

Some upper bounds for the rate of convergence of penalized likelihood context tree estimators

Florencia Leonardi*

*Instituto de Matemática e Estatística, Universidade de São Paulo
Rua do Matão 1010 - Cidade Universitária
CEP 05508-090 - São Paulo - SP - Brazil
e-mail: florencia@usp.br*

Abstract: We find upper bounds for the probability of underestimation and overestimation errors in penalized likelihood context tree estimation. The bounds are explicit and applies to processes of not necessarily finite memory. We allow for general penalizing terms and we give conditions over the maximal depth of the estimated trees in order to get strongly consistent estimates. This generalizes previous results obtained in the case of estimation of the order of a Markov chain.

AMS 2000 subject classifications: Primary 60G10; secondary 62M09.
Keywords and phrases: context tree, penalized maximum likelihood estimation, Bayesian Information Criterion (BIC), rate of convergence.

1. Introduction

In this paper we obtain an exponential upper bound for the underestimation of the *context tree* of a variable memory process by penalized likelihood (PL) criteria and a sub-exponential upper bound for the overestimation event. Our result applies to processes of not necessarily finite memory that satisfies some continuity requirements, generalizing the bound obtained in Dorea and Zhao (2006) for the estimation of the order of a Markov chain by similar methods (EDC criterion).

The concept of context tree was first introduced by Rissanen (1983) to denote the minimum set of sequences that are necessary to predict the next symbol in a finite memory stochastic chain. A particular case of context tree is the set of all sequences of length k , representing a Markov chain of order k . For that reason, context trees allow a more detailed and parsimonious representation of processes than finite order Markov chains do.

In the statistical literature, the processes allowing a context tree representation are called Variable Length Markov Chains (Bühlmann and Wyner; 1999).

*This work is part of PRONEX/FAPESP's project *Stochastic behavior, critical phenomena and rhythmic pattern identification in natural languages* (grant number 03/09930-9), CNRS-FAPESP project *Probabilistic phonology of rhythm* and CNPq project *Rhythmic patterns, prosodic domains and probabilistic modeling in Portuguese Corpora* (grant number 485999/2007-2).

This class of models has shown to be useful in real data modeling, as for example, for the case of protein classification into families (Bejerano and Yona; 2001; Leonardi; 2006).

Historically, the estimation of the context tree of a process has been addressed by different versions of the algorithm *Context*, introduced by Rissanen in its seminal paper. This algorithm was proven to be weak consistent in the case of bounded memory (Bühlmann and Wyner; 1999) and also in the case of unbounded memory (Ferrari and Wyner; 2003; Duarte et al.; 2006). Recently, in Galves et al. (2008) it was obtained an upper bound for the rate of convergence of the algorithm *Context* in the case of bounded memory processes. A generalization of this result to the case of unbounded memory processes was given in Galves and Leonardi (2008).

The estimation of context trees by PL criteria had not been addressed in the literature until the recent work by Csiszár and Talata (2006). The reason for that was the exponential cost of the estimation, due to the number of trees that had to be considered in order to find the optimal one. In their article, Csiszár and Talata showed that the Bayesian Information Criterion (BIC), which is a particular case of the PL estimators (using a penalizing term growing logarithmically), is strongly consistent and can be computed in linear time, using a suitable version of the Context Tree Weighting method of Willems, Shtarkov and Tjalkens (Willems et al.; 1995; Willems; 1998). Their result applies to unbounded memory processes and the depth of the estimated tree is allowed to grow with the sample size as a sub-logarithmic function. This last condition was proven to be unnecessary in the case of finite memory processes, as proven in Garivier (2006). An explicit bound on the rate of convergence of the PL context tree estimators had remained until now as an open question.

The paper is organized as follows. In Section 2 we introduce some definitions and state the main result. In Section 3 we present the proofs and in Section 4 we do some final remarks. Finally, Section 5 constitutes an appendix that contains some results needed in our proofs and obtained elsewhere in the literature.

2. Definitions and results

In what follows A will represent a finite alphabet of size $|A|$. Given two integers $m \leq n$, we will denote by w_m^n the sequence (w_m, \dots, w_n) of symbols in A . The length of the sequence w_m^n is denoted by $\ell(w_m^n)$ and is defined by $\ell(w_m^n) = n - m + 1$. Any sequence w_m^n with $m > n$ represents the empty string and is denoted by λ . The length of the empty string is $\ell(\lambda) = 0$. In the sequel A^j will denote the set of all sequences of length j over A .

Given two sequences $w = w_m^n$ and $v = v_j^k$, we will denote by vw the sequence of length $\ell(v) + \ell(w)$ obtained by concatenating the two strings. In particular, $\lambda w = w \lambda = w$. The concatenation of sequences is also extended to the case in which v denotes a semi-infinite sequence, that is $v = (\dots, v_{-2}, v_{-1})$, denoted by $v = v_{-\infty}^{-1}$.

We say that the sequence s is a *suffix* of the sequence w if there exists a

sequence u , with $\ell(u) \geq 1$, such that $w = us$. In this case we write $s \prec w$. When $s \prec w$ or $s = w$ we write $s \preceq w$.

Definition 2.1. A set \mathcal{T} of finite or semi-infinite sequences is a *tree* if no sequence $s \in \mathcal{T}$ is a suffix of another sequence $w \in \mathcal{T}$. This property is called the *suffix property*.

We define the *height* of the tree \mathcal{T} as

$$h(\mathcal{T}) = \sup\{\ell(w) : w \in \mathcal{T}\}.$$

In the case $h(\mathcal{T}) < +\infty$ we say that \mathcal{T} is *bounded* and we denote by $|\mathcal{T}|$ the number of sequences in \mathcal{T} . On the other hand, if $h(\mathcal{T}) = +\infty$ we say that the tree \mathcal{T} is *unbounded*.

Given a tree \mathcal{T} and an integer K we will denote by $\mathcal{T}|_K$ the tree \mathcal{T} *truncated* to level K , that is

$$\mathcal{T}|_K = \{w \in \mathcal{T} : \ell(w) \leq K\} \cup \{w : \ell(w) = K \text{ and } w \prec u, \text{ for some } u \in \mathcal{T}\}.$$

The expression $\text{Int}(\mathcal{T})$ will denote the set of all sequences that are suffixes of some $u \in \mathcal{T}$, that is

$$\text{Int}(\mathcal{T}) = \{w : w \prec u, \text{ for some } u \in \mathcal{T}\}.$$

We will say that a tree \mathcal{T} is *complete* if for every semi-infinite sequence w_∞^{-1} there exists a sequence $s \in \mathcal{T}$ such that $s \preceq w_\infty^{-1}$.

Consider a stationary ergodic stochastic chain $\{X_t : t \in \mathbb{Z}\}$ over A . Given a sequence $w \in A^j$ we denote by

$$p(w) = \mathbb{P}(X_1^j = w)$$

the stationary probability of the cylinder defined by the sequence w . If $p(w) > 0$ we write

$$p(a|w) = \mathbb{P}(X_0 = a | X_{-j}^{-1} = w).$$

In the sequel we will use the simpler notation X_t for the process $\{X_t : t \in \mathbb{Z}\}$.

Definition 2.2. A sequence $w \in A^j$ is a *context* for the process X_t if it satisfies

1. For any semi-infinite sequence x_∞^{-1} having w as a suffix

$$\mathbb{P}(X_0 = a | X_\infty^{-1} = x_\infty^{-1}) = p(a|w), \quad \text{for all } a \in A.$$

2. No suffix of w satisfies (1).

An *infinite context* is a semi-infinite sequence w_∞^{-1} such that any of its suffixes w_{-j}^{-1} , $j = 1, 2, \dots$ is a context.

Definition 2.2 implies that the set of all contexts (finite or infinite) satisfies the suffix property and hence it is a tree. This tree is called *the context tree* of the process X_t and will be denoted by \mathcal{T}_0 .

Remark 2.3. In this paper we will also consider i.i.d. processes. We will assume that these processes are compatible with a particular tree, given by the set $\{\lambda\}$.

Define the sequence $\{\alpha_k\}_{k \in \mathbb{N}}$ as

$$\begin{aligned}\alpha_0 &:= \inf_{w \in \mathcal{T}_0, a \in A} \{p(a|w)\}, \\ \alpha_k &:= \inf_{u \in A^k} \sum_{a \in A} \inf_{w \in \mathcal{T}_0, w \succ u} \{p(a|w)\}.\end{aligned}\tag{2.4}$$

Assumption 1. From now on we will assume the process X_t satisfies

1. $\alpha_0 > 0$ and
2. $\alpha := \sum_{k \in \mathbb{N}} (1 - \alpha_k) < +\infty$.

The positivity assumption over α_0 implies that the context tree of the process X_t is complete, i.e., any semi-infinite sequence $w_{-\infty}^{-1}$ belongs to \mathcal{T}_0 or has a suffix that belongs to \mathcal{T}_0 . The second assumption is related to the loss of memory of a process of infinite order. (see Galves and Leonardi (2008) for more details).

In what follows we will assume x_1, x_2, \dots, x_n is a sample of the process X_t . Let $d(n) < n$ be a function taking integer values and growing to infinity with n . This will denote the maximal height of the estimated context trees (and will be denoted simply by d). Then, given a sequence w , with $1 \leq \ell(w) \leq d$, and a symbol $a \in A$ we denote by $N_n(w, a)$ the number of occurrences of symbol a preceded by the sequence w , starting at $d+1$, that is,

$$N_n(w, a) = \sum_{t=d+1}^n \mathbf{1}\{x_{t-\ell(w)}^{t-1} = w, x_t = a\}.\tag{2.5}$$

On the other hand, $N_n(w)$ will denote the sum $\sum_{a \in A} N_n(w, a)$.

Definition 2.6. We will say that the tree \mathcal{T} is *feasible* if it is complete, $h(\mathcal{T}) \leq d$, $N_n(w) \geq 1$ for all $w \in \mathcal{T}$ and any string w' with $N_n(w') \geq 1$ either belongs to \mathcal{T} , is a suffix of some $w \in \mathcal{T}$ or has a suffix w that belongs to \mathcal{T} .

We will denote by $\mathcal{F}^d(x_1^n)$ the set of all feasible trees. Then, given a tree $\mathcal{T} \in \mathcal{F}^d(x_1^n)$, the maximum likelihood of the sequence x_1, \dots, x_n is given by

$$\hat{\mathbb{P}}_{\text{ML}, \mathcal{T}}(x_1^n) = \prod_{w \in \mathcal{T}} \prod_{a \in A} \hat{p}_n(a|w)^{N_n(w, a)},\tag{2.7}$$

where the empirical probabilities $\hat{p}_n(a|w)$ are given by

$$\hat{p}_n(a|w) = \frac{N_n(w, a)}{N_n(w)}.\tag{2.8}$$

Here and in the sequel we use the convention $0^0 = 1$, for example in the case of $N_n(w, a) = 0$ in expression 2.7. Note that by Definition 2.6, as $N_n(w) \geq 1$ for any $w \in \mathcal{T}$, it is not necessary to give an extra definition of $\hat{p}_n(a|w)$ in the case $N_n(w) = 0$.

Given a sequence w , with $N_n(w) \geq 1$, we will denote by

$$\hat{\mathbb{P}}_{\text{ML},w}(x_1^n) = \prod_{a \in A} \hat{p}_n(a|w)^{N_n(w,a)}.$$

Hence, we have

$$\hat{\mathbb{P}}_{\text{ML},\mathcal{T}}(x_1^n) = \prod_{w \in \mathcal{T}} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n).$$

Let $f(n)$ be any positive function such that $f(n) \rightarrow +\infty$, when $n \rightarrow +\infty$, and $n^{-1}f(n) \rightarrow 0$, when $n \rightarrow +\infty$. This function will represent the generic penalizing term of our estimator, replacing the function $\frac{|A|-1}{2} \log n$ in the classical definition of BIC (Csiszár and Talata; 2006). A function satisfying these conditions will be called *penalizing term*.

Definition 2.9. Given a penalizing term $f(n)$, the PL context tree estimator is given by

$$\hat{\mathcal{T}}(x_1^n) = \arg \min_{\mathcal{T} \in \mathcal{F}^d(x_1^n)} \{ -\log \hat{\mathbb{P}}_{\text{ML},\mathcal{T}}(x_1^n) + |\mathcal{T}|f(n) \}. \quad (2.10)$$

As can be seen, the computation of the estimated context tree using its raw definition would imply a search for the optimal tree on the set of all feasible trees. This was the biggest drawback of this approach, because the size of this set grows extremely fast as a function of the maximal height d . Fortunately, there is a way of computing the PL estimator without exploring the set of all trees, as shown by Csiszár and Talata (2006). The details of this algorithm are given in the Appendix and will be used in the proof of our main result.

Let $K \in \mathbb{N}$. Define the underestimation event with respect to the truncated tree $\mathcal{T}_0|_K$ by

$$U_n^K = \bigcup_{w \in \text{Int}(\mathcal{T}_0|_K)} \{ w \in \hat{\mathcal{T}}_n(x_1^n) \}$$

and the overestimation event by

$$O_n^K = \bigcup_{w \succ v \in \mathcal{T}_0, \ell(v) < K} \{ w \in \hat{\mathcal{T}}_n(x_1^n) \}.$$

We are ready to present the main result in this paper. It establishes upper bounds for the probability of occurrence of the underestimation and overestimation events.

Theorem 2.11. *Let x_1, x_2, \dots be a sample of the stationary ergodic stochastic process X_t having context tree \mathcal{T}_0 and satisfying Assumption 1. For any constant $K \in \mathbb{N}$ there exist an integer n_0 and positive constants c_1, c_2, c_3 and c_4 depending on the process X_t such that for any $n \geq n_0$*

- (a) $\mathbb{P}[U_n^K] \leq c_1 e^{-c_2(n-d)}$;
- (b) $\mathbb{P}[O_n^K] \leq c_3 |A|^d e^{-c_4 f(n)(\alpha_0^2/|A|)^d/d}$.

Corollary 2.12. *For any penalizing term $f(n)$ and any function $d(n)$ such that for any constant $c > 0$,*

$$\sum_{n \in \mathbb{N}} |A|^{d(n)} \exp\left[-\frac{f(n)c^{d(n)}}{d(n)}\right] < +\infty \quad (2.13)$$

we have that there exists an integer n_0 depending on the process X_t such that $\hat{\mathcal{T}}_n(x_1^n)|_K = \mathcal{T}_0|_K$ for any $n \geq n_0$.

3. Proof of Theorem 2.11

Using Definition 5.5 and Lemma 5.7 and we see that the tree in (2.10) can be written as

$$\hat{\mathcal{T}}(x_1^n) = \{w \in \cup_{j=1}^d A^j : \mathcal{X}_w(x_1^n) = 0, \mathcal{X}_v(x_1^n) = 1 \text{ for all } v \prec w\}$$

if $\mathcal{X}_\lambda(x_1^n) = 1$, and to $\{\lambda\}$ if $\mathcal{X}_\lambda(x_1^n) = 0$. Then, for n sufficiently large in order to guarantee that $\mathcal{T}_0|_K$ will be in $\mathcal{F}^d(x_1^n)$ we have that

$$U_n^K = \bigcup_{w \in \text{Int}(\mathcal{T}_0|_K)} \{\mathcal{X}_w(x_1^n) = 0\}$$

and

$$O_n^K \subset \bigcup_{v \in \mathcal{T}_0, \ell(v) < K} \{\mathcal{X}_v(x_1^n) = 1\}.$$

To prove (a) let $w \in \text{Int}(\mathcal{T}_0|_K)$, then using Definition 5.4 and Lemma 5.6 we have that

$$\mathbb{P}[\mathcal{X}_w(x_1^n) = 0] = \mathbb{P}\left[\prod_{a \in A} V_{aw}(x_1^n) \leq e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n)\right]$$

and for any $a \in A$

$$V_{aw}(x_1^n) = \max_{\mathcal{T} \in \mathcal{F}_{aw}^d(x_1^n)} \prod_{s \in \mathcal{T}} e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},s}(x_1^n),$$

where $\mathcal{F}_{aw}^d(x_1^n)$ is the set containing all trees \mathcal{T} that have the form $\mathcal{T} = \mathcal{T}' \cap \{u : u \succeq aw\}$, with $\mathcal{T}' \in \mathcal{F}^d(x_1^n)$. Then

$$\mathbb{P}[\mathcal{X}_w(x_1^n) = 0] = \mathbb{P}\left[\max_{\mathcal{T} \in \mathcal{F}_w^d(x_1^n)} \prod_{s \in \mathcal{T}} e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},s}(x_1^n) \leq e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n)\right].$$

For a tree $\mathcal{T} \in \mathcal{F}_w^d(x_1^n)$ define the quantity

$$\delta_{\mathcal{T}}(w) = \sum_{a \in A} \left[\sum_{u \in \mathcal{T}} p(ua) \log p(a|u) - p(wa) \log p(a|w) \right]. \quad (3.1)$$

Using Jensen's inequality we can see that $\delta_{\mathcal{T}}(w) > 0$ unless $p(a|w) = p(a|u)$ for all $a \in A$ and all $u \in \mathcal{T}$. Therefore, for a sufficiently large n there must be

a tree $\mathcal{T}'_w \in \mathcal{F}_w^d(x_1^n)$ such that $\delta_{\mathcal{T}'_w}(w) > 0$; if not we contradict the fact that $w \in \text{Int}(\mathcal{T}_0)$ and it is not a context in the sense of Definition 2.2. Therefore

$$\mathbb{P}[\mathcal{X}_w(x_1^n) = 0] \leq \mathbb{P}\left[\prod_{u \in \mathcal{T}'_w} e^{-f(n)\hat{\mathbb{P}}_{\text{ML},u}(x_1^n)} \leq e^{-f(n)\hat{\mathbb{P}}_{\text{ML},w}(x_1^n)}\right].$$

Now we can apply the logarithm function on both sides inside the probability obtaining that the right hand side equals

$$\mathbb{P}\left[\sum_{u \in \mathcal{T}'_w} \log \hat{\mathbb{P}}_{\text{ML},u}(x_1^n) - \log \hat{\mathbb{P}}_{\text{ML},w}(x_1^n) \leq (|\mathcal{T}'_w| - 1)f(n)\right].$$

Dividing by $n - d$ and subtracting on both sides the term $\delta_{\mathcal{T}'_w}(w)$ we have that for a sufficiently large n such that

$$\frac{|\mathcal{T}'_w|f(n)}{n - d} < \frac{\delta_{\mathcal{T}'_w}(w)}{2}$$

we can bound above the last expression by

$$\mathbb{P}\left[|L_n(w)| > \frac{\delta_{\mathcal{T}'_w}(w)}{4}\right] + \sum_{u \in \mathcal{T}'_w} \mathbb{P}\left[|L_n(u)| > \frac{\delta_{\mathcal{T}'_w}(w)}{4}\right],$$

where for any finite sequence s

$$L_n(s) = \sum_{a \in A} p(sa) \log p(a|s) - \frac{N_n(s, a)}{n - d} \log \hat{p}_n(a|s).$$

Using Corollary 5.9 we can bound above this expression by

$$3e^{\frac{1}{2}}|A|^2(1 + |\mathcal{T}'_w|) \exp\left[-\frac{(n - d) \min(\delta_{\mathcal{T}'_w}, \delta_{\mathcal{T}'_w}^2) \alpha_0^{2(h(\mathcal{T}'_w)+1)}}{1024e|A|^3(\alpha + \alpha_0) \log^2 \alpha_0 h(\mathcal{T}'_w)}\right].$$

We conclude the proof of part (a) by observing that we only have a finite number of sequences $w \in \text{Int}(\mathcal{T}_0|_K)$, so we can take

$$c_1 = \max_{w \in \text{Int}(\mathcal{T}_0|_K)} \{3e^{\frac{1}{2}}|A|^2(1 + |\mathcal{T}'_w|)\}$$

and

$$c_2 = \min_{w \in \text{Int}(\mathcal{T}_0|_K)} \left\{ \frac{\min(\delta_{\mathcal{T}'_w}, \delta_{\mathcal{T}'_w}^2) \alpha_0^{2(h(\mathcal{T}'_w)+1)}}{1024e|A|^3(\alpha + \alpha_0) \log^2 \alpha_0 h(\mathcal{T}'_w)} \right\}.$$

To prove part (b) observe that for any $w \in \mathcal{T}_0$ with $\ell(w) < K$

$$\mathbb{P}[\mathcal{X}_w(x_1^n) = 1] = \mathbb{P}\left[\prod_{a \in A} V_{aw}(x_1^n) > e^{-f(n)\hat{\mathbb{P}}_{\text{ML},w}(x_1^n)}\right]. \quad (3.2)$$

Using Lemma 5.6 we have that

$$\prod_{a \in A} V_{aw}(x_1^n) = \prod_{u \in \mathcal{T}_w(x_1^n)} e^{-f(n) \hat{\mathbb{P}}_{\text{ML},u}(x_1^n)}.$$

Then, applying the logarithm function the probability (3.2) is equal to

$$\begin{aligned} & \mathbb{P} \left[\sum_{u \in \mathcal{T}_w(x_1^n)} \log e^{-f(n) \hat{\mathbb{P}}_{\text{ML},u}(x_1^n)} > \log e^{-f(n) \hat{\mathbb{P}}_{\text{ML},w}(x_1^n)} \right] \\ &= \mathbb{P} \left[\log \hat{\mathbb{P}}_{\text{ML},w}(x_1^n) - \sum_{u \in \mathcal{T}_w(x_1^n)} \log \hat{\mathbb{P}}_{\text{ML},u}(x_1^n) < (1 - |\mathcal{T}_w(x_1^n)|)f(n) \right]. \end{aligned} \quad (3.3)$$

We know, by the maximum likelihood estimator of the transition probabilities that

$$\hat{\mathbb{P}}_{\text{ML},w}(x_1^n) \geq \prod_{a \in A} p(a|w)^{N_n(w,a)}. \quad (3.4)$$

Therefore, we can bound above the right hand side of (3.3) by

$$\begin{aligned} & \mathbb{P} \left[\sum_{a \in A} N_n(w,a) \log p(a|w) - \sum_{u \in \mathcal{T}_w(x_1^n)} \log \hat{\mathbb{P}}_{\text{ML},u}(x_1^n) < (1 - |\mathcal{T}_w(x_1^n)|)f(n) \right] \\ &= \mathbb{P} \left[\sum_{a \in A} \sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u,a) \log \frac{p(a|u)}{\hat{p}_n(a|u)} < (1 - |\mathcal{T}_w(x_1^n)|)f(n) \right]. \end{aligned}$$

This equality follows by substituting $N_n(w,a)$ by $\sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u,a)$ and the fact that $p(a|u) = p(a|w)$ for all $u \in \mathcal{T}_w(x_1^n)$, remembering that $w \in \mathcal{T}_0$. Observe that

$$\begin{aligned} \sum_{a \in A} \sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u,a) \log \frac{p(a|u)}{\hat{p}_n(a|u)} &= \sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u) \sum_{a \in A} \hat{p}_n(a|u) \log \frac{p(a|u)}{\hat{p}_n(a|u)} \\ &= - \sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u) D(\hat{p}_n(\cdot|u) \| p(\cdot|u)), \end{aligned}$$

where D is the *Kullback-Leibler divergence* between the two distributions $\hat{p}_n(\cdot|u)$ and $p(\cdot|u)$ (see the Appendix). Using Lemma 5.2 and dividing by $n-d$ we have that

$$\begin{aligned} & \mathbb{P} \left[- \sum_{u \in \mathcal{T}_w(x_1^n)} N_n(u) D(\hat{p}_n(\cdot|u) \| p(\cdot|u)) < (1 - |\mathcal{T}_w(x_1^n)|)f(n) \right] \\ & \leq \mathbb{P} \left[- \sum_{u \in \mathcal{T}_w(x_1^n)} \frac{N_n(u)}{n-d} \sum_{a \in A} \frac{[\hat{p}_n(a|u) - p(a|u)]^2}{p(a|u)} < \frac{(1 - |\mathcal{T}_w(x_1^n)|)f(n)}{n-d} \right]. \end{aligned}$$

As $\mathcal{X}_w(x_1^n) = 1$ it follows that $|\mathcal{T}_w(x_1^n)| > 1$. On the other hand, $N_n(u) \leq n-d$ and $f(n) > 0$. Therefore, we can bound above the right hand side of the last

expression by

$$\sum_{u \in \mathcal{T}_w(x_1^n)} \sum_{a \in A} \mathbb{P} \left[|\hat{p}_n(a|u) - p(a|u)| > \sqrt{\frac{f(n)p(a|u)}{(n-d)|A||\mathcal{T}_w(x_1^n)|}} \right].$$

Hence, using Corollary 5.9 we can bound above this expression by

$$2 e^{\frac{1}{e}} |A|^{d+2} \exp \left[- \frac{f(n) \alpha_0^{2(d+1)}}{32e(\alpha + \alpha_0) |A|^{d+3} d} \right].$$

This finishes the proof of Theorem 2.11, by taking

$$c_3 = 2e^{\frac{1}{e}} |A|^2 \quad \text{and} \quad c_4 = \frac{\alpha_0^2}{32e(\alpha + \alpha_0) |A|^3}.$$

Proof of Corollary 2.12. It follows from the Borel-Cantelli Lemma and Theorem 2.11, by noting that

$$\mathbb{P}[\hat{\mathcal{T}}(x_1^n)|_K \neq \mathcal{T}_0|_K] \leq \mathbb{P}[U_n^K] + \mathbb{P}[O_n^K]$$

and the right hand side is summable in n when condition (2.13) is satisfied.

4. Final Remarks

The present paper presents upper bounds for the rate of convergence of penalized likelihood context tree estimators. We obtain an exponential bound for the underestimation event and an under-exponential bound in the case of the overestimation event. These results generalizes the previous work by Dorea and Zhao (2006), who obtained similar bounds in the case of the estimation of the order of a Markov chain, using also penalized likelihood criteria. One question that still remains open is if these bounds are optimal, as in the case of an estimator introduced in Finesso et al. (1996) for the estimation of the order of a Markov chain. They prove that in the case of their estimator, the constant appearing in the underestimation bound is optimal, and that the overestimation bound can not be exponential if the estimator is universal, as in our case. The answer to these questions are important subjects for future work in this area.

5. Appendix

5.1. The context tree maximizing principle

The following definitions and results were taken from Csiszár and Talata (2006) and were included for completeness. Definitions 5.4 and 5.5 and Lemmas 5.6 and 5.7 were originally proven for the usual penalizing term $f(n) = \frac{|A|-1}{2} \log n$, but can be adapted in a straightforward way to our setting.

Given two probability distributions p and q over A , the *Kullback-Leibler divergence* is defined by

$$D(p\|q) = \sum_{a \in A} p(a) \log \frac{p(a)}{q(a)}, \quad (5.1)$$

where, by convention, $p(a) \log \frac{p(a)}{q(a)}$ equals 0 if $p(a) = 0$ and $+\infty$ if $p(a) > q(a) = 0$.

Lemma 5.2. *If p and q are two probability distributions over A then*

$$D(p\|q) \leq \sum_{a \in A} \frac{[p(a) - q(a)]^2}{q(a)}. \quad (5.3)$$

Proof. See Csiszár and Talata (2006, Lemma 6.3). \square

Consider the full tree A^d , and let S^d denote the set of all sequences of length at most d , that is $S^d = \cup_{j=0}^d A^j$.

Definition 5.4. Given a sequence $w \in S^d$ with $N_n(w) \geq 1$, we define recursively, starting from the sequences of the full tree A^d , the value

$$V_w(x_1^n) = \begin{cases} \max\{e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n), \prod_{a \in A} V_{aw}(x_1^n)\}, & \text{if } 0 \leq \ell(w) < d, \\ e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n), & \text{if } \ell(w) = d \end{cases}$$

and the indicator

$$\mathcal{X}_w(x_1^n) = \begin{cases} 1, & \text{if } 0 \leq \ell(w) < d \text{ and } \prod_{a \in A} V_{aw}(x_1^n) > e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n), \\ 0, & \text{if } 0 \leq \ell(w) < d \text{ and } \prod_{a \in A} V_{aw}(x_1^n) \leq e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},w}(x_1^n), \\ 0, & \text{if } \ell(w) = d. \end{cases}$$

Definition 5.5. Given $w \in S^d$ with $N_n(w) \geq 1$, the maximizing tree assign to the sequence w is the tree

$$\mathcal{T}_w(x_1^n) = \{u \in S^d : \mathcal{X}_u(x_1^n) = 0, \mathcal{X}_v(x_1^n) = 1 \text{ for all } w \preceq v \prec u\}$$

if $\mathcal{X}_w(x_1^n) = 1$ and $\mathcal{T}_w(x_1^n) = \{w\}$ if $\mathcal{X}_w(x_1^n) = 0$.

For a sequence $w \in S^d$, with $N_n(w) \geq 1$, define $\mathcal{F}_w^d(x_1^n)$ as the set containing all trees \mathcal{T} that have the form $\mathcal{T} = \mathcal{T}' \cap \{u : u \succeq w\}$, with $\mathcal{T}' \in \mathcal{F}^d(x_1^n)$.

Lemma 5.6. *For any $w \in S^d$ with $N_n(w) \geq 1$,*

$$V_w(x_1^n) = \max_{\mathcal{T} \in \mathcal{F}_w^d(x_1^n)} \prod_{u \in \mathcal{T}} e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},u}(x_1^n) = \prod_{u \in \mathcal{T}_w(x_1^n)} e^{-f(n)} \hat{\mathbb{P}}_{\text{ML},u}(x_1^n).$$

Proof. See Csiszár and Talata (2006, Lemma 4.4). \square

Lemma 5.7. *The context tree estimator $\hat{T}(x_1^n)$ in (2.10) equals the maximizing tree assigned to the empty string λ , that is,*

$$\hat{T}(x_1^n) = \mathcal{T}_\lambda(x_1^n).$$

Proof. See Csiszár and Talata (2006, Proposition 4.3). \square

From this result it follows that in order to obtain the tree maximizing the penalized maximum likelihood criteria it is sufficient to assign to each sequence $w \in S^d$, with $N_n(w) \geq 1$, the indicator $\mathcal{X}_w(x_1^n)$ and then to get the maximizing tree $\mathcal{T}_\lambda(x_1^n)$. The computational cost of this algorithm is linear in n if $d(n) = o(n)$, as proven by Csiszár and Talata (2006).

5.2. Exponential inequalities for empirical probabilities

The following result was proven in Galves and Leonardi (2008), we omit its proof here.

Theorem 5.8. *Assume the process X_t satisfies Assumption 1, then for any finite sequence w , any symbol $a \in A$ and any $t > 0$ the following inequality holds*

$$\mathbb{P}(|N_n(w, a) - (n-d)p(wa)| > t) \leq e^{\frac{1}{\varepsilon}} \exp\left[\frac{-t^2 C}{(n-d)\ell(wa)}\right],$$

where

$$C = \frac{\alpha_0}{8e(\alpha + \alpha_0)}.$$

As a consequence of Theorem 5.8 we obtain the following corollary.

Corollary 5.9. *For any finite sequence w , with $p(w) > 0$, any $t > 0$ and any sufficiently large n such that $N_n(w) \geq 1$ we have*

- (a) $\max_{a \in A} \mathbb{P}(|\hat{p}_n(a|w) - p(a|w)| > t) \leq 2e^{\frac{1}{\varepsilon}} |A| \exp\left[-\frac{(n-d)t^2 p(w)^2 \alpha_0}{32e|A|^2(\alpha + \alpha_0)\ell(wa)}\right];$
- (b) $\mathbb{P}[|L_n(w)| > t] \leq 3e^{\frac{1}{\varepsilon}} |A|^2 \exp\left[-\frac{(n-d)\min(t, t^2)p(w)^2 \alpha_0^2}{64e|A|^3(\alpha + \alpha_0)\log^2 \alpha_0 \ell(wa)}\right],$

$$\text{where } L_n(w) = \sum_{a \in A} p(wa) \log p(a|w) - \frac{N_n(w, a)}{n-d} \log \hat{p}_n(a|w).$$

Proof. To prove (a) observe that

$$p(a|w) = \frac{(n-d)p(wa)}{(n-d)p(w)}.$$

Then, summing and subtracting the term $\frac{N_n(w, a)}{(n-d)p(w)}$ we obtain

$$\begin{aligned} \left| \frac{N_n(w, a)}{N_n(w)} - \frac{(n-d)p(wa)}{(n-d)p(w)} \right| &\leq \frac{N_n(w, a)}{N_n(w)(n-d)p(w)} |(n-d)p(w) - N_n(w)| \\ &\quad + \frac{1}{(n-d)p(w)} |N_n(w, a) - (n-d)p(wa)|. \end{aligned}$$

Therefore, as $\frac{N_n(w,a)}{N_n(w)} \leq 1$ we have

$$\begin{aligned} \mathbb{P}(|\hat{p}_n(a|w) - p(a|w)| > t) &\leq \mathbb{P}\left(|(n-d)p(w) - N_n(w)| > \frac{t(n-d)p(w)}{2}\right) \\ &\quad + \mathbb{P}\left(|N_n(w,a) - (n-d)p(wa)| > \frac{t(n-d)p(w)}{2}\right) \end{aligned}$$

We can write $N_n(w) = \sum_{b \in A} N_n(w,b)$ and $p(w) = \sum_{b \in A} p(wb)$, then the right hand side of the last expression can be bounded above by the sum

$$\begin{aligned} \sum_{b \in A} \mathbb{P}\left(|N_n(w,b) - (n-d)p(wb)| > \frac{t(n-d)p(w)}{2|A|}\right) + \\ \mathbb{P}\left(|N_n(w,a) - (n-d)p(wa)| > \frac{t(n-d)p(w)}{2}\right). \end{aligned}$$

Using Theorem 5.8 we can bound above this expression by

$$e^{\frac{1}{e}} (|A| + 1) \exp\left[-(n-d) \frac{t^2 p(w)^2 C}{4|A|^2 \ell(wa)}\right].$$

This finishes the proof of (a). To prove (b) observe that

$$\begin{aligned} \mathbb{P}[|L_n(w)| > t] &\leq \mathbb{P}\left[\left|\sum_{a \in A} \log p(a|w)(p(wa) - \frac{N_n(w,a)}{n-d})\right| > \frac{t}{2}\right] \\ &\quad + \mathbb{P}\left[\left|\sum_{a \in A} \frac{N_n(w,a)}{n-d} \log \frac{p(a|w)}{\hat{p}_n(a|w)}\right| > \frac{t}{2}\right]. \end{aligned}$$

Using Theorem 5.8 we have that

$$\begin{aligned} \mathbb{P}\left[\left|\sum_{a \in A} \log p(a|w)(p(wa) - \frac{N_n(w,a)}{n-d})\right| > \frac{t}{2}\right] \\ \leq \sum_{a \in A} \mathbb{P}\left[|N_n(w,a) - (n-d)p(wa)| > \frac{(n-d)t}{2|\log p(a|w)||A|}\right] \\ \leq e^{\frac{1}{e}} |A| \exp\left[\frac{-(n-d)t^2 C}{4|A|^2 \log^2 \alpha_0 \ell(wa)}\right]. \end{aligned} \quad (5.10)$$

On the other hand, using the definition of the *Kullback-Leibler divergence*, Lemma 5.2 and part (a) of this Corollary we obtain

$$\begin{aligned} \mathbb{P}\left[\left|\sum_{a \in A} \frac{N_n(w,a)}{n-d} \log \frac{p(a|w)}{\hat{p}_n(a|w)}\right| > \frac{t}{2}\right] &\leq \mathbb{P}[D(\hat{p}(\cdot|w)||p(\cdot|w)) > \frac{t}{2}] \\ &\leq \sum_{a \in A} \mathbb{P}\left[|p(a|w) - \hat{p}_n(a|w)| > \sqrt{\frac{tp(a|w)}{2|A|}}\right] \\ &\leq 2e^{\frac{1}{e}} |A|^2 \exp\left[-\frac{(n-d)tp(w)^2 \alpha_0^2}{64e|A|^3(\alpha + \alpha_0)\ell(wa)}\right]. \end{aligned} \quad (5.11)$$

Summing (5.10) and (5.11) we obtain the bound in part (b) and we conclude the proof of Corollary 5.9. \square

Acknowledgments

The author is thankful to Antonio Galves, Aurélien Garivier, Eric Moulines and Bernard Prum for interesting suggestions to improve the presentation of the results in this paper.

References

- Bejerano, G. and Yona, G. (2001). Variations on probabilistic suffix trees: statistical modeling and prediction of protein families, *Bioinformatics* **17**(1): 23–43.
- Bühlmann, P. and Wyner, A. J. (1999). Variable length Markov chains, *Ann. Statist.* **27**: 480–513.
- Csiszár, I. and Talata, Z. (2006). Context tree estimation for not necessarily finite memory processes, via BIC and MDL, *IEEE Trans. Inform. Theory* **52**(3): 1007–1016.
- Dorea, C. and Zhao, L. (2006). Exponential bounds for the probability of wrong determination of the order of a markov chain by using the edc criterion, *J. Stat. Plan. Infer.* **136**: 3689–3697.
- Duarte, D., Galves, A. and Garcia, N. (2006). Markov approximation and consistent estimation of unbounded probabilistic suffix trees, *Bull. Braz. Math. Soc.* **37**(4): 581–592.
- Ferrari, F. and Wyner, A. (2003). Estimation of general stationary processes by variable length Markov chains, *Scand. J. Statist.* **30**(3): 459–480.
- Finesso, L., Liu, C.-C. and Narayan, P. (1996). The optimal error exponent for Markov order estimation, *IEEE Trans. Inform. Theory* **42**(5): 1488–1497.
- Galves, A. and Leonardi, F. (2008). *Exponential inequalities for empirical unbounded context trees*, Vol. 60 of *Progress in Probability*, Birkhauser, pp. 257–270.
- Galves, A., Maume-Deschamps, V. and Schmitt, B. (2008). Exponential inequalities for VLMC empirical trees, *ESAIM Probab. Stat* **12**: 43–45.
- Garivier, A. (2006). Consistency of the unlimited BIC context tree estimator, *IEEE Trans. Inform. Theory* **52**(10): 4630–4635.
- Leonardi, F. (2006). A generalization of the PST algorithm: modeling the sparse nature of protein sequences, *Bioinformatics* **22**(11): 1302–1307.
- Rissanen, J. (1983). A universal data compression system, *IEEE Trans. Inform. Theory* **29**(5): 656–664.
- Willems, F. (1998). The context-tree weighting method: Extensions, *IEEE Trans. Inf. Theory* **44**(2): 792–798.
- Willems, F., Shtarkov, Y. and Tjalkens, T. (1995). The context-tree weighting method: Basic properties, *IEEE Trans. Inf. Theory* **41**(3): 653–664.