



Stochastic chains with memory of variable length

Antonio Galves, Eva Loecherbach

► To cite this version:

Antonio Galves, Eva Loecherbach. Stochastic chains with memory of variable length. Festschrift in honour of the 75th birthday of Jorma Rissanen, International Statistical Review, pp.329-463, 2008. hal-00798528

HAL Id: hal-00798528

<https://hal.science/hal-00798528>

Submitted on 11 Mar 2013

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.

STOCHASTIC CHAINS WITH MEMORY OF VARIABLE LENGTH

A. GALVES AND E. LÖCHERBACH

Dedicated to Jorma Rissanen on his 75'th birthday

ABSTRACT. Stochastic chains with memory of variable length constitute an interesting family of stochastic chains of infinite order on a finite alphabet. The idea is that for each past, only a finite suffix of the past, called *context*, is enough to predict the next symbol. These models were first introduced in the information theory literature by Rissanen (1983) as a universal tool to perform data compression. Recently, they have been used to model up scientific data in areas as different as biology, linguistics and music. This paper presents a personal introductory guide to this class of models focusing on the algorithm Context and its rate of convergence.

1. INTRODUCTION

Chains with memory of variable length appear in Rissanen's 1983 paper called *A universal data compression system*. His idea was to model a string of symbols as a realization of a stochastic chain where the length of the memory needed to predict the next symbol is not fixed, but is a deterministic function of the string of the past symbols.

Considering a memory of variable length is a practical way to overcome the well known difficulty of the exponentially growing number of parameters which are needed to describe a Markov chain when its order increases. However if one wants to fit accurately complex data using a Markov chain of fixed order, one has to use a very high order. And this means that to estimate the parameters of the model we need huge samples, which makes this approach unsuitable for many practical issues.

Date: January 7, 2008.

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), CNPq's project *Stochastic modeling of speech* (grant number 475177/2004-5) and CNRS-FAPESP project *Probabilistic phonology of rhythm*. AG is partially supported by a CNPq fellowship (grant 308656/2005-9).

It turns out that in many important scientific data, the length of the relevant portion of the past is not fixed, on the contrary it depends on the past. For instance, in molecular biology, the translation of a gene into a protein is initiated by a fixed specific sequence of nucleotide bases called *start codon*. In other words, the start codon designs the end of the relevant portion of the past to be considered in the translation.

The same phenomenon appears in other scientific domains. For instance in linguistics, both in phonology and in syntax, there is the notion of domains in which the grammar operates to define admissible strings of forthcoming symbols. In other terms, the boundary of the linguistic domain defines the relevant part of the past for the processing of the next linguistic units. Rissanen's ingenious idea was to construct a stochastic model that generalizes this notion of relevant domain to any kind of symbolic strings.

To be more precise, Rissanen (1983) called *context* the relevant part of the past. The stochastic model is defined by the set of all contexts and an associated family of transition probabilities.

Models with memory of variable length are not only less expensive than the classical fixed order Markov chains, but also much more clever since they take into account the structural dependencies present in the data. This is precisely what the set of contexts expresses.

Rissanen has introduced models having memory of variable length as a universal system of data compression. His goal was to compress in real time a string of symbols generated by an unknown source. To do this, we have to estimate at each step the length of the context of the string observed until that time step, as well as the associated transition probabilities.

If we knew the contexts, then the estimation of the associated transition probabilities could be done using a classical procedure such as maximum likelihood estimation. Therefore, the main point is to put hands on the context length. In his seminal 1983 paper, Rissanen solved this problem by introducing the algorithm *Context*. This algorithm estimates in a consistent way both the length of the context as well as the associated transition probability.

The class of models with memory of variable length raises interesting questions from the point of view of statistics. Examples are the rate of convergence and the fluctuations

of the algorithm Context and other estimators of the model. Another challenging question would be how to produce a robust version of the algorithm Context.

But also from the point of view of probability theory, this class of models is interesting. In effect, if the length of the contexts is not bounded, then chains with memory of variable length are chains of infinite order. Existence, uniqueness, phase-transitions, perfect simulation are deep mathematical questions that should be addressed to in this new and challenging class of models.

Last but not least, models of variable length revealed to be very performing tools in applied statistics, by achieving in an efficient way classification tasks in proteomics, genomics, linguistics, classification of musical styles, and much more.

In what follows we present a personal introductory guide to this class of models with no attempt to give a complete survey of the subject. We will mainly focus on the algorithm Context and present some recent results obtained by our research team.

2. PROBABILISTIC CONTEXT TREES

In what follows A will represent a finite alphabet of size $|A|$. Given two integers $m \leq n$, we will denote by x_m^n the sequence (x_m, \dots, x_n) of symbols in A . Let A_+^* be the set of all finite sequences, that is

$$A_+^* = \bigcup_{k=1}^{\infty} A^{\{-k, \dots, -1\}}.$$

We shall write $\underline{A} = A^{\{\dots, -n, \dots, -2, -1\}}$, and denote by $x_{-\infty}^{-1}$ any element of \underline{A} .

Our main object of interest is what we shall call *context length function*.

Definition 2.1. A *context length function* l is a function $l : A_+^* \rightarrow \{1, 2, \dots\} \cup \{\infty\}$ satisfying the following two properties.

(i) For any $k \geq 1$, for any $x_{-k}^{-1} \in A_+^*$, we have

$$l(x_{-k}^{-1}) \in \{1, \dots, k\} \cup \{+\infty\}.$$

(ii) For any $x_{-\infty}^{-1} \in \underline{A}$, if $l(x_{-k}^{-1}) = k$ for some $k \geq 1$, then

$$\begin{aligned} l(x_{-i}^{-1}) &= \infty, & \text{for any } i < k \\ l(x_{-i}^{-1}) &= k, & \text{for any } i > k. \end{aligned}$$

Intuitively, given a sequence $x_{-\infty}^{-1}$, the function l tells us, at which position in the past we can stop since we have reached the end of the context. The first condition is a kind of adaptivity condition. It tells us that we can decide whether the end of the context has already been reached at step k just by inspecting the past sequence up to that step. If l equals $+\infty$, we have to look further back in the past. The second condition is a consistency condition. It tells us that once we have reached the bound of the context, we do not have to look further back in the past, and that the context of a longer sequence x_{-i}^{-1} , $i > k$, is also the context of x_{-k}^{-1} . In other terms, once the identification of the context is made at a given step k , this decision will not be changed by any further data present in the past before k .

By abuse of notation, we shall also call l the natural extension of the context length function to \underline{A} , given by

$$l(x_{-\infty}^{-1}) = \inf \{k \geq 1 : l(x_{-k}^{-1}) < +\infty\},$$

with the convention that $\inf \emptyset = +\infty$.

Definition 2.2. For any $x_{-\infty}^{-1} \in \underline{A}$, we shall call $x_{-l(x_{-\infty}^{-1})}^{-1}$ the *context* associated to l of the infinite sequence $x_{-\infty}^{-1}$.

Definition 2.3. Let l be a given context length function. A stationary stochastic chain $(X_n)_{n \in \mathbb{Z}}$ taking values in A is a chain having memory of variable length, if for any infinite past $x_{-\infty}^{-1} \in \underline{A}$ and any symbol $a \in A$, we have

$$P(X_0 = a \mid X_{-\infty}^{-1} = x_{-\infty}^{-1}) = P\left(X_0 = a \mid X_{-l(x_{-\infty}^{-1})}^{-1} = x_{-l(x_{-\infty}^{-1})}^{-1}\right). \quad (2.4)$$

We shall use the short hand notation

$$p(a|x_{-k}^{-1}) = P(X_0 = a \mid X_{-k}^{-1} = x_{-k}^{-1}). \quad (2.5)$$

We are mainly interested in those values of $p(a|x_{-k}^{-1})$ where $k = l(x_{-\infty}^{-1})$.

Observe that the set $\{\ell(X_{-\infty}^{-1}) = k\}$ is measurable with respect to the σ -algebra generated by X_{-k}^{-1} . Thus we have

Proposition 2.6. *Let (X_n) be a stationary chain as in definition 2.3, having context length function l . Put $\mathcal{F}_k = \sigma\{X_{-k}, \dots, X_{-1}\}, k \geq 1$. Then $l(X_{-\infty}^{-1})$ is a $(\mathcal{F}_k)_k$ -stopping time.*

Given a context length function l , we define an associated countable subset $\tau \subset A_+^*$ by

$$\tau = \tau^l = \{x_{-k}^{-1} : k = l(x_{-k}^{-1}), k \geq 1\}.$$

To simplify notation, we will denote by \underline{x} and \underline{y} generic elements of τ .

Definition 2.7. Given a finite sequence x_{-k}^{-1} , we shall call *suffix* of x_{-k}^{-1} each string x_{-j}^{-1} with $j \leq k$. If $j < k$ we call x_{-j}^{-1} *proper suffix* of x_{-k}^{-1} . Now let $\mathcal{S} \subset A_+^*$. We say that \mathcal{S} satisfies the *suffix property* if for no $x_{-k}^{-1} \in \mathcal{S}$, there exists a proper suffix $x_{-j}^{-1} \in \mathcal{S}$ of x_{-k}^{-1} .

The following proposition follows immediately from property (ii) of definition 2.1.

Proposition 2.8. *Given a context length function l , the associated set τ^l satisfies the suffix property.*

As a consequence, τ can be identified with the set of leaves of a rooted tree with a countable set of finite labeled branches.

Definition 2.9. We call *probabilistic context tree on A* the ordered pair (τ, p) , where

$$p = \{p(\cdot | \underline{x}), \underline{x} \in \tau\}$$

is the family of transition probabilities of (2.5). We say that the probabilistic context tree (τ, p) is unbounded if the function l is unbounded.

Definition 2.10. Let $(X_n)_{n \in \mathbb{Z}}$ be a stationary chain and let (τ, p) be a probabilistic context tree. We shall say that $(X_n)_{n \in \mathbb{Z}}$ is *compatible* with (τ, p) , if (2.5) holds for all $\underline{x} \in \tau$.

In order to illustrate these mathematical concepts, let us consider the following example.

Example 2.11. *Consider a two-symbol alphabet $A = \{0, 1\}$ and the following context length function*

$$l(x_{-\infty}^{-1}) = \inf \{k : x_{-k} = 1\}.$$

Then the associated tree τ is given by

$$\tau = \{10^k, k \geq 0\},$$

where 10^k represents the sequence $(x_{-k-1}, x_{-k}, \dots, x_{-1})$ such that $x_{-i} = 0$ for all $1 \leq i \leq k$ and $x_{-k-1} = 1$.

The associated transition probabilities are defined by

$$P(X_0 = 1 | X_{-1} = \dots = X_{-k} = 0, X_{-k-1} = 1) = q_k, k \geq 0,$$

with $0 \leq q_k \leq 1$.

Clearly, the stochastic chain associated to this context length function l is a chain of infinite order. This raises the mathematical question of existence of such a process. It is straightforward to see that the following proposition holds true.

Proposition 2.12. *Suppose that $X_0 = 1$. Put $T_1 = \inf\{k \geq 1 : X_k = 1\}$. A necessary and sufficient condition for $T_1 < +\infty$ almost surely is*

$$\sum_{k \geq 0} q_k = +\infty. \quad (2.13)$$

This means that if (2.13) is satisfied, then – provided the chain starts from 1 – almost surely there will be appearance of an infinite number of the symbol 1. This implies that there exists a non-trivial stationary chain associated to this probabilistic context tree.

Observe that the process X_n is actually a renewal process with the renewal times defined as follows.

$$T_0 = \sup\{n < 0 : X_n = 1\},$$

and for $k \geq 1$,

$$T_k = \inf\{n > T_{k-1} : X_n = 1\} \quad \text{and} \quad T_{-k} := \sup\{n < T_{-(k-1)} : X_n = 1\}.$$

This example shows clearly that the tree of contexts defines a partition of all possible pasts with the exception of the single string composed of all symbols identical to 0. The condition (2.13) shows that it is possible to construct the chain taking values in the set of all sequences having an infinite number of symbols 1 both to the left and to the right of the origin.

However, we could also include this exceptional string to the set of possible contexts by defining an extra parameter q_∞ . This is the choice of Csiszár and Talata (2006). If q_∞ is strictly positive, then condition (2.13) implies that after a finite time, there will appearance of the symbol 1, even if we start with an infinity of symbols 0. In other terms, this exceptional string does not have to be considered if we are interested in the stationary regime of the chain.

In case that $q_\infty = 0$ and (2.13) holds, we have the phenomenon of phase transition. One of the phases is composed of only one string having only the symbol 0.

The renewal process is an interesting example of a chain having memory of unbounded variable length. In the case where the probabilistic context tree is bounded, the corresponding chain is in fact a Markov chain whose order is equal to the maximal context length. However, the tree of contexts provides interesting additional information concerning the dependencies in the data and the structure of the chain. This raises the issue how to estimate the context tree out of the data. This was originally solved in Rissanen's 1983 paper using the algorithm Context.

At this point it is important to discuss the following minimality issue. Among all possible context trees fitting the data, we want of course to identify the smallest one. This is the tree corresponding to the smallest context length function. More precisely, if l and l' are context length functions, we shall say that $l \leq l'$ if $l(x_{-\infty}^{-1}) \leq l'(x_{-\infty}^{-1})$ for any string $x_{-\infty}^{-1} \in \underline{A}$. From now on we shall call *context* of a string $x_{-\infty}^{-1}$ the context associated to the minimal context length function. Estimating this minimal context is precisely the goal of the algorithm Context.

3. THE ALGORITHM CONTEXT

We now present the algorithm Context introduced by Rissanen (1983). The goal of the algorithm is to estimate adaptively the context of the next symbol X_n given the past symbols X_0^{n-1} . The way the algorithm Context works can be summarized as follows. Given a sample produced by a chain with variable memory, we start with a maximal tree of candidate contexts for the sample. The branches of this first tree are then pruned starting from the leaves towards the root until we obtain a minimal tree of contexts well adapted to the sample. We associate to each context an estimated probability transition

defined as the proportion of time the context appears in the sample followed by each one of the symbols in the alphabet. We stop pruning once the gain function exceeds a given threshold.

Let X_0, X_1, \dots, X_{n-1} be a sample from the finite probabilistic tree (τ, p) . For any finite string x_{-j}^{-1} with $j \leq n$, we denote $N_n(x_{-j}^{-1})$ the number of occurrences of the string in the sample

$$N_n(x_{-j}^{-1}) = \sum_{t=0}^{n-j} \mathbf{1} \left\{ X_t^{t+j-1} = x_{-j}^{-1} \right\}. \quad (3.1)$$

Rissanen first constructs a maximal candidate context $X_{n-M(n)}^{n-1}$ where $M(n)$ is a random length defined as follows

$$M(n) = \min \left\{ i = 0, 1, \dots, \lfloor C_1 \log n \rfloor : N_n(X_{n-i}^{n-1}) > \frac{C_2 n}{\sqrt{\log n}} \right\}. \quad (3.2)$$

Here C_1 and C_2 are arbitrary positive constants. In the case the set is empty we take $M(n) = 0$.

Rissanen then shortens this maximal candidate context by successively pruning the branches according to a sequence of tests based on the likelihood ratio statistics. This is formally done as follows.

If $\sum_{b \in A} N_n(x_{-k}^{-1}b) > 0$, define the estimator of the transition probability p by

$$\hat{p}_n(a|x_{-k}^{-1}) = \frac{N_n(x_{-k}^{-1}a)}{\sum_{b \in A} N_n(x_{-k}^{-1}b)} \quad (3.3)$$

where $x_{-j}^{-1}a$ denotes the string $(x_{-j}, \dots, x_{-1}, a)$, obtained by concatenating x_{-j}^{-1} and the symbol a . If $\sum_{b \in A} N_n(x_{-k}^{-1}b) = 0$, define $\hat{p}_n(a|x_{-k}^{-1}) = 1/|A|$.

For $i \geq 1$ we define

$$\Lambda_n(x_{-i}^{-1}) = 2 \sum_{y \in A} \sum_{a \in A} N_n(yx_{-i}^{-1}a) \log \left[\frac{\hat{p}_n(a|x_{-i}^{-1}y)}{\hat{p}_n(a|x_{-i}^{-1})} \right], \quad (3.4)$$

where yx_{-i}^{-1} denotes the string $(y, x_{-i}, \dots, x_{-1})$, and where

$$\hat{p}_n(a|x_{-i}^{-1}y) = \frac{N_n(yx_{-i}^{-1}a)}{\sum_{b \in A} N_n(yx_{-i}^{-1}b)}.$$

Notice that $\Lambda_n(x_{-i}^{-1})$ is the log-likelihood ratio statistic for testing the consistency of the sample with a probabilistic suffix tree (τ, p) against the alternative that it is consistent with (τ', p') where τ and τ' differ only by one set of sibling nodes branching from x_{-i}^{-1} .

$\Lambda_n(x_{-i}^{-1})$ plays the role of a gain function telling us whether it is worth or not taking a next step further back in the past.

Rissanen then defines the length of the estimated current context $\hat{\ell}_n$ as

$$\hat{\ell}_n(X_0^{n-1}) = 1 + \max \{i = 1, \dots, M(n) - 1 : \Lambda_n(X_{n-i}^{n-1}) > C_2 \log n\}, \quad (3.5)$$

where C_2 is any positive constant.

Then, the result in Rissanen (1983) is the following.

Theorem 3.6. *Given a realization X_0, \dots, X_{n-1} of a probabilistic suffix tree (τ, p) with finite height, then*

$$P\left(\hat{\ell}_n(X_0^{n-1}) \neq \ell(X_0^{n-1})\right) \longrightarrow 0 \quad (3.7)$$

as $n \rightarrow \infty$.

Rissanen proves this result in a very short and elegant way. His starting point is the following upper bound.

$$\begin{aligned} P\left(\hat{\ell}_n(X_0^{n-1}) \neq \ell(X_0^{n-1})\right) &\leq \\ P\left(\hat{\ell}_n(X_0^{n-1}) \neq \ell(X_0^{n-1}) \mid N_n\left(X_{n-\ell(X_0^{n-1})}^{n-1}\right) > \frac{C_2 n}{\sqrt{\log n}}\right) &P\left(N_n\left(X_{n-\ell(X_0^{n-1})}^{n-1}\right) > \frac{C_2 n}{\sqrt{\log n}}\right) \\ &+ P\left(\bigcup_{w \in \tau} \left\{N_n(w) \leq \frac{C_2 n}{\sqrt{\log n}}\right\}\right). \end{aligned} \quad (3.8)$$

Then he provides the following explicit upper bound for the conditional probability in the right-hand side of (3.8)

$$P\left(\hat{\ell}_n(X_0^{n-1}) \neq \ell(X_0^{n-1}) \mid N_n\left(X_{n-\ell(X_0^{n-1})}^{n-1}\right) > \frac{C_2 n}{\sqrt{\log n}}\right) \leq C_1 \log n e^{-C'_2 \sqrt{\log n}}, \quad (3.9)$$

where C_1 , C_2 and C'_2 are positive constants independent of the maximum of the context length function.

With respect to the second term he only observes that, by ergodicity, for each $x_{-k}^{-1} \in \tau$ we have

$$P\left(N_n(x_{-k}^{-1}) \leq \frac{C_2 n}{\sqrt{\log n}}\right) \longrightarrow 0 \quad (3.10)$$

as $n \rightarrow \infty$. Since τ is finite the convergence in (3.10) implies the desired result.

4. THE UNBOUNDED CASE

In his original paper, Rissanen was only interested in the case of bounded context trees. However, from the mathematical point of view, it is interesting to consider also the case of unbounded probabilistic context trees corresponding to chains of infinite order. It can be argued that also from an applied point of view the unbounded case must be considered as noisy observation of Markov chains generically have infinite order memory.

The unbounded case raises immediately the preliminary question of existence and uniqueness of the corresponding chain. This issue can be addressed by adapting to probabilistic context trees the conditions for existence and uniqueness that have already been proved for infinite order chains. This is precisely what is done in the paper by Duarte et al. (2006) who adapt the type A condition presented in Fernández and Galves (2002) in the following way.

To simplify the presentation, let us introduce some extra notation. Recall that \underline{x} and \underline{y} denote generic elements of τ . Given $\underline{x} = x_{-i}^{-1}$ and $\underline{y} = y_{-j}^{-1}$, we shall write $\underline{x} \stackrel{k}{=} \underline{y}$ if and only if $k \leq \min\{i, j\}$ and $x_{-1} = y_{-1}, \dots, x_{-k} = y_{-k}$.

Definition 4.1. A probabilistic suffix tree (τ, p) on A is of *type A* if its transition probabilities p satisfy the following conditions.

- (1) **Weakly non-nullness**, that is

$$\sum_{a \in A} \inf_{\underline{x} \in \tau} p(a \mid \underline{x}) > 0; \quad (4.2)$$

- (2) **Continuity**, that is

$$\beta(k) = \max_{a \in A} \sup\{|p(a \mid \underline{x}) - p(a \mid \underline{y})|, \underline{y} \in \tau, \underline{x} \in \tau \text{ with } \underline{x} \stackrel{k}{=} \underline{y}\} \rightarrow 0 \quad (4.3)$$

as $k \rightarrow \infty$. We also define

$$\beta(0) = \max_{a \in A} \sup\{|p(a \mid \underline{x}) - p(a \mid \underline{y})|, \underline{y} \in \tau, \underline{x} \in \tau \text{ with } x_{-1} \neq y_{-1}\}.$$

The sequence $\{\beta(k)\}_k \in \mathbb{N}$ is called the **continuity rate**.

For a probabilistic suffix tree of type A with summable continuity rate, the maximal coupling argument used in Fernández and Galves (2002) implies the uniqueness of the law of the chain consistent with it.

We now present a slightly different version of the algorithm Context using the same gain function Λ_n but in which the length of the maximum context candidate is now deterministic and not more random. More precisely, we define the length of the biggest candidate context now as

$$k(n) = C_1 \log n \quad (4.4)$$

with a suitable positive constant C_1 .

The intuitive reason behind the choice of the upper bound length $C_1 \log n$ is the impossibility of estimating the probability of sequences of length much longer than $\log n$ based on a sample of length n . Recent versions of this fact can be found in Marton and Shields (1994, 1996) and Csiszár (2002).

Now, the definition of $\hat{\ell}_n$ is similar to the one in the original algorithm of Rissanen, that is

$$\hat{\ell}_n(X_0^{n-1}) = 1 + \max \{i = 1, \dots, k(n) - 1 : \Lambda_n(X_{n-i}^{n-1}) > C_2 \log n\} , \quad (4.5)$$

where C_2 is any positive constant.

The reason for taking the length of the maximum context candidate deterministic and no more random is to be able to use the classical results on the convergence of the law of $\Lambda_n(x_{-i}^{-1})$ to a chi-square distribution. However, we are not in a Markov setup since the probabilistic context tree is unbounded, and the chi-square approximation only works for Markov chains of fixed finite order.

To overcome this difficulty, we use the canonical Markov approximation of chains of infinite order presented in Fernández and Galves (2002) that we recall now by adapting the definitions and theorem to the framework of probabilistic context trees. The goal is to approximate a chain compatible with an unbounded probabilistic context tree by a sequence of chains compatible with bounded probabilistic context trees.

Definition 4.6. For all $k \geq 1$, the *canonical Markov approximation of order k* of a chain $(X_n)_{n \in \mathbb{Z}}$ is the chain with memory of variable length bounded by k compatible with the probabilistic context tree $(\tau^{[k]}, p^{[k]})$ where

$$\tau^{[k]} = \{\underline{x} \in \tau; l(\underline{x}) \leq k\} \cup \{x_{-k}^{-1}; \underline{x} \in \tau, l(\underline{x}) \geq k\} \quad (4.7)$$

for all $a \in A$, $\underline{x} \in \tau$, and where

$$p^{[k]}(a \mid x_{-j}^{-1}) := P(X_0 = a \mid X_{-j}^{-1} = x_{-j}^{-1}) \quad (4.8)$$

for all $x_{-j}^{-1} \in \tau^{[k]}$.

Observe that for contexts $\underline{x} \in \tau$ which length does not exceed k , we have $p^{[k]}(a \mid \underline{x}) = p(a \mid \underline{x})$. However, for sequences x_{-k}^{-1} which are internal nodes of τ , there is no easy explicit formula expressing $p^{[k]}(\cdot \mid x_{-k}^{-1})$ in terms of the family $\{p(\cdot \mid \underline{y}), \underline{y} \in \tau\}$.

The main result of Fernández and Galves (2002) that will be used in the proof of the consistency of the algorithm Context can be stated as follows.

Theorem 4.9. *Let $(X_n)_{n \in \mathbb{Z}}$ be a chain compatible with a type A probabilistic context tree (τ, p) with summable continuity rate, and let $(X_n^{[k]})_{n \in \mathbb{Z}}$ be its canonical Markov approximation of order k . Then there exists a coupling between $(X_n)_{n \in \mathbb{Z}}$ and $(X_n^{[k]})_{n \in \mathbb{Z}}$ and a constant $C > 0$ such that*

$$P(X_0 \neq X_0^{[k]}) \leq C\beta(k). \quad (4.10)$$

Using this result and the classical chi-square approximation for Markov chains, Duarte et al. (2006) proved the consistency of their version of the algorithm Context in the unbounded case and also provided an upper bound for the rate of convergence. Their result is the following.

Theorem 4.11. *Let X_0, X_2, \dots, X_{n-1} be a sample from a type A unbounded probabilistic suffix tree (τ, p) with continuity rate $\beta(j) \leq f(j) \exp\{-j\}$, with $f(j) \rightarrow 0$ as $j \rightarrow \infty$. Then, for any choice of positive constants C_1 and C_2 in (4.4) and (4.5), there exist positive constants C and D such that*

$$P(\hat{\ell}_n(X_0^{n-1}) \neq \ell(X_0^{n-1})) \leq C_1 \log n (n^{-C_2} + D/n) + Cf(C_1 \log n).$$

The proof can be sketched very easily. Take $k = k(n) = C_1 \log(n)$ and construct a coupled version of the processes $(X_t)_{t \in \mathbb{Z}}$ and $(X_t^{[k(n)]})_{t \in \mathbb{Z}}$. First of all notice that for

$$k = k(n),$$

$$\begin{aligned} P\left(\hat{\ell}_n(X_0, \dots, X_{n-1}) \neq \ell(X_0, \dots, X_{n-1})\right) \leq \\ P\left(\hat{\ell}_n(X_0^{[k]}, \dots, X_{n-1}^{[k]}) \neq \ell(X_0^{[k]}, \dots, X_{n-1}^{[k]})\right) + P\left(\bigcup_{i=1}^n \{X_i \neq X_i^{[k]}\}\right). \end{aligned} \quad (4.12)$$

Using the inequality (4.10) of Fernández and Galves (2002), the second term in (4.12) can be bounded above as

$$P\left(\bigcup_{i=1}^n \{X_i \neq X_i^{[k]}\}\right) \leq n C \beta(k(n)).$$

The first term in (4.12) can be treated using the classical chi-square approximation for the log-likelihood ratio test for Markov chains of fixed order k .

More precisely, we know that for fixed x_{-i}^{-1} , under the null hypothesis, the statistics $\Lambda_n(x_{-i}^{-1})$, given by (3.4), has asymptotically chi-square distribution with $|A| - 1$ degrees of freedom (see, for example, van der Vaart (1998)). We recall that, for each x_{-i}^{-1} the null hypothesis (H_0^i) is that the true context is x_{-i}^{-1} .

Since we are going to perform a sequence of $k(n)$ sequential tests where $k(n) \rightarrow \infty$ as n diverges, we need to control the error in the chi-square approximation. For this, we use a well-known asymptotic expansion for the distribution of $\Lambda_n(x_{-i}^{-1})$ due to Hayakawa (1977) which implies that

$$P\left(\Lambda_n(x_{-i}^{-1}) \leq t \mid H_0^i\right) = P\left(\chi^2 \leq t\right) + D/n, \quad (4.13)$$

where D is a positive constant and χ^2 is random variable with distribution chi-square with $|A| - 1$ degrees of freedom.

Therefore, it is immediate that

$$P\left(\Lambda_n(x_{-i}^{-1}) > C_2 \log n\right) \leq e^{-C_2 \log n} + D/n.$$

By the way we defined $\hat{\ell}_n$ in (4.5), in order to find $\hat{\ell}_n(X_0^{n-1})$ we have to perform at most $k(n)$ tests. We want to give an upper bound for the overall probability of type I error in a sequence of $k(n)$ sequential tests. An upper bound is given by the Bonferroni inequality, which in our case can be written as

$$P\left(\bigcup_{i=2}^{k(n)} \{\Lambda_n(x_{-i}^{-1}) > C_2 \log n\} \mid H_0^i\right) \leq \sum_{i=2}^{k(n)} P(\Lambda_n(x_{-i}^{-1}) > C_2 \log n \mid H_0^i).$$

This last term is bounded above by $C_1 \log n(n^{-C_2} + D/n)$. This concludes the proof.

Theorem 4.11 not only proves the consistency of the algorithm Context, but it also gives an upper bound for the rate of convergence. The estimation of the rate of convergence is crucial because it gives a bound on the minimum size of a sample required to guarantee, with a given probability, that the estimated tree is the good one. This is the issue we address to in the next section.

5. RATE OF CONVERGENCE OF THE ALGORITHM CONTEXT

Note that Rissanen's original theorem 3.6 as well as theorem 4.11 only show that all the contexts identified are true contexts with high probability. In other words, the estimated tree is a subtree of the true tree with high probability. In the case of bounded probabilistic context trees this missing point was handled with in Weinberger et al. (1995). This paper not only proves that the set of all contexts is reached, but also gives a bound for the rate of convergence.

More precisely, let us define the empirical tree

$$\hat{\tau}_n = \left\{ X_{j-\hat{\ell}_j(X_0^{j-1})}^{j-1} : j = n/2, \dots, n \right\}. \quad (5.1)$$

Actually, this is a slightly simplified version of the empirical tree defined in Weinberger et al. (1995). In particular, we are neglecting all the computational aspects considered there. But from the mathematical point of view, this definition perfectly does the job. Their convergence result is the following.

Theorem 5.2. *Let (τ, p) be a bounded probabilistic context tree and let X_0, \dots, X_n be compatible with (τ, p) . Then we have*

$$\sum_{n \geq 1} P(\hat{\tau}_n \neq \tau) \log n < +\infty.$$

In the unbounded case, this issue was treated without estimation of the rate of convergence in Ferrari and Wyner (2003) and including estimation of the rate of convergence in Galves and Leonardi (2008).

This last paper considers another slightly modified version of the algorithm Context using a different gain function, which has been introduced in Galves et al. (2007). More

precisely, let us define for any finite string $x_{-k}^{-1} \in A_+^*$ the gain function

$$\Delta_n(x_{-k}^{-1}) = \max_{a \in A} |\hat{p}_n(a|x_{-k}^{-1}) - \hat{p}_n(a|x_{-(k-1)}^{-1})|.$$

This gain function is well adapted to use exponential inequalities for the empirical transition probabilities in the pruning procedure rather than the chi-square approximation of the log-likelihood ratio as in theorems 3.6 and 4.11.

The theorem is stated in the following framework. Consider a stationary chain $(X_n)_{n \in \mathbb{Z}}$ compatible with an unbounded probabilistic context tree (τ, p) . For this chain, we define the sequence $(\alpha_n)_{n \in \mathbb{N}}$ by

$$\begin{aligned} \alpha_0 &= \sum_{a \in A} \inf_{\underline{x} \in \tau} p(a|\underline{x}), \\ \alpha_n &= \inf_{x_{-n}^{-1}} \sum_{a \in A} \inf_{\underline{y} \in \tau: l(\underline{y}) \geq n, \underline{y} = x_{-n}^{-1}} p(a|\underline{y}). \end{aligned}$$

We assume that the probabilistic context tree (τ, p) satisfies the condition (4.2) of weakly non-nullness, that is $\alpha_0 > 0$. We assume also the following summability condition

$$\alpha = \sum_{n \geq 0} (1 - \alpha_n) < +\infty. \quad (5.3)$$

Given a sequence $x_1^j = (x_1, \dots, x_j) \in A^j$ we denote by

$$p(x_1^j) = \mathbb{P}(X_1^j = x_1^j).$$

Then for an integer $m \geq 1$, we define

$$D_m = \min_{x_{-k}^{-1} \in \tau: k \leq m} \max_{a \in A} \{ |p(a|x_{-k}^{-1}) - p(a|x_{-(k-1)}^{-1})| \}, \quad (5.4)$$

and

$$\epsilon_m = \min \{ p(x_{-k}^{-1}) : k \leq m \text{ and } p(x_{-k}^{-1}) > 0 \}. \quad (5.5)$$

Intuitively, D_m tells us how distinctive is the difference between transition probabilities associated to the exact contexts and those associated to a string shorter one step in the past. We do not want to impose restrictions on the transition probabilities elsewhere than at the end of the branches of the context tree. This has to do with the pruning procedure which goes from the leaves to the root of the tree.

In the unbounded case, a natural way to state the convergence results is to consider truncated trees. The definition is the following. Given an integer K we will denote by $\tau|_K$ the tree τ *truncated* to level K , that is

$$\tau|_K = \{x_{-k}^{-1} \in \tau : k \leq K\} \cup \{x_{-K}^{-1} \text{ such that } x_k^{-1} \in \tau \text{ for some } k \geq K\}.$$

Actually, this is exactly the same tree which was called $\tau^{[K]}$ in (4.7). The notation $\tau|_K$ is more suitable for what follows.

The associated empirical tree of height k is defined in the following way.

Definition 5.6. Given $\delta > 0$ and $k < n$, the empirical tree is defined as

$$\hat{\tau}_n^{\delta,k} = \{x_{-r}^{-1}, 1 \leq r \leq k : \Delta_n(x_{-r}^{-1}) > \delta \wedge \Delta_n(y_{-(r+j)}^{-(r+1)} x_{-r}^{-1}) \leq \delta, \forall y_{-j}^{-(r+1)}, 1 \leq j \leq k-r\}.$$

In case $r = k$, the string $y_{-(r+j)}^{-(r+1)}$ is empty.

Note that in this definition, the parameter δ expresses the coarseness of the pruning criterion and k is the maximal length of the estimated contexts.

Now, Galves and Leonardi (2008) obtain the following result on the rate of convergence for the truncated context tree.

Theorem 5.7. *Let (τ, p) be a probabilistic context tree satisfying (4.2) and (5.3). Let X_0, \dots, X_n be a stationary stochastic chain compatible with (τ, p) . Then for any integer K , any k satisfying*

$$k > \max_{\underline{x} \in \tau|_K} \min \{\ell(\underline{y}) : \underline{y} \in \tau, \underline{x} \stackrel{K}{=} \underline{y}\}, \quad (5.8)$$

for any $\delta < D_k$ and for each

$$n > \frac{2(|A| + 1)}{\min(\delta, D_k - \delta)\epsilon_k} + k \quad (5.9)$$

we have that

$$P(\hat{\tau}_n^{\delta,k}|_K \neq \tau|_K) \leq 4e^{\frac{1}{e}} |A|^{k+2} \exp[-(n-k) \frac{[\min(\frac{\delta}{2}, \frac{D_k - \delta}{2}) - \frac{|A|+1}{(n-k)\epsilon_k}]^2 \epsilon_k^2 C}{4|A|^2(k+1)}],$$

where

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

In this theorem, the empirical trees have to be of height $k \geq K$ for the following reason. Truncating τ at level K implies that contexts longer than K are cut before reaching their end, and associated transition probabilities might not differ when comparing them at length K and $K-1$. That's why we consider the bigger empirical tree of height k satisfying condition (5.8). This guaranties that for each element \underline{x} of the truncated empirical tree there is at least one real context \underline{y} which has \underline{x} as its suffix.

As a consequence of theorem 5.7, Galves and Leonardi (2008) obtain the following strong consistency result.

Corollary 5.10. *Let (τ, p) be a probabilistic context tree satisfying the conditions of theorem 5.7. Then*

$$\hat{\tau}_n^{\delta, k}|_K = \tau|_K, \quad (5.11)$$

eventually almost surely, as $n \rightarrow \infty$.

The main ingredient of the proof of theorem 5.7 is an exponential upper bound for the deviations of the empirical transition probabilities. More precisely, Galves and Leonardi (2008) prove the following result.

Theorem 5.12. *For any finite sequence x_{-k}^{-1} with $p(x_{-k}^{-1}) > 0$, any symbol $a \in A$, any $t > 0$ and any $n > \frac{|A|+1}{tp(x_{-k}^{-1})} + k$ the following inequality holds.*

$$P(|\hat{p}_n(a|x_{-k}^{-1}) - p(a|x_{-k}^{-1})| > t) \leq 2|A|e^{\frac{1}{e}} \exp\left[-(n-k) \frac{[t - \frac{|A|+1}{(n-k)p(x_{-k}^{-1})}]^2 p(x_{-k}^{-1})^2 C}{4|A|^2(k+1)}\right], \quad (5.13)$$

where

$$C = \frac{\alpha_0}{8e(\alpha + \alpha_0)}. \quad (5.14)$$

The proof of this theorem is inspired by recent exponential upper bounds obtained by Dedecker and Doukhan (2003), Dedecker and Prieur (2005) and Maume-Deschamps (2006). It is based on the following loss-of-memory inequality of Comets et al. (2002).

Theorem 5.15. *Let $(X_n)_{n \in \mathbb{Z}}$ be a stationary stochastic chain compatible with the probabilistic context tree (τ, p) of theorem 5.7. Then, there exist a sequence $\{\rho_l\}_{l \in \mathbb{N}}$ such that*

for any $i \geq 1$, any $k > i$, any $j \geq 1$ and any finite sequence x_1^j , the following inequality holds

$$\sup_{x_1^i \in A^i} |P(X_k^{k+j-1} = x_1^j \mid X_1^i = y_1^i) - p(x_1^j)| \leq j \rho_{k-i-1}. \quad (5.16)$$

Moreover, the sequence $\{\rho_l\}_{l \in \mathbb{N}}$ is summable and

$$\sum_{l \in \mathbb{N}} \rho_l \leq 1 + \frac{2\alpha}{\alpha_0}.$$

Theorem 5.7 generalizes to the unbounded case previous results in Galves et al. (2008) for the case of bounded context trees. Note that the definition of the context tree estimator depends on the parameter δ , the same appearing in the constants of the exponential bound. To assure the consistency of the estimator we have to choose a δ sufficiently small, depending on the true probabilities of the process. The same thing happens to the parameter k . Therefore, this estimator is not universal, meaning that for fixed δ and k it fails to be consistent for all variable memory processes for which conditions (5.8) and (5.9) are not satisfied. We could try to overcome this difficulty by letting $\delta = \delta(n) \rightarrow 0$ and $k = k(n) \rightarrow +\infty$ as n increases. But doing this, we loose the exponential character of the upper bound. This could be considered as an illustration of the result in Finesso et al. (1996) who proved that in the simpler case of estimating the order of a Markov chain, it is not possible to have a universal estimator with exponential bounds for the probability of overestimation.

6. SOME FINAL COMMENTS AND BIBLIOGRAPHIC REMARKS

Chains with memory of variable length were introduced in the information theory literature by Rissanen (1983) as a universal system for data compression. Originally called by Rissanen *tree machine*, *tree source*, *context models*, etc., this class of models recently became popular in the statistics literature under the name of *Variable Length Markov Chains (VLMC)*, coined by Bühlmann and Wyner (1999).

Rissanen (1983) not only introduced the notion of variable memory models but he also introduced the algorithm Context to estimate the probabilistic context tree. From Rissanen (1983) to Galves et al. (2008), passing by Ron et al. (1996) and Bühlmann and Wyner (1999), several variants of the algorithm Context have been presented in the

literature. In all the variants the decision to prune a branch is taken by considering a *gain* function.

Rissanen (1983), Bühlmann and Wyner (1999) and Duarte et al. (2006) all defined the gain function in terms of the log likelihood ratio function. Rissanen (1983) proved the weak consistency of the algorithm Context in the case where the contexts have a bounded length. Bühlmann and Wyner (1999) proved the weak consistency of the algorithm also in the finite case without assuming a prior known bound on the maximal length of the memory but using a bound allowed to grow with the size of the sample.

A different gain function was introduced in Galves et al. (2008), considering differences between successive empirical transition probabilities and comparing them with a given threshold δ . An interesting consequence of the use of this different gain function was obtained by Collet et al. (2007). They proved that in the case of a binary alphabet and when taking δ within a suitable interval, it is possible to recover the context tree in the bounded case out from a noisy sample where each symbol can be flipped with small probability independently of the others.

The case of unbounded probabilistic context trees as far as we know was first considered by Ferrari and Wyner (2003) who also proved a weak consistency result for the algorithm Context in this more general setting. The unbounded case was also considered by Csiszár and Talata (2006) who introduced a different approach for the estimation of the probabilistic context tree using the Bayesian Information Criterion (BIC) as well as the Minimum Description Length Principle (MDL). We refer the reader to this last paper for a nice description of other approaches and results in this field, including the context tree maximizing algorithm by Willems et al. (1995). We also refer the reader to Garivier (2006a, b) for recent and elegant results on the BIC and the Context Tree Weighting Method (CTW). Garivier (2006c) is a very good presentation of models having memory of variable length, BIC, MDL, CTW and related issues in the framework of information theory.

With exception of Weinberger et al. (1995), the issue of the rate of convergence of the algorithm estimating the probabilistic context tree was not addressed in the literature until recently. Weinberger et al. (1995) proved in the bounded case that the probability that the estimated tree differs from the finite context tree is summable as a function of the

sample size. Assuming weaker hypotheses than Ferrari and Wyner (2003), Duarte et al. (2006) proved in the unbounded case that the probability of error decreases as the inverse of the sample size.

Leonardi (2007) obtained an upper bound for the rate of convergence of penalized likelihood context tree estimators. It showed that the estimated context tree truncated at any fixed height approximates the real truncated tree at a rate that decreases faster than the inverse of an exponential function of the penalizing term. The proof mixes the approaches of Galves et al. (2008) and Csiszár and Talata (2006).

Several interesting papers have recently addressed the question of classification of proteins and DNA sequences using models with memory of variable length, which in bio-informatics are often called prediction suffix trees (PST). Many of these papers have been written from a bio-informatics point of view focusing on the development of new tools rather than being concerned with mathematically rigorous proofs. The interested reader can find a starting point to this literature for instance in the papers by Bejerano et al. (2001), Bejerano and Yona (2001), Eskin et al. (2000), Leonardi (2006) and Miele et al. (2005). The same type of analysis has been used successfully to classification tasks in other domains like musicology (Lartillot et al. 2003), linguistics (Selding et al. 2001), etc.

This presentation did not intend to be exhaustive and the bibliography in many cases only gives a few hints about possible starting points to the literature. However, we think we have presented the state of the art concerning the rate of convergence of context tree estimators.

In the introduction we said that *Rissanen's ingenious idea was to construct a stochastic model that generalizes the notion of relevant domain (in biology or linguistics) to any kind of symbolic strings*. Actually, God only knows what Jorma had in mind when he invented this class of models. The French poet Paul Eluard wrote a book called *Les frères voyants*. This was the name given in the middle-age to people guiding blind persons. So maybe Rissanen acted as a *frère voyant* using his intuition to push mathematics and statistics into a challenging new direction.

REFERENCES

- [1] Bejerano, G., Seldin, Y., Margalit, H. and Tishby, N., Markovian domain fingerprinting: statistical segmentation of protein sequences, *Bioinformatics* **17**, Number 10, 927–934, 2001.
- [2] Bejerano, G. and Yona, G., Variations on probabilistic suffix trees: statistical modeling and prediction of protein families, *Bioinformatics* **17**, Number 1, 23–43, 2001.
- [3] Bühlmann, P. and Wyner, A. J., Variable length Markov chains, *Ann. Statist.* **27**, 480–513, 1999.
- [4] Collet, P., Galves, A. and Leonardi, F., Random perturbations of stochastic chains with unbounded variable length memory, Manuscript 2007, can be downloaded from ArXiv: math/0707.2796.
- [5] Comets, F., Fernández, R. and Ferrari, P., Processes with long memory: Regenerative construction and perfect simulation, *Ann. of Appl. Probab.* **12**, Number 3, 921–943, 2002.
- [6] Csiszár, I., Large-scale typicality of Markov sample paths and consistency of MDL order estimators. Special issue on Shannon theory: perspective, trends, and applications, *EEE Trans. Inform. Theory* **48**, Number 6, 1616–1628, 2002.
- [7] Csiszár, I. and Talata, Z., Context tree estimation for not necessarily finite memory processes, via BIC and MDL, *IEEE Trans. Inform. Theory* **52**, Number 3, 1007–1016, 2006.
- [8] Dedecker, J. and Doukhan, P., A new covariance inequality and applications, *Stochastic Process. Appl.*, **106**, Number 1, 63–80, 2003.
- [9] Dedecker, J. and Priour, C., New dependence coefficients. Examples and applications to statistics, *Probab. Theory Related Fields*, **132**, 203–236, 2005.
- [10] Duarte, D., Galves, A. and Garcia, N.L., Markov approximation and consistent estimation of unbounded probabilistic suffix trees, *Bull. Braz. Math. Soc.* **37**, Number 4, 581–592, 2006.
- [11] Eskin, E., Grundy, W. N., Singer, Y., Protein family classification using sparse Markov transducers, manuscript 2000, can be downloaded from <http://citeseer.ist.psu.edu/328658.html>
- [12] Fernández, R. and Galves, A., Markov approximations of chains of infinite order, Fifth Brazilian School in Probability (Ubatuba, 2001), *Bull. Braz. Math. Soc. (N.S.)* **33**, Number 3, 295–306, 2002.
- [13] Ferrari, F. and Wyner, A., Estimation of general stationary processes by variable length Markov chains, *Scand. J. Statist.* **30**, Number 3, 459–480, 2003.
- [14] Finesso, L., Liu, C.-C. and Narayan, P., The optimal error exponent for Markov order estimation, *IEEE Trans. Inform. Theory* **42**, Number 5, 1488–1497, 1996.
- [15] Galves, A. and Leonardi, F., Exponential inequalities for empirical unbounded context trees. Manuscript 2007, to appear in a special issue of “Progress in Probability”, Birkhäuser, eds. V. Sidoravicius and M. E. Vares, can be downloaded from ArXiv: math/0710.5900 .
- [16] Galves, A., Maume-Deschamps, V. and Schmitt, B., Exponential inequalities for VLMC empirical trees, *ESAIM. Probability and Statistics*, to appear 2008.
- [17] Garivier, A., Consistency of the unlimited BIC Context Tree estimator. *IEEE Trans. Inform. Theory* **52**, Number 10, 4630–4635, 2006.

- [18] Garivier, A., Redundancy of the context-tree weighting method on renewal and Markov renewal processes. *IEEE Trans. Inform. Theory* **52**, Number 12, 5579–5586, 2006.
- [19] Garivier, A., Modèles contextuels et alphabets infinis en théorie de l’information. PhD thesis, Université de Paris Sud, November 2006, can be downloaded from <http://www.math.u-psud.fr/~garivier/>
- [20] Hayakawa, T., The likelihood ratio criterion and the asymptotic expansion of its distribution, *Ann. Inst. Statist. Math.* **29**, Number 3, 359–378, 1977.
- [21] Lartillot, O., Dubnov, S., Assayag, G., Bejerano, G., A system for computer music generation by learning and improvisation in a particular style, *IEEE Computer J.* **36**, Number 10, 73–80, 2003.
- [22] Leonardi, F., A generalization of the PST algorithm: modeling the sparse nature of protein sequences, *Bioinformatics* **22**, Number 7, 1302–1307, 2006.
- [23] Leonardi, F., Rate of convergence of penalized likelihood context tree estimators. Manuscript 2007, can be downloaded from ArXiv: [math/0701810v2](http://arxiv.org/abs/math/0701810v2).
- [24] Marton, K. and Shields, P. C., Entropy and the consistent estimation of joint distributions, *Ann. Probab.* **22**, Number 2, 960–977, 1994.
- [25] Marton, K. and Shields, P. C., Correction: “Entropy and the consistent estimation of joint distributions” [*Ann. Probab.* **22** (1994), no. 2, 960–977; MR1288138 (95g:94004)], *Ann. Probab.* **24**, Number 1, 541–545, 1996.
- [26] Maume-Deschamps, V., *Exponential inequalities and estimation of conditional probabilities*, Lecture Notes in Statist. **187**, Springer 2006.
- [27] Miele, V., Bourguignon, P. Y., Robelin, D., Nuel, G., Richard, H., seq++: a package for biological sequences analysis with a range of Markov-related models. *Bioinformatics* **21**, Number 11, 2783–2874, 2005.
- [28] Rissanen, J., A universal data compression system, *IEEE Trans. Inform. Theory* **29**, Number 5, 656–664, 1983.
- [29] Ron, D., Singer, Y. and Tishby, N., The power of amnesia: learning probabilistic automata with variable memory length, *Machine Learning* **25**, 117–149, 1996.
- [30] Seldin, Y., Bejerano, G. and Tishby, N., Unsupervised sequence segmentation by a mixture of variable memory Markov models. Manuscript 2001, can be downloaded from <http://citeseer.ist.psu.edu/449505.html>
- [31] van der Vaart, A. W., *Asymptotic statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 1998.
- [32] Weinberger, M.J., Rissanen, J., Feder, M., A universal finite memory source, *IEEE Trans. Inform. Theory* **41**, Number 3, 643–652, 1995.
- [33] Willems, E. M. J., Shtarkov, Y. M. and Tjalkens, T. J., The Context-Tree weighting method: Basic properties, *IEEE Trans. Inform. Theory* **41**, 653–664, 1995.

Antonio Galves

Instituto de Matemática e Estatística

Universidade de São Paulo

BP 66281

05315-970 São Paulo, Brasil

e-mail: galves@ime.usp.br

Eva Löcherbach

Université de Paris 12 Val de Marne

LAMA – UMR CNRS 8050

61, Avenue du Général de Gaulle

94000 Créteil, France

e-mail: locherbach@univ-paris12.fr

UNIVERSIDADE DE SÃO PAULO AND UNIVERSITÉ DE PARIS 12 VAL DE MARNE