, y_k \in K_1$ and $z_1, \dots, z_k \in K_2$, the following inequality holds*

$$\sum_{i=1}^k L(y_i, z_i + \lambda z) \leq \sum_{i=1}^k L(y_i, z_i), \quad \text{with } z = -\frac{\sum_{i=1}^k \frac{\partial L}{\partial z}(y_i, z_i)}{\sum_{i=1}^k \frac{\partial^2 L}{\partial z^2}(y_i, z_i)}.$$

Proof. Fix $y_1, \dots, y_k \in K_1$ and $z_1, \dots, z_k \in K_2$ and define

$$h(u) = \sum_{i=1}^k L(y_i, z_i + u), \quad u \in \mathbb{R}.$$

By Taylor's approximation, we have

$$h(u) - h(0) = uh'(0) + \frac{u^2}{2} h''(\theta u),$$

for some $\theta \in [0, 1]$. Therefore, for $z = -h'(0)/h''(0)$,

$$\begin{aligned} h(\lambda z) - h(0) &= -\lambda \frac{h'(0)^2}{h''(0)} + \frac{\lambda^2}{2} \frac{h'(0)^2}{h''(0)^2} h''(\theta \lambda z) \\ &= -\lambda \frac{h'(0)^2}{h''(0)} \left(1 - \frac{\lambda}{2} \frac{h''(\theta \lambda z)}{h''(0)} \right), \end{aligned}$$

for some $\theta \in [0, 1]$. We need to prove that uniformly in $k \geq 1$, $y_1, \dots, y_k \in K_1$, $z_1, \dots, z_k \in K_2$ and $\theta \in [0, 1]$, the last expression is non-positive for λ small enough, which amounts to showing that

$$\lambda \leq 2 \frac{h''(\theta\lambda z)}{h(0)} = 2 \frac{\sum_{i=1}^k \frac{\partial^2 L}{\partial z^2} \left(y_i, z_i - \theta\lambda \frac{\sum_i \frac{\partial L}{\partial z}(y_i, z_i)}{\sum_i \frac{\partial^2 L}{\partial z^2}(y_i, z_i)} \right)}{\sum_{i=1}^k \frac{\partial^2 L}{\partial z^2}(y_i, z_i)}. \quad (30)$$

Using the fact that $L(y, z)$ is C^2 in z with $\frac{\partial^2 L}{\partial z^2}(y, z) > 0$, first let

$$M := \max_{y \in K_1, z \in K_2} \left| \frac{\frac{\partial L}{\partial z}(y, z)}{\frac{\partial^2 L}{\partial z^2}(y, z)} \right| < \infty$$

and

$$\lambda_0 := \min_{y \in K_1, z \in K_2, \theta \in [-1, 1]} 2 \frac{\frac{\partial^2 L}{\partial z^2}(y, z + \theta M)}{\frac{\partial^2 L}{\partial z^2}(y, z)} > 0.$$

Clearly, for all $k \geq 1$, $y_1, \dots, y_k \in K_1$, $z_1, \dots, z_k \in K_2$ and $\lambda \leq \lambda_0 \wedge 1$, the inequality (30) is satisfied, concluding the proof. \square

Proof of Proposition 4.5. The proof relies on Lemma 5.2. Let K be a compact set containing the interval $[F_0 - \delta, F_0 + \delta]$ with F_0 the initial value of the boosting procedure and $\delta > 0$. Let λ_0 as in Lemma 5.2 and $\lambda \in (0, \lambda_0]$. We prove that $\|\hat{F}_m^\lambda - F_0\|_\infty \leq \delta$ implies $L_n(\hat{F}_{m+1}^\lambda) \leq L_n(\hat{F}_m^\lambda)$. Consider a leaf of the tree $T(\cdot; \hat{F}_m^\lambda, \zeta_{m+1})$ that contains the values $(x_i)_{i \in I}$ for some $I \subset \llbracket 1, n \rrbracket$. Applying Lemma 5.2 with $k = \text{card}(I)$ and $z_i = \hat{F}_m^\lambda(x_i) \in K$ for $i \in I$, we obtain

$$\sum_{i \in I} L(y_i, \hat{F}_{m+1}^\lambda(x_i)) \leq \sum_{i \in I} L(y_i, \hat{F}_m^\lambda(x_i)).$$

In words, the updated model reduces the error within each leaf. Summing over all leaves, we get $L_n(\hat{F}_{m+1}^\lambda) \leq L_n(\hat{F}_m^\lambda)$ so that the updated model reduces the training error. This shows that the training error is non-increasing as long as $(\hat{F}_m^\lambda)_{m \geq 0}$ remains in the ball $\Delta = \{F : \|F - F_0\|_\infty \leq \delta\}$.

We next evaluate the time needed to exit Δ . Let $C = L_n(F_0)$ be the initial value of the training error. By Lemma 5.1, $L_n(\hat{F}_m^\lambda) \leq C$ implies

$$\|T(\cdot; \hat{F}_m^\lambda, \zeta_{m+1})\|_\infty \leq M := M(C) \quad \text{a.s.},$$

whence we deduce

$$\|\hat{F}_{m+1}^\lambda - \hat{F}_m^\lambda\|_\infty = \lambda \|T(\cdot; \hat{F}_m^\lambda, \zeta_{m+1})\|_\infty \leq \lambda M \quad \text{a.s.}$$

As long as $(\hat{F}_m^\lambda)_{m \geq 0}$ remains in Δ , the training error is non-increasing and hence lower than its initial value $L_n(F_0) = C$ and this implies $\|F_{m+1}^\lambda - F_m^\lambda\|_\infty \leq \lambda M$. To exit Δ , the chain must travel the distance δ at speed less than λM so that at least $\delta/(\lambda M)$ iterations are needed. Renormalizing time, we get that $t \mapsto L_n(\hat{F}_{\lfloor t/\lambda \rfloor}^\lambda)$ does not exit Δ and is non-increasing on $[0, T]$ with $T = \delta/M$. Since $\delta > 0$ is arbitrary and M depends only on $C = L_n(F_0)$, we conclude the proof by taking $\delta = MT$ and choosing λ_0 accordingly. \square

Proof of Theorem 4.4 (i).

The proof of Proposition 4.1 is easily adapted to prove that \mathcal{T}^+ and \mathcal{T}^- are locally Lipschitz from $(\mathbb{B}, \|\cdot\|_\infty)$ into $(\mathbb{T}, \|\cdot\|_{\text{TV}})$. Indeed, with exactly the same proof, Equation (27) can be modified into

$$\begin{aligned} \|\tilde{T}^+(\cdot; F, \xi) - \tilde{T}^+(\cdot; F', \xi)\|_{\text{TV}} &= \|\mu_{\tilde{T}^+(\cdot; F, \xi)}^\pm - \mu_{\tilde{T}^+(\cdot; F', \xi)}^\pm\|_{\text{TV}} \\ &\leq 4^d C \|F - F'\|_\infty, \end{aligned}$$

and similarly for the negative part $\tilde{T}^-(\cdot; F, \xi) - \tilde{T}^-(\cdot; F', \xi)$.

Because a function $F \in L_{\mathbf{x}}^2$ is well defined at $(x_i)_{1 \leq i \leq n}$ and because the softmax gradient trees $\tilde{T}^\pm(\cdot; F, \xi)$ depends on F only through the values $(F(x_i))_{1 \leq i \leq n}$, the operators $\mathcal{T}^\pm, \mathcal{T}^\pm : L_{\mathbf{x}}^2 \rightarrow L_{\mathbf{x}}^2$ are well defined and also locally Lipschitz as

$$\begin{aligned} \|\mathcal{T}^\pm(F) - \mathcal{T}^\pm(F')\|_{L_{\mathbf{x}}^2} &\leq \|\mathcal{T}^\pm(F) - \mathcal{T}^\pm(F')\|_{\text{TV}} \\ &\leq C \max_{1 \leq i \leq n} |F(x_i) - F'(x_i)| \\ &\leq nC \|F - F'\|_{L_{\mathbf{x}}^2}. \end{aligned}$$

The fact that \mathcal{T}^\pm are locally Lipschitz operators implies the existence and uniqueness of *local solutions* to the ODE (20). Proving that the local solution started at (F_0^+, F_0^-) at time $t = 0$ can be extended uniquely into a maximal solution on $[0, \infty)$ is done exactly as in the proof of Theorem 4.3 and the existence and uniqueness of *global solutions* follows.

When $(F_0^+, F_0^-) \in \mathbb{T}^+ \times \mathbb{T}^+$, the integral representation

$$F^\pm(t) = F_0^\pm + \int_0^t \mathcal{T}^\pm(F(s)) ds, \quad t \geq 0,$$

implies that $F^\pm(t) \in \mathbb{T}^+$ because $\mathcal{T}^\pm(F(s)) \in \mathbb{T}^+$ for all $s \in [0, t]$. Considering the difference between the positive and negative components, we get

$$F(t) = F(0) + \int_0^t \mathcal{T}(F(s)) ds, \quad t \geq 0,$$

proving that $(F(t))_{t \geq 0}$ is the solution to ODE (15) started at F_0 . \square

Proof of Theorem 4.4 (ii). Our strategy is the following:

- Step 1 — tightness. We consider a sequence $\lambda = \lambda(N) \rightarrow 0$ and prove that the sequence of processes $(\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-})_{t \geq 0}$ is tight in $\mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$. We use here the fact that the processes take their values in $\bar{\mathbb{T}}^+ \subset L_{\mathbf{x}}^2$ which is a proper metric space, i.e. in which bounded sets are relatively compact.
- Step 2 — identification of the limit. We show that the weak convergence

$$(\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-})_{t \geq 0} \xrightarrow{d} (\tilde{F}_t^+, \tilde{F}_t^-)_{t \geq 0} \quad \text{in } \mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2) \quad (31)$$

along a subsequence $\lambda(N) \rightarrow 0$ implies that the limit $(\tilde{F}_t^+, \tilde{F}_t^-)$ is the unique solution of the ODE in $L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2$ considered in Theorem 4.4 (i). We use here the convergence of martingales associated with the Markov chain.

Steps 1 and 2 together imply the announced convergence in $\mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$.

• Proof of Step 1 — tightness.

Fix $T > 0$. By Proposition 4.5, for $\lambda \in (0, \lambda_0]$, where λ_0 may depend on T , the training error $t \mapsto L_n(\hat{F}_{[t/\lambda]}^\lambda)$ is a.s. non-increasing on $[0, T]$ and hence bounded from above by its initial value $C = L_n(F_0)$. Then Lemma 5.1 implies the following control of the increments: for $0 \leq t_1 \leq t_2 \leq T$,

$$\left\| \hat{F}_{[t_2/\lambda]}^{\lambda\pm} - \hat{F}_{[t_1/\lambda]}^{\lambda\pm} \right\|_{\text{TV}} \leq \lambda ([t_2/\lambda] - [t_1/\lambda]) 4^d M. \quad (32)$$

Indeed, there are $[t_2/\lambda] - [t_1/\lambda]$ increments, each with norm less than $\lambda 4^d M$. This implies the following global control in total variation norm:

$$\left\| \hat{F}_{[t/\lambda]}^{\lambda\pm} \right\|_{\text{TV}} \leq \|F_0\|_{\text{TV}} + 4^d MT, \quad t \in [0, T].$$

We deduce that the Markov chain remains up to time T in the set

$$\mathcal{K} := \{F \in \overline{\mathbb{T}}^+ : \|F\|_{L_{\mathbf{x}}^2} \leq \|F_0\|_{\text{TV}} + 4^d MT\},$$

which is a compact subset of $\overline{\mathbb{T}}^+$ according to Proposition 3.6. This proves a compact containment condition: for all $\lambda \in (0, \lambda_0]$,

$$\mathbb{P}((\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-}) \in \mathcal{K} \times \mathcal{K}, t \in [0, T]) = 1.$$

Furthermore, Inequality (32) allows to control the modulus of continuity of the path $t \mapsto F_{[t/\lambda]}^{\lambda\pm}$. We define, for $\delta > 0$,

$$\omega_{\hat{F}^{\lambda\pm}}(\delta) = \sup_{\substack{0 \leq t_1 \leq t_2 \leq T \\ |t_1 - t_2| < \delta}} \left\| \hat{F}_{[t_2/\lambda]}^{\lambda\pm} - \hat{F}_{[t_1/\lambda]}^{\lambda\pm} \right\|_{\text{TV}}.$$

Inequality (32) implies

$$\omega_{\hat{F}^{\lambda\pm}}(\delta) \leq (\delta + \lambda) 4^d M \quad \text{a.s.}$$

Taking $\delta < \eta/(2 \cdot 4^d M)$ and $\lambda \leq \inf(\delta, \lambda_0)$, we have

$$\mathbb{P}(\omega_{\hat{F}^{\lambda\pm}}(\delta) \geq \eta) = 0.$$

According to Ethier and Kurtz (1986, Chapter 3, Corollary 7.4), we obtain the tightness in $\mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$ of the sequence of processes $(\hat{F}_{[t/\lambda]}^{\lambda+}, \hat{F}_{[t/\lambda]}^{\lambda-})_{t \geq 0}$, where $\lambda = \lambda(N) \rightarrow 0$ is an arbitrary sequence.

• Proof of Step 2 — identification of the limit.

Assume the weak convergence (31) along a subsequence $\lambda = \lambda(N) \rightarrow 0$. Consider the processes $(G_t^{\lambda\pm})_{t \geq 0}$ in $\mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$ defined by

$$G_t^{\lambda\pm} = \hat{F}_{[t/\lambda]}^{\lambda\pm} - F_0 - \int_0^t \mathcal{T}^\pm(\hat{F}_{[s/\lambda]}^\lambda) ds, \quad t \geq 0.$$

We take here the (Bochner-)integral of an $L_{\mathbf{x}}^2$ -valued function and refer to Diestel and Uhl (1977) for integration in Banach spaces. We will only integrate locally

bounded $L_{\mathbf{x}}^2$ -valued functions, which are Bochner-integrable simply because $L_{\mathbf{x}}^2$ is separable. In this context the integral behaves mostly as usual: in particular, for a continuous integrand $t \mapsto \varphi(t)$, the function $t \mapsto \int_0^t \varphi(s) ds$ is C^1 with derivative φ .

Consider the map $\Phi : \mathbb{D}([0, \infty) L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2) \rightarrow \mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$ defined by $\Phi(F^\pm) = G^\pm$ with

$$G_t^\pm = F_t^\pm - F_0^\pm - \int_0^t \mathcal{T}^\pm(F_s^+ - F_s^-) ds, \quad t \geq 0.$$

Because Φ is continuous and $(\hat{F}_{[t/\lambda]}^{\lambda\pm})_{t \geq 0} \xrightarrow{d} (\tilde{F}_t^\pm)_{t \geq 0}$, the continuous mapping theorem (Billingsley, 1999, Theorem 2.7) implies the convergence

$$(G_t^{\lambda\pm})_{t \geq 0} = \Phi((\hat{F}_{[t/\lambda]}^{\lambda\pm})_{t \geq 0}) \xrightarrow{d} \Phi((\tilde{F}_t^\pm)_{t \geq 0}). \quad (33)$$

On the other hand, consider the discrete time processes

$$M_m^{\lambda\pm} = \hat{F}_m^{\lambda\pm} - F_0^\pm - \lambda \sum_{k=1}^m \mathcal{T}^\pm(\hat{F}_k^\lambda) ds, \quad m \geq 0. \quad (34)$$

Its increments are given by

$$\begin{aligned} M_{m+1}^{\lambda\pm} - M_m^{\lambda\pm} &= \hat{F}_{m+1}^{\lambda\pm} - \hat{F}_m^{\lambda\pm} - \lambda \mathcal{T}^\pm(\hat{F}_m^\lambda) \\ &= \lambda T^\pm(\cdot; \hat{F}_m^\lambda, \zeta_{m+1}) - \lambda \mathbb{E}[T^\pm(\cdot; \hat{F}_m^\lambda, \zeta_{m+1}) \mid \hat{F}_m^\lambda], \quad m \geq 0. \end{aligned}$$

This shows that $(M_m^{\lambda\pm})_{m \geq 0}$ are martingales with respect to the filtration $\mathcal{F}_m = \sigma(\zeta_k, 1 \leq k \leq m)$, $m \geq 0$, with values in $L_{\mathbf{x}}^2$. More precisely, $(M_m^{\lambda\pm})_{m \geq 0}$ are square-integrable martingales with values in the separable Hilbert space $L_{\mathbf{x}}^2$. For the existence of conditional expectation and properties of square-integrable martingales in Hilbert spaces, we refer to Métivier (1982). Here we will use only basic properties that are satisfied exactly as in the scalar case, namely the fact that

$$\mathbb{E}[\|M_m^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2] = \sum_{k=1}^m \mathbb{E}[\|M_k^{\lambda\pm} - M_{k-1}^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2], \quad (35)$$

and Doob's inequality

$$\mathbb{E} \left[\sup_{0 \leq k \leq m} \|M_k^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2 \right] \leq 4 \mathbb{E}[\|M_m^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2]. \quad (36)$$

We prove below that, as $N \rightarrow \infty$,

$$\sup_{0 \leq t \leq T} \|G_t^{\lambda\pm} - M_{[t/\lambda]}^{\lambda\pm}\|_{L_{\mathbf{x}}^2} \longrightarrow 0 \quad \text{a.s.}, \quad (37)$$

and

$$\sup_{0 \leq t \leq T} \|M_{[t/\lambda]}^{\lambda\pm}\|_{L_{\mathbf{x}}^2} \longrightarrow 0 \quad \text{in probability.} \quad (38)$$

Equation (37) is rather straightforward and consists in handling a boundary term. Indeed, writing the sum in Equation (34) as an integral, we get

$$M_{[t/\lambda]}^{\lambda\pm} = \hat{F}_{[t/\lambda]}^{\lambda\pm} - F_0^\pm - \int_0^{\lambda[t/\lambda]} \mathcal{T}^\pm(\cdot; \hat{F}_{[s/\lambda]}^\lambda) ds$$

and

$$G_t^{\lambda\pm} - M_{[t/\lambda]}^{\lambda\pm} = \int_{\lambda[t/\lambda]}^t \mathcal{T}^\pm(\cdot; \hat{F}_{[s/\lambda]}^\lambda) ds.$$

By the triangle inequality, we deduce

$$\sup_{0 \leq t \leq T} \|G_t^{\lambda\pm} - M_{[t/\lambda]}^{\lambda\pm}\|_{L_{\mathbf{x}}^2} \leq \lambda \sup_{0 \leq s \leq T} \|\mathcal{T}^\pm(\cdot; \hat{F}_{[s/\lambda]}^\lambda)\|_{L_{\mathbf{x}}^2}.$$

Since $\|\cdot\|_{L_{\mathbf{x}}^2} \leq \|\cdot\|_{\text{TV}}$, by Proposition 4.5 and the proof of tightness above, the supremum remains bounded for $\lambda \in (0, \lambda_0]$ whence we deduce the almost sure convergence to 0 as $\lambda(N) \rightarrow 0$.

For the proof of (38), consider (35) with $m = [T/\lambda]$ and $\lambda \in (0, \lambda_0]$. We get

$$\begin{aligned} \mathbb{E}[\|M_m^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2] &= \sum_{k=1}^m \mathbb{E}[\|M_k^{\lambda\pm} - M_{k-1}^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2], \\ &= \lambda^2 \sum_{k=1}^m \mathbb{E}[\|T^\pm(\cdot; \hat{F}_m^\lambda, \zeta_{m+1}) - \mathcal{T}^\pm(\hat{F}_m^\lambda)\|_{L_{\mathbf{x}}^2}^2]. \end{aligned}$$

Again because $\|\cdot\|_{L_{\mathbf{x}}^2} \leq \|\cdot\|_{\text{TV}}$ and since $T^\pm(\cdot; \hat{F}_m^\lambda, \zeta_{m+1})$ and $\mathcal{T}^\pm(\hat{F}_m^\lambda)$ are almost surely bounded in total variation by a deterministic constant, we have $\mathbb{E}[\|M_m^{\lambda\pm}\|_{L_{\mathbf{x}}^2}^2] = O(\lambda) \rightarrow 0$. An application of Doob's inequality (36) is enough to conclude.

Putting (37) and (38) together, we get

$$(G_t^{\lambda\pm})_{t \geq 0} \xrightarrow{d} 0 \quad \text{in } \mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2),$$

so that (33) becomes

$$\tilde{F}_t^\pm = \tilde{F}_0^\pm + \int_0^t \mathcal{T}^\pm(\tilde{F}_s^+ - \tilde{F}_s^-) ds, \quad t \geq 0.$$

In other words, $(\tilde{F}^\pm)_{t \geq 0}$ is the solution to the ODE (20). \square

Proof of Theorem 1.9. We assume here that, for all $t \in [0, T]$, \hat{F}_t^+ and \hat{F}_t^- are continuous on $[0, 1]^p$ and that $t \mapsto \hat{F}_t^\pm$ are continuous for the uniform norm — this is a consequence of Proposition 4.11 below. Let us fix $T > 0$, $\varepsilon > 0$ and $\delta > 0$ such that

$$\sup_{\substack{t, s \in [0, T] \\ |t-s| \leq \delta}} \|\hat{F}_t^\pm - \hat{F}_s^\pm\|_\infty \leq \varepsilon. \quad (39)$$

Let us also fix a subdivision $0 = t_0 < t_1 < \dots < t_k = T$ with $t_i - t_{i-1} \leq \delta$.

Using Skorokod's representation theorem in the separable space $\mathbb{D}([0, \infty), L_{\mathbf{x}}^2 \times L_{\mathbf{x}}^2)$, let us assume that $(\hat{F}_{[t/\lambda]}^{\lambda\pm})_{t \geq 0} \rightarrow (\hat{F}_t^\pm)_{t \geq 0}$ almost surely. Because the $F_{t_i}^\pm$ are continuous and in \mathbb{T}^+ , Proposition 3.6 (ii) implies

$$\|\hat{F}_{[t_i/\lambda]}^{\lambda\pm} - \hat{F}_{t_i}^\pm\|_\infty \longrightarrow 0, \quad 0 \leq i \leq k \quad (40)$$

almost surely as $\lambda \rightarrow 0$. By construction $\hat{F}_{[t/\lambda]}^{\lambda\pm}(x)$ is non-decreasing in t for all x , so that we have

$$\hat{F}_{[t_{i-1}/\lambda]}^{\lambda\pm} \leq \hat{F}_{[t/\lambda]}^{\lambda\pm} \leq \hat{F}_{[t_i/\lambda]}^{\lambda\pm}, \quad \text{for } t \in [t_{i-1}, t_i].$$

By the construction of the subdivision $(t_i)_{0 \leq i \leq k}$ based on Equation (39), we deduce

$$\begin{aligned} \|\hat{F}_{[t/\lambda]}^{\lambda\pm} - \hat{F}_t^\pm\|_\infty &\leq \max\left(\|\hat{F}_{[t_{i-1}/\lambda]}^{\lambda\pm} - \hat{F}_t^\pm\|_\infty, \|\hat{F}_{[t/\lambda]}^{\lambda\pm} - \hat{F}_t^\pm\|_\infty\right) \\ &\leq \max\left(\|\hat{F}_{[t_{i-1}/\lambda]}^{\lambda\pm} - \hat{F}_{t_{i-1}}^\pm\|_\infty, \|\hat{F}_{[t/\lambda]}^{\lambda\pm} - \hat{F}_{t_i}^\pm\|_\infty\right) + 2\varepsilon \end{aligned}$$

Therefore Equation (40) implies $\limsup_{\lambda \rightarrow 0} \|\hat{F}_{[t/\lambda]}^{\lambda\pm} - \hat{F}_t^\pm\|_\infty \leq \varepsilon$. Since $\varepsilon > 0$ is arbitrary, this concludes the proof of Theorem 1.9. \square

5.3.3 Proofs related to Section 4.3

Proof of Proposition 4.6.

- (i) This monotonicity property is established in the proof of Theorem 4.3, see Equation (29).
- (ii) Because L is convex and differentiable with respect to the second variable, the initialization (5) implies that $\bar{r}_t = 0$ at time $t = 0$. Similarly as in Equation (29), the derivative with respect to $t \geq 0$ is given by

$$\begin{aligned} \frac{d}{dt} \bar{r}_t &= \frac{1}{n} \sum_{i=1}^n \mathcal{T}(F_t)(x_i) \frac{\partial^2 L}{\partial z^2}(y_i, \hat{F}_t(x_i)) \\ &= -\frac{1}{n} \mathbb{E}_\zeta \left[\sum_{v \in \{0,1\}^d} \sum_{i=1}^n \frac{\partial L}{\partial z}(y_i, \hat{F}_t(x_i)) \mathbf{1}_{A_v}(x_i) \right] \\ &= -\bar{r}_t. \end{aligned} \tag{41}$$

In the second line, for each term indexed by v , a factor $\sum_{i=1}^n \frac{\partial L^2}{\partial z^2}(y_i, \hat{F}_t(x_i)) \mathbf{1}_{A_v}(x_i)$ is canceling out. The third line follows because the sum of all terms is equal to \bar{r}_t and does not depend on the randomness ζ .

Equation (41) implies that $\bar{r}_t = \bar{r}_0 e^{-t}$. This is equal to 0 for initialization (5). Interestingly, a different initialization would provide an exponentially fast decay of the residuals mean. \square

The following Lemma will be used in the proof of Proposition 4.8.

Lemma 5.3. *Let $(x_i, r_i)_{1 \leq i \leq n}$ and (β, K, d) be fixed. Consider $J \subset \llbracket 1, p \rrbracket$ with $|J| \leq d$ and define*

$$A(h) = \{x \in [0, 1]^p : \forall j \in J, x^j < h^j\}, \quad h \in [0, 1]^J.$$

Then for any $h_0 \in [0, 1]^p$ and $\varrho > 0$, there exists $q > 0$ such that

$$P_{\beta, K} \left(A_0(\xi) = A(h) \text{ for some } h \text{ such that } \|h - h_0\|_\infty \leq \varrho \right) \geq q,$$

where $P_{\beta, K}$ is the splitting scheme distribution in Equation (19) and $A_0(\xi)$ is the region containing 0 in the partition associated to ξ . Furthermore, the constant q depends only on p, β, K, d, ϱ and $R := \max_i |r_i|$.

Proof of Lemma 5.3. We first consider the case $\beta = 0$, i.e. $P_{\beta,K} = P_0$, and prove that the event

$$E_{h_0,\varrho} = \{A_0(\xi) = A(h) \text{ for some } h \text{ such that } \|h - h_0\|_\infty \leq \varrho\}$$

satisfies $P_0(E_{h_0,\varrho}) \geq q$ with $q > 0$ depending only on p , d and ϱ . Possibly shifting h_0 by a small distance and reducing ϱ , we may assume without loss of generality that $h_0 \in [\varrho, 1 - \varrho]^J$. For simplicity, assume first that $|J| = d$, and let $j_1 < \dots < j_d$ be the increasing enumeration of J . We recall the notation $\xi = (j_v, u_v)_{v \in \mathcal{T}_{d-1}}$ and observe that the region $A_0(\xi)$ containing 0 depends only of the values (j_v, u_v) for the d leftmost nodes in \mathcal{T}_{d-1} that we denote by v_1, v_2, \dots, v_d . We then have

$$\begin{aligned} P_0(E_{h_0,\varrho}) &\geq P_0(\forall k \in \llbracket 1, d \rrbracket, j_{v_k} = j_k \text{ and } u_{v_k} \in [h_0^{j_k} - \varrho, h_0^{j_k} + \varrho]) \\ &= \left(\frac{2\varrho}{p}\right)^d =: q. \end{aligned}$$

This simple argument is adapted to the case when $|J| = r < d$ by considering the sequence $j_1 < \dots < j_r = \dots = j_d$ and noting that $\xi \in E_{h_0,\varrho}$ if and only if

$$j_{v_k} = j_k \quad \text{and} \quad u_{v_k} \in [h_0^{j_k} - \varrho, h_0^{j_k} + \varrho], \quad \text{for } k = 1, \dots, r-1,$$

and

$$j_{v_r} = \dots = j_{v_d} = j_r, \quad \prod_{k=r}^d u_{v_k} \in [h_0^{j_r} - \varrho, h_0^{j_r} + \varrho].$$

Further details are left to the reader.

Next we consider the case $\beta > 0$. Since $P_{\beta,K}$ is absolutely continuous with respect to P_0 according to Proposition 2.1, we can write

$$P_{\beta,K}(E_{h_0,\varrho}) = \int_{E_{h_0,\varrho}} \frac{dP_{\beta,K}}{dP_0}(\xi) P_0(d\xi).$$

We prove that the Radon-Nykodym derivative (19) is bounded from below by a positive constant. By definition each of the $\Delta(s_v^k, u_v^k; A_v(\xi))$ is a nonnegative polynomial with degree 2 in the (r_i) and with coefficients bounded by 2. Therefore there is a uniform bound

$$0 \leq \Delta(s_v^k, u_v^k; A_v(\xi)) \leq C,$$

where C depends only on $R := \max_i |r_i|$. This implies that each factor in (19) satisfies

$$\frac{\exp(\beta \Delta(s_v^1, u_v^1; A_v(\xi)))}{\sum_{k=1}^K \exp(\beta \Delta(s_v^k, u_v^k; A_v(\xi)))} \geq \frac{1}{1 + (K-1)e^{\beta C}}.$$

We deduce that the Radon-Nykodym derivative (19) is bounded from below by a positive constant depending only on β, K, d and R . Since $P_0(E_{h_0,\varrho}) \geq q$, this yields a lower bound for $P_{\beta,K}(E_{h_0,\varrho})$ and concludes the proof. \square

Proof of Proposition 4.8. It is trivially checked that, for $F \in \mathbb{B}$, the nullity of the residuals $r_i = \frac{\partial L}{\partial z}(y_i, F(x_i))$ implies $\mathcal{T}(F) = 0$. We prove here the converse implication under assumption (A5). We consider $F \in \mathbb{B}$ such that $(r_i)_{1 \leq i \leq n} \neq 0$, and our goal is to show that $\mathcal{T}(F) \neq 0$.

Recall from (29) that we can write

$$\sum_{i=1}^n \mathcal{T}(F)(x_i) r_i = -\mathbb{E}_\xi \left[\sum_{v \in \{0,1\}^d} \frac{\left(\sum_{i=1}^n r_i \mathbb{1}_{A_v}(x_i) \right)^2}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{A_v}(x_i)} \right].$$

It is therefore sufficient to show that the sum under the expectation is not null with positive probability. All the terms being nonnegative, we show that the term corresponding to the region A_0 is positive with positive probability. Consider the set of indices $i \in \llbracket 1, n \rrbracket$ such that $r_i \neq 0$, and among those, fix i_0 such that x_{i_0} is minimal for the lexicographic order on $[0, 1]^p$. Under assumption (A5), there is J such that the $(x_i^J)_i$ are pairwise distinct, which ensures the positivity of $\varrho = \min_{1 \leq i_1 < i_2 \leq n} \|x_{i_1}^J - x_{i_2}^J\|_\infty > 0$. By definition of i_0 and ϱ , for $h \in (x_{i_0}^J, x_{i_0}^J + \varrho)$ (addition by ϱ meant componentwise), the region $A(h)$ contains only one point x_i corresponding to a non-null residual and this point is x_{i_0} . Lemma 5.3 shows that

$$P_{\beta, K}(A_0(\xi) = A(h) \text{ for some } h \in (x_{i_0}^J, x_{i_0}^J + \varrho)) > 0.$$

On this event, we have $(\sum_{i=1}^n r_i \mathbb{1}_{A_0}(x_i))^2 = r_{i_0}^2 > 0$ and the sum is positive. This implies that $\sum_{i=1}^n \mathcal{T}(F)(x_i) r_i \neq 0$ and hence $\mathcal{T}(F) \neq 0$. \square

Proof of Proposition 4.10. According to Proposition 4.6, the non-negative function $t \mapsto L_n(\hat{F}_t)$ is non-increasing. We denote by $\ell \geq 0$ its limit as $t \rightarrow +\infty$ and we prove that $\ell = 0$. We proceed by contradiction and assume that $\ell > 0$. Then we will prove below that

$$\sup_{t \geq 0} \frac{d}{dt} L_n(\hat{F}_t) \leq -\eta \quad \text{for some } \eta > 0,$$

whence $L_n(\hat{F}_t) \leq L_n(\hat{F}_0) - \eta t$, yielding a contradiction because the function is non-negative.

We will use the following properties of L_n :

- (i) for fixed $(y_i)_{1 \leq i \leq n}$ and $\ell_1 > 0$, there is a constant $R > 0$ such that $L_n(F) \leq \ell_1$ implies

$$\max_{1 \leq i \leq n} \left| \frac{\partial L}{\partial z}(y_i, F(x_i)) \right| \leq R;$$

- (ii) for fixed $(y_i)_{1 \leq i \leq n}$ and $\ell_0 > 0$, there is a constant $\delta > 0$ such that $L_n(F) \geq \ell_0$ implies

$$\max_{1 \leq i \leq n} \left| \frac{\partial L}{\partial z}(y_i, F(x_i)) \right| \geq \delta.$$

The first point relies on the fact that $z \mapsto L(y, z)$ is convex and non-negative for all y (Assumption 1.8). The second point uses the fact that $z \mapsto L(y, z)$ has infimum 0 (Assumption (A3)); then $L_n(F) \geq \ell_0$ implies $L(y_i, F(x_i)) \geq \ell_0/n$ for some i , so that $F(x_i)$ is far from the minimizer of the convex function $z \mapsto L(y_i, z)$ and the (non-increasing) derivative must be bounded away from 0.

We let $\ell_0 = \ell > 0$ and $\ell_1 = L_n(\hat{F}_0)$. Note that $\ell_0 \leq L_n(\hat{F}_t) \leq \ell_1$ for all $t \geq 0$ so that we can make use of the bound (i) and (ii) above. By (29), $\frac{d}{dt}L_n(\hat{F}_t) = \sum_i \mathcal{T}(\hat{F}_t)(x_i)r_i$ where

$$\sum_{i=1}^n \mathcal{T}(\hat{F}_t)(x_i)r_i = -\mathbb{E}_\xi \left[\sum_{v \in \{0,1\}^d} \frac{\left(\sum_{i=1}^n r_i \mathbb{1}_{A_v}(x_i)\right)^2}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{A_v}(x_i)} \right],$$

where we use again, for conciseness, the notation $r_i = \frac{\partial L}{\partial z}(y_i, \hat{F}_t(x_i))$. We need to provide a lower bound for $\sum_{i=1}^n \mathcal{T}(F)(x_i)r_i$ when $\ell_0 \leq L_n(F) \leq \ell_1$ and use a similar argument as in the proof of Proposition 4.8. Let J be the set of coordinates appearing in (A5) and set $\varrho = \frac{1}{2} \min_{i_1 \neq i_2} \|x_{i_1}^J - x_{i_2}^J\|_\infty > 0$. Consider the sets

$$A(h) = \{x \in [0, 1]^p : x^j < h^j \text{ for all } j \in J\}, \quad h \in [0, 1]^J.$$

By (ii) above, $L_n(F) \geq \ell_0$ implies that there is an index i_0 such that $|r_{i_0}| \leq \delta$. Introducing the function

$$S(h) = \sum_{i=1}^n r_i \mathbb{1}_{A(h)}(x_i),$$

we have

$$r_{i_0} = \sum_{v \in \{-1, 1\}^J} \sigma_v S(x_i + \varrho v) \quad \text{with} \quad \sigma_v = (-1)^{\sum_k \mathbb{1}_{v_k = -1}}. \quad (42)$$

This is due to the fact that S is a tree function on $[0, 1]^J$ with $\mu_S = \sum_{i=1}^n r_i \delta_{x_i^J}$ and that $x_{i_0}^J$ is the only atom in the hypercube $[x_{i_0}^J - \varrho, x_{i_0}^J + \varrho]$ — addition by ϱ meant componentwise. Then the bound $r_{i_0} \geq \delta$ together with Equation (42) imply that

$$|S(x_{i_0} + \varrho v)| \geq 2^{-|J|} |r_{i_0}| \geq 2^{-d} \delta$$

for some $v \in \{-1, 1\}^J$. Furthermore, by construction we have $S(h) = S(x_i + \varrho v)$ for all h such that $\|h - x_{i_0}^J - \varrho v\|_\infty < \varrho$. Now we use Lemma 5.3 to bound from below the probability of the event

$$E = \{A_0(\xi) = A(h) \text{ for some } h \in [x_{i_0}^J - \varrho, x_{i_0}^J + \varrho]\}$$

More precisely, we have $P_{\beta, K}(E) \geq q > 0$, where q depends only on p, β, K, d, ϱ and $R \geq \max_i |r_i|$. By point (i) above, R depends on ℓ_1 but not on F satisfying $L_n(F) \leq \ell_1$. On the event E , we have $\left(\sum_{i=1}^n r_i \mathbb{1}_{A(h)}\right)^2 \geq 4^{-d} \delta^2$ and, by (A4),

$$\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{A(h)} \leq nC,$$

where C depends only on R and the (y_i) . Therefore on the event E , we have

$$\frac{\left(\sum_{i=1}^n r_i \mathbb{1}_{A(h)}\right)^2}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{A(h)}} \geq \frac{4^{-d} \delta^2}{nC},$$

whence we deduce

$$\mathbb{E}_\xi \left[\sum_{v \in \{0,1\}^d} \frac{\left(\sum_{i=1}^n r_i \mathbb{1}_{A_v}(x_i)\right)^2}{\sum_{i=1}^n \frac{\partial^2 L}{\partial z^2}(y_i, F(x_i)) \mathbb{1}_{A_v}(x_i)} \right] \geq q \frac{4^{-d} \delta^2}{nC} > 0.$$

This proves that $L_n(\hat{F}_t) \rightarrow 0$. It follows that $L(y_i, \hat{F}_t(x_i)) \rightarrow 0$ for all i . Since 0 is the infimum of the convex function $z \mapsto L(y_i, z)$, the convergence $\frac{\partial L}{\partial z}(y_i, \hat{F}_t(x_i)) \rightarrow 0$ holds for the derivative as well, ending the proof. \square

Our proof of Proposition 4.11 relies on the following lemma.

Lemma 5.4. *For all $F \in \mathbb{B}$, $\mathcal{T}(F) \in \mathbb{W}$ and the mapping $\mathcal{T} : (\mathbb{B}, \|\cdot\|_\infty) \rightarrow (\mathbb{W}, \|\cdot\|_W)$ is locally Lipschitz. Furthermore, the same is true for \mathcal{T}^+ and \mathcal{T}^- .*

Similarly as in Remark 4.2, the inclusion $\mathbb{W} \subset \mathbb{B}$ with norm inequality $\|F\|_\infty \leq \pi_0([0, 1]^p) \|F\|_W$ for all $F \in \mathbb{W}$ implies that \mathcal{T} can be seen as a locally Lipschitz operator on $(\mathbb{W}, \|\cdot\|_W)$.

Proof of Lemma 5.4. We consider only the proof for \mathcal{T} which is easily adapted for \mathcal{T}^+ and \mathcal{T}^- by considering the positive and negative part of the gradient trees. It is enough to prove the following properties:

(i) for all $F \in \mathbb{B}$, there is $M \geq 0$ such that

$$|\mu_{\mathcal{T}(F)}(A)| \leq M\pi_0(A), \quad \text{for all borel set } A \subset [0, 1]^p;$$

(ii) for all $R > 0$, there exists $M \geq 0$ such that for all F and $F' \in \mathbb{B}$ satisfying $\max(\|F\|_\infty, \|F'\|_\infty) \leq R$,

$$|\mu_{\mathcal{T}(F)}(A) - \mu_{\mathcal{T}(F')}(A)| \leq M\|F - F'\|_\infty \pi_0(A), \quad \text{for all borel set } A \subset [0, 1]^p.$$

Proof of (i). According to Proposition 3.3 and its proof, for each splitting scheme ξ , the measure associated to the gradient tree $\tilde{T}(\cdot; F, \xi)$ is a point measure with atoms at the vertices in $[0, 1]^p$ of the partition $(A_v(\xi))_{v \in \{0, 1\}^d}$ and each atom has a mass with absolute value less than $\|\tilde{T}(\cdot; F, \xi)\|_{\text{TV}}$. By the definition of π_ξ , this implies

$$\mu_{\tilde{T}(\cdot; F, \xi)} \ll \pi_\xi \quad \text{with} \quad \left| \frac{d\mu_{\tilde{T}(\cdot; F, \xi)}}{d\pi_\xi}(x) \right| \leq \|\tilde{T}(\cdot; F, \xi)\|_{\text{TV}} \leq C,$$

for some $C > 0$ that depends only on $\|F\|_\infty$ — this last inequality is a consequence of Equations (27) with $F' = 0$. We deduce

$$\begin{aligned} |\mu_{\mathcal{T}(F)}(A)| &= \left| \int \mu_{\tilde{T}(\cdot; F, \xi)}(A) P_{\beta, K}(d\xi) \right| \\ &\leq \int \left(\int_A \left| \frac{d\mu_{\tilde{T}(\cdot; F, \xi)}}{d\pi_\xi}(x) \right| \pi_\xi(dx) \right) \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) \\ &\leq \int C \pi_\xi(A) \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi). \end{aligned}$$

Since $\frac{dP_{\beta, K}}{dP_0}(\xi) \leq K^{2^d - 1}$ according to Equation (19), we get

$$\begin{aligned} |\mu_{\mathcal{T}(F)}(A)| &\leq CK^{2^d - 1} \int \pi_\xi(A) P_0(d\xi) \\ &= CK^{2^d - 1} \pi_0(A). \end{aligned}$$

We deduce $\mu_{\mathcal{T}(F)} \ll \pi_0$ with Radon-Nykodym derivative bounded by CK^{2^d-1} , so that $\mathcal{T}(F) \in \mathbb{W}$.

Proof of (ii). With similar notation as in the proof of Proposition 4.1, we compute

$$\begin{aligned}
& |\mu_{\mathcal{T}(F)}(A) - \mu_{\mathcal{T}(F')}(A)| \\
& \leq \int \left(\int_A \left| \frac{d\mu_{\tilde{T}(\cdot; F, \xi)}(x)}{d\pi_\xi} - \frac{d\mu_{\tilde{T}(\cdot; F', \xi)}(x)}{d\pi_\xi} \right| \pi_\xi(dx) \right) \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) \\
& \quad + \int \left(\int_A \left| \frac{d\mu_{\tilde{T}(\cdot; F', \xi)}(x)}{d\pi_\xi} \right| \pi_\xi(dx) \right) \left| \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) - \frac{dP'_{\beta, K}}{dP_0}(\xi) \right| P_0(d\xi) \\
& \leq \int \left(\int_A \|\tilde{T}(\cdot; F, \xi) - \tilde{T}(\cdot; F', \xi)\|_{\text{TV}} \pi_\xi(dx) \right) \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) \\
& \quad + \int C\pi_\xi(A) \left| \frac{dP_{\beta, K}}{dP_0}(\xi) P_0(d\xi) - \frac{dP'_{\beta, K}}{dP_0}(\xi) \right| P_0(d\xi) \\
& \leq (4^d CK^{2^d-1} + CR) \|F - F'\|_\infty \pi_0(A),
\end{aligned}$$

where the constant C depends on R but not on F or F' . The last inequality relies on Equations (27) and (23). \square

Proof of Proposition 4.11. We consider the proof for $(\hat{F}_t)_{t \geq 0}$ only since it can be easily adapted to $(\hat{F}_t^+)_{t \geq 0}$ and $(\hat{F}_t^-)_{t \geq 0}$. By Theorem 4.3, infinitesimal gradient boosting is a solution of the ODE (15) in the space $(\mathbb{T}, \|\cdot\|_{\text{TV}})$ so that

$$\hat{F}_t = \hat{F}_0 + \int_0^t \mathcal{T}(\hat{F}_s) ds, \quad t \geq 0.$$

By Lemma 5.4, \mathcal{T} takes its value in $\mathbb{W} \subset \mathbb{T}$ and, since $\hat{F}_0 \in \mathbb{W}$, it implies that $\hat{F}_t \in \mathbb{W}$ for all $t \geq 0$. Furthermore, $\mathcal{T} : \mathbb{W} \rightarrow \mathbb{W}$ is locally Lipschitz so that $(\hat{F}_t)_{t \geq 0}$ is the unique solution of the ODE (15) in $(\mathbb{W}, \|\cdot\|_{\mathbb{W}})$. Therefore infinitesimal gradient boosting defines a smooth path in \mathbb{W} . \square

References

- Billingsley, P. (1995). *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, third edition. A Wiley-Interscience Publication.
- Billingsley, P. (1999). *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York. doi:[10.1002/9780470316962](https://doi.org/10.1002/9780470316962), second edition. A Wiley-Interscience Publication.
- Billingsley, P. and Topsøe, F. (1967). Uniformity in weak convergence. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 7:1–16. doi:[10.1007/BF00532093](https://doi.org/10.1007/BF00532093).
- Blanchard, G., Lugosi, G., and Vayatis, N. (2004). On the rate of convergence of regularized boosting classifiers. *J. Mach. Learn. Res.*, 4(5):861–894. doi:[10.1162/1532443041424319](https://doi.org/10.1162/1532443041424319).

- Breiman, L. (2004). Population theory for boosting ensembles. *Ann. Statist.*, 32(1):1–11. doi:[10.1214/aos/1079120126](https://doi.org/10.1214/aos/1079120126).
- Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and Regression Trees*. Chapman & Hall.
- Bühlmann, P. and Yu, B. (2003). Boosting with the L_2 loss: regression and classification. *J. Amer. Statist. Assoc.*, 98(462):324–339. doi:[10.1198/016214503000125](https://doi.org/10.1198/016214503000125).
- 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, San Francisco California USA. ACM. doi:[10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785).
- Diestel, J. and Uhl, J. J. (1977). *Vector Measures*, volume 15 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, Rhode Island. doi:[10.1090/surv/015](https://doi.org/10.1090/surv/015).
- Dombry, C. and Esstafa, Y. (2020). Behaviour of linear l^2 -boosting algorithm in the vanishing learning rate asymptotic. arXiv:[2012.14657](https://arxiv.org/abs/2012.14657). preprint.
- Ethier, S. N. and Kurtz, T. G., editors (1986). *Markov Processes*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ, USA. doi:[10.1002/9780470316658](https://doi.org/10.1002/9780470316658).
- Freund, Y. and Schapire, R. (1999). Adaptive game playing using multiplicative weights. volume 29, pages 79–103. doi:[10.1006/game.1999.0738](https://doi.org/10.1006/game.1999.0738). Learning in games: a symposium in honor of David Blackwell.
- Friedman, J., Hastie, T., and Tibshirani, R. (2000). Additive logistic regression: a statistical view of boosting. *Ann. Statist.*, 28(2):337–407. doi:[10.1214/aos/1016218223](https://doi.org/10.1214/aos/1016218223). With discussion and a rejoinder by the authors.
- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232. doi:[10.1214/aos/1013203451](https://doi.org/10.1214/aos/1013203451).
- Friedman, J. H. (2002). Stochastic gradient boosting. *Computational statistics & data analysis*, 38(4):367–378. doi:[10.1016/S0167-9473\(01\)00065-2](https://doi.org/10.1016/S0167-9473(01)00065-2).
- Geurts, P., Ernst, D., and Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, 63(1):3–42. doi:[10.1007/s10994-006-6226-1](https://doi.org/10.1007/s10994-006-6226-1).
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data mining, inference, and prediction*. Springer Series in Statistics. Springer, New York. doi:[10.1007/978-0-387-84858-7](https://doi.org/10.1007/978-0-387-84858-7), second edition.
- Jiang, W. (2004). Process consistency for AdaBoost. *Ann. Statist.*, 32(1):13–29. doi:[10.1214/aos/1079120128](https://doi.org/10.1214/aos/1079120128).
- Lugosi, G. and Vayatis, N. (2004). On the Bayes-risk consistency of regularized boosting methods. *Ann. Statist.*, 32(1):30–55. doi:[10.1214/aos/1079120129](https://doi.org/10.1214/aos/1079120129).

- Métivier, M. (1982). *Semimartingales: A Course on Stochastic Processes*. De Gruyter, Berlin, New York. doi:[10.1515/9783110845563](https://doi.org/10.1515/9783110845563).
- Neuhaus, G. (1971). On weak convergence of stochastic processes with multidimensional time parameter. *The Annals of Mathematical Statistics*, 42(4):1285–1295. doi:[10.1214/aoms/1177693241](https://doi.org/10.1214/aoms/1177693241).
- Ridgeway, G. (2007). Generalized boosting models: a guide to the gbm package. url:<https://cran.r-project.org/web/packages/gbm/vignettes/gbm.pdf>.
- Zhang, T. and Yu, B. (2005). Boosting with early stopping: convergence and consistency. *Ann. Statist.*, 33(4):1538–1579. doi:[10.1214/009053605000000255](https://doi.org/10.1214/009053605000000255).