

On estimating the memory for finitarily Markovian processes

Gusztáv Morvai^{a,*}, Benjamin Weiss^b

^a *Research Group of the Hungarian Academy of Sciences, Budapest, 1521 Goldmann György tér 3, Hungary*

^b *Hebrew University of Jerusalem, Institute of Mathematics, Jerusalem 91904, Israel*

Received 24 January 2005; received in revised form 12 August 2005; accepted 8 November 2005

Available online 12 July 2006

Abstract

Finitarily Markovian processes are those processes $\{X_n\}_{n=-\infty}^{\infty}$ for which there is a finite K ($K = K(\{X_n\}_{n=-\infty}^0)$) such that the conditional distribution of X_1 given the entire past is equal to the conditional distribution of X_1 given only $\{X_n\}_{n=1-K}^0$. The least such value of K is called the memory length. We give a rather complete analysis of the problems of universally estimating the least such value of K , both in the backward sense that we have just described and in the forward sense, where one observes successive values of $\{X_n\}$ for $n \geq 0$ and asks for the least value K such that the conditional distribution of X_{n+1} given $\{X_i\}_{i=n-K+1}^n$ is the same as the conditional distribution of X_{n+1} given $\{X_i\}_{i=-\infty}^n$. We allow for finite or countably infinite alphabet size.

© 2006 Elsevier Masson SAS. All rights reserved.

Résumé

Les processus Markoviens finitaires sont des processus $\{X_n\}_{n=-\infty}^{\infty}$ pour lesquels il existe un entier K fini ($K = K(\{X_n\}_{n=-\infty}^0)$) telle que la distribution conditionnelle de X_1 étant donné tout le passé soit égale à la distribution conditionnelle de X_1 étant donné seulement $\{X_n\}_{n=1-K}^0$. La plus petite valeur d'un tel K est appelée la longueur de la mémoire. Nous donnons une analyse complète du problème de l'estimation de la plus petite de ces valeurs de K , aussi bien en remontant dans le passé qu'en allant vers le futur, c'est à dire quand on observe les valeurs successives de $\{X_n\}$ pour $n \geq 0$ et qu'on recherche la plus petite valeur de K telle que la distribution conditionnelle de X_{n+1} étant donné $\{X_i\}_{i=n-K+1}^n$ soit la même que la distribution conditionnelle de X_{n+1} étant donné $\{X_i\}_{i=-\infty}^n$. La taille des alphabets peut être choisie finie ou infinie.

© 2006 Elsevier Masson SAS. All rights reserved.

MSC: 62G05; 60G25; 60G10

Keywords: Nonparametric estimation; Stationary processes

1. Introduction

An important class of stationary ergodic processes that greatly extends the finite order Markov chains is the finitarily Markovian class. Informally, these are those processes $\{X_n\}_{n=-\infty}^{\infty}$ for which there is a finite K (that depends on the past $\{X_n\}$, $n \leq 0$) such that the conditional distribution of X_1 given the entire past is equal to the conditional distribution of X_1 given only $\{X_n\}$, $-K < n \leq 0$. When the process is a Markov chain of order L then one can simply

* Corresponding author.

E-mail addresses: morvai@szit.bme.hu (G. Morvai), weiss@math.huji.ac.il (B. Weiss).

take $K = L$ independent of the values that the process takes. However, even for such Markov chains, quite often a smaller value may exist for certain realizations of the process. Our main goal here is to give a rather complete analysis of the problems of universally estimating the least such value of K , both in the backward sense that we have just described and in the forward sense, where one observes successive values of $\{X_n\}$, $n \geq 0$. For the case of finite alphabet finite order Markov chains similar questions have been studied by Bühlmann and Wyner in [2]. However, the fact that we want to treat countable alphabets complicates matters significantly. The point is that while finite alphabet Markov chains have exponential rates of convergence of empirical distributions, for countable alphabet Markov chains no universal rates are available at all.

We encountered this problem in [16] where we gave a universal estimator for the order of a Markov chain on a countable state space, and some of the techniques that we use here have their origin in that paper. Before describing our results in more detail let us define more precisely the class of processes that we are considering.

First let us fix the notation. Let $\{X_n\}_{n=-\infty}^{\infty}$ be a stationary and ergodic time series taking values from a discrete (finite or countably infinite) alphabet \mathcal{X} . (Note that all stationary time series $\{X_n\}_{n=0}^{\infty}$ can be thought to be a two sided time series, that is, $\{X_n\}_{n=-\infty}^{\infty}$.) For notational convenience, let $X_m^n = (X_m, \dots, X_n)$, where $m \leq n$. Note that if $m > n$ then X_m^n is the empty string.

For convenience let $p(x_{-k}^0)$ and $p(y|x_{-k}^0)$ denote the distribution $P(X_{-k}^0 = x_{-k}^0)$ and the conditional distribution $P(X_1 = y | X_{-k}^0 = x_{-k}^0)$, respectively.

Definition 1. For a stationary time series $\{X_n\}$ the (random) length $K(X_{-\infty}^0)$ of the memory of the sample path $X_{-\infty}^0$ is the smallest possible $0 \leq K < \infty$ such that for all $i \geq 1$, all $y \in \mathcal{X}$, all $z_{-K-i+1}^{-K} \in \mathcal{X}^i$

$$p(y|X_{-K+1}^0) = p(y|z_{-K-i+1}^{-K}, X_{-K+1}^0)$$

provided $p(z_{-K-i+1}^{-K}, X_{-K+1}^0, y) > 0$, and $K(X_{-\infty}^0) = \infty$ if there is no such K .

Definition 2. The stationary time series $\{X_n\}$ is said to be finitarily Markovian if $K(X_{-\infty}^0)$ is finite (though not necessarily bounded) almost surely.

This class includes of course all finite order Markov chains but also many other processes such as the finitarily determined processes of Kalikow, Katznelson and Weiss [10], which serve to represent all isomorphism classes of zero entropy processes. For some concrete examples that are not Markovian consider the following example:

Example 1. Let $\{M_n\}$ be any stationary and ergodic first order Markov chain with finite or countably infinite state space S . Let $s \in S$ be an arbitrary state with $P(M_1 = s) > 0$. Now let $X_n = I_{\{M_n=s\}}$. By Shields [24] Chapter I.2.c.1, the binary time series $\{X_n\}$ is stationary and ergodic. It is also finitarily Markovian. Indeed, the conditional probability $P(X_1 = 1 | X_{-\infty}^0)$ does not depend on values beyond the first (going backwards) occurrence of one in $X_{-\infty}^0$ which identifies the first (going backwards) occurrence of state s in the Markov chain $\{M_n\}$. The resulting time series $\{X_n\}$ is not a Markov chain of any order in general. Indeed, consider the Markov chain $\{M_n\}$ with state space $S = \{0, 1, 2\}$ and transition probabilities $P(M_2 = 1 | M_1 = 0) = P(M_2 = 2 | M_1 = 1) = 1$, $P(M_2 = 0 | M_1 = 2) = P(M_2 = 1 | M_1 = 2) = 0.5$. This yields a stationary and ergodic Markov chain $\{M_n\}$, cf. Example I.2.8 in Shields [24]. Clearly, the resulting time series $X_n = I_{\{M_n=0\}}$ will not be Markov of any order. The conditional probability $P(X_1 = 0 | X_{-\infty}^0)$ depends on whether until the first (going backwards) occurrence of one you see even or odd number of zeros. These examples include all stationary and ergodic binary renewal processes with finite expected inter-arrival times, a basic class for many applications. (A stationary and ergodic binary renewal process is defined as a stationary and ergodic binary process such that the times between occurrences of ones are independent and identically distributed with finite expectation, cf. Chapter I.2.c.1 in Shields [24].)

We note that Morvai and Weiss [18] proved that there is no classification rule for discriminating the class of finitarily Markovian processes from other ergodic processes.

For the finitarily Markovian processes an important notion is that of a *memory word* which is defined as follows.

Definition 3. We say that w_{-k+1}^0 is a memory word if $p(w_{-k+1}^0) > 0$ and for all $i \geq 1$, all $y \in \mathcal{X}$, all $z_{-k-i+1}^{-k} \in \mathcal{X}^i$

$$p(y|w_{-k+1}^0) = p(y|z_{-k-i+1}^{-k}, w_{-k+1}^0)$$

provided $p(z_{-k-i+1}^{-k}, w_{-k+1}^0, y) > 0$.

Define the set \mathcal{W}_k of those memory words w_{-k+1}^0 with length k , that is,

$$\mathcal{W}_k = \{w_{-k+1}^0 \in \mathcal{X}^k: w_{-k+1}^0 \text{ is a memory word}\}.$$

Our first result is a solution of the backward estimation problem, namely determining the value of $K(X_{-\infty}^0)$ from observations of increasing length of the data segments X_{-n}^0 . We will give in the next section a universal consistent estimator which will converge almost surely to the memory length $K(X_{-\infty}^0)$ for any ergodic finitarily Markovian process on a countable state space. The proofs that we give are pretty explicit and give some information on the average length of a memory word and the extent to which the stationary distribution diffuses over the state space one could extract rates for the convergence of the estimators from our estimates. We concentrate however, on the more universal aspects of the problem.

As is usual in these kinds of questions, the problem of forward estimation, namely trying to determine $K(X_{-\infty}^n)$ from successive observations of X_0^n is more difficult. The stationarity means that results in probability can be carried over automatically. However, almost sure results present serious problems. For example, while Ornstein in [21] (cf. Morvai et al. [12] also) showed that there is a universal consistent estimator for the conditional probability of X_1 given $X_{-\infty}^0$ based on successive observations of the past, Bailey [1] showed that one simply cannot estimate the forward conditional probabilities in a similar universal way. One can obtain results modulo a zero density set of moments, but if one wants to be sure that when one is giving an estimate that eventually the estimate converges one is forced to resort to estimating along a sequence of stopping times (cf. Morvai [11], Morvai and Weiss [13,14,19]). For some more results in this circle of ideas of what can be learned about processes by forward observations see Ornstein and Weiss [22], Dembo and Peres [6], Nobel [20], and Csiszár [3].

Recently in Csiszár and Talata [5] the authors define a finite context to be a memory word w of minimal length, that is, no proper suffix of w is a memory word. An infinite context for a process is an infinite string with all finite suffix having positive probability but none of them being a memory word. They treat there the problem of estimating the entire context tree in case the size of the alphabet is finite. For a bounded depth context tree, the process is Markovian, while for an unbounded depth context tree the universal pointwise consistency result there is obtained only for the truncated trees which are again finite in size. This is in contrast to our results which deal with infinite alphabet size and consistency in estimating memory words of arbitrary length. This is what forces us to consider estimating at specially chosen times.

In the succeeding two Sections 3 and 4 we will present two such schemes which depend upon a positive parameter ϵ , and we guarantee that sequence of times along which the estimates are being given have density at least $1 - \epsilon$. The purpose of the next two sections is to show that this result is sharp in that the ϵ cannot be removed even in more restricted classes of processes. In Section 5 we show that you cannot achieve density one in forward estimation of the memory in the class of Markov chains on countable alphabets, while in Section 6 we prove a similar negative result for binary valued finitarily Markovian processes.

The last part of the paper is devoted to seeing how this memory length estimation can be applied to estimating conditional probabilities. In Section 7 we do this for finitarily Markovian processes along a sequence of stopping times which achieve density $1 - \epsilon$. We do not know if the ϵ can be dropped in this case for the estimation of conditional probabilities.

We can dispense with ϵ in the Markovian case. In Section 8 we use an earlier result of ours on a universal estimator for the order of a finite order Markov chain on a countable alphabet in order to estimate the conditional probabilities along a sequence of stopping times of density one.

2. Backward estimation of the memory length for finitarily Markovian processes

In order to estimate $K(X_{-\infty}^0)$ we need to define some explicit statistics. The first is a measurement of the failure of w_{-k+1}^0 to be a memory word.

For w_{-k+1}^0 of positive probability define

$$\Delta_k(w_{-k+1}^0) = \sup_{1 \leq i} \sup_{\{z_{-k-i+1}^{-k} \in \mathcal{X}^i, x \in \mathcal{X}: p(z_{-k-i+1}^{-k}, w_{-k+1}^0, x) > 0\}} |p(x|w_{-k+1}^0) - p(x|z_{-k-i+1}^{-k}, w_{-k+1}^0)|.$$

Clearly this will vanish precisely when w_{-k+1}^0 is a memory word. We need to define an empirical version of this based on the observation of a finite data segment X_{-n}^0 . To this end first define the empirical version of the conditional probability as

$$\hat{p}_n(x|w_{-k+1}^0) = \frac{\#\{-n+k-1 \leq t \leq -1: X_{t-k+1}^{t+1} = (w_{-k+1}^0, x)\}}{\#\{-n+k-1 \leq t \leq -1: X_{t-k+1}^t = w_{-k+1}^0\}}.$$

These empirical distributions, as well as the sets we are about to introduce are functions of X_{-n}^0 , but we suppress the dependence to keep the notation manageable.

For a fixed $0 < \gamma < 1$ let \mathcal{L}_k^n denote the set of strings with length $k+1$ which appear more than $n^{1-\gamma}$ times in X_{-n}^0 . That is,

$$\mathcal{L}_k^n = \{x_{-k}^0 \in \mathcal{X}^{k+1}: \#\{-n+k \leq t \leq 0: X_{t-k}^t = x_{-k}^0\} > n^{1-\gamma}\}.$$

Finally, define the empirical version of Δ_k as follows:

$$\hat{\Delta}_k^n(w_{-k+1}^0) = \max_{1 \leq i \leq n} \max_{\{z_{-k-i+1}^{-k}, w_{-k+1}^0, x\} \in \mathcal{L}_{k+i}^n} |\hat{p}_n(x|w_{-k+1}^0) - \hat{p}_n(x|z_{-k-i+1}^{-k}, w_{-k+1}^0)|.$$

Let us agree by convention that if the smallest of the sets over which we are maximizing is empty then $\hat{\Delta}_k^n = 0$. Observe, that by ergodicity, the ergodic theorem implies that almost surely the empirical distributions \hat{p} converge to the true distributions p and so for any $w_{-k+1}^0 \in \mathcal{X}^k$,

$$\liminf_{n \rightarrow \infty} \hat{\Delta}_k^n(w_{-k+1}^0) \geq \Delta_k(w_{-k+1}^0) \quad \text{almost surely.}$$

With this in hand we can give a test for w_{-k+1}^0 to be a memory word. Let $0 < \beta < \frac{1-\gamma}{2}$ be arbitrary. Let $NTEST_n(w_{-k+1}^0) = YES$ if $\hat{\Delta}_k^n(w_{-k+1}^0) \leq n^{-\beta}$ and NO otherwise. Note that $NTEST_n$ depends on X_{-n}^0 .

Theorem 1. *Eventually almost surely, $NTEST_n(w_{-k+1}^0) = YES$ if and only if w_{-k+1}^0 is a memory word.*

We define an estimate χ_n for $K(X_{-\infty}^0)$ from samples X_{-n}^0 as follows. Set $\chi_0 = 0$, and for $n \geq 1$ let χ_n be the smallest $0 \leq k < n$ such that $NTEST_n(X_{-k+1}^0) = YES$ if there is such and n otherwise.

Theorem 2. $\chi_n = K(X_{-\infty}^0)$ eventually almost surely.

In order to prove these theorems we need some lemmas. The first is a variant of the simple fact that the states u_i that follow the successive occurrences of a fixed memory word w are independent and identically distributed random variables. We cannot use such a naive version because we are dealing with a countable alphabet, and thus even the collection of memory words of a fixed length is infinite. In order to cut down to a manageable set we would like to consider only those words that appear in the sample X_{-n}^0 , but now the independence becomes a little subtler. This is the reason for the rather forbidding looking formulas in the proof of the next lemma. What we do is fix a location $(l-k, l]$ in the index set and then fix a memory word w_{-k+1}^0 that occurs there together with a particular state x that follows it. The random times $l + \lambda_{l,k,i}^+$ and $l - \lambda_{l,k,i}^-$ are the other occurrences of this memory word in the process. Here is the formal definition. Set $\lambda_{l,k,0}^+ = 0, \lambda_{l,k,0}^- = 0$ and define

$$\lambda_{l,k,i}^+ = \lambda_{l,k,i-1}^+ + \min\{t > 0: X_{l+\lambda_{l,k,i-1}^+-k+1+t}^{l+\lambda_{l,k,i-1}^+} = X_{l+\lambda_{l,k,i-1}^+-k+1}^{l+\lambda_{l,k,i-1}^+}\} \tag{1}$$

and

$$\lambda_{l,k,i}^- = \lambda_{l,k,i-1}^- + \min\{t > 0: X_{l-\lambda_{l,k,i-1}^- -k+1-t}^{l-\lambda_{l,k,i-1}^-} = X_{l-\lambda_{l,k,i-1}^- -k+1}^{l-\lambda_{l,k,i-1}^-}\} \tag{2}$$

Lemma 1. Assume w_{-k+1}^0 is a memory word and x is a letter. Then for any $i, j \geq 1$,

$$X_{l-\lambda_{l,k,i}^-+1}, \dots, X_{l-\lambda_{l,k,1}^-+1}, X_{l+\lambda_{l,k,1}^++1}, \dots, X_{l+\lambda_{l,k,j}^++1}$$

are conditionally independent and identically distributed random variables given $X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x$, where the identical distribution is $p(\cdot | w_{-k+1}^0)$.

Proof. Fix the values $z_{-i}, \dots, z_{-1}, u_1, \dots, u_j$ and x in the alphabet and calculate

$$\begin{aligned} & P(X_{l-\lambda_{l,k,i}^-+1} = z_{-i}, \dots, X_{l-\lambda_{l,k,1}^-+1} = z_{-1}, \\ & \quad X_{l+\lambda_{l,k,1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,j}^++1} = u_j | X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x) \\ &= P(X_{l-\lambda_{l,k,i}^-+1} = z_{-i}, \dots, X_{l-\lambda_{l,k,1}^-+1} = z_{-1}, X_{l+\lambda_{l,k,1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,j}^++1} = u_j, \\ & \quad X_{l-k+1}^l = w_{-k+1}^0 | X_{l+1} = x) / P(X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x). \end{aligned}$$

In order to be able to use the fact that w_{-k+1}^0 is a memory word we will shift back to the first occurrence at $l - \lambda_{l,k,i}^-$ and use the stationarity.

$$\begin{aligned} & P(X_{l-\lambda_{l,k,i}^-+1} = z_{-i}, \dots, X_{l-\lambda_{l,k,1}^-+1} = z_{-1}, X_{l+\lambda_{l,k,1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,j}^++1} = u_j, \\ & \quad X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x, \lambda_{l,k,i}^- = t) \\ &= P(X_{l-t+\lambda_{l-t,k,i-i}^++1} = z_{-i}, \dots, X_{l-t+\lambda_{l-t,k,i-1}^++1} = z_{-1}, X_{l-t+\lambda_{l-t,k,i+1}^++1} = u_1, \dots, X_{l-t+\lambda_{l-t,k,i+j}^++1} = u_j, \\ & \quad X_{l-t+\lambda_{l-t,k,i}^++1-k+1} = w_{-k+1}^0, X_{l-t+\lambda_{l-t,k,i}^++1} = x, \lambda_{l-t,k,i}^+ = t) \\ &= P(T^{-t} \{ X_{l-t+\lambda_{l-t,k,i-i}^++1} = z_{-i}, \dots, X_{l-t+\lambda_{l-t,k,i-1}^++1} = z_{-1}, \\ & \quad X_{l-t+\lambda_{l-t,k,i+1}^++1} = u_1, \dots, X_{l-t+\lambda_{l-t,k,i+j}^++1} = u_j, X_{l-t+\lambda_{l-t,k,i}^++1-k+1} = w_{-k+1}^0, \\ & \quad X_{l-t+\lambda_{l-t,k,i}^++1} = x, \lambda_{l-t,k,i}^+ = t \}) \\ &= P(X_{l+\lambda_{l,k,0}^++1} = z_{-i}, \dots, X_{l+\lambda_{l,k,i-1}^++1} = z_{-1}, X_{l+\lambda_{l,k,i+1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,i+j}^++1} = u_j, \\ & \quad X_{l+\lambda_{l,k,i}^++1-k+1} = w_{-k+1}^0, X_{l+\lambda_{l,k,i}^++1} = x, \lambda_{l,k,i}^+ = t). \end{aligned}$$

Summing over t ,

$$\begin{aligned} & P(X_{l-\lambda_{l,k,i}^-+1} = z_{-i}, \dots, X_{l-\lambda_{l,k,1}^-+1} = z_{-1}, X_{l+\lambda_{l,k,1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,j}^++1} = u_j, X_{l-k+1}^l = w_{-k+1}^0 | X_{l+1} = x) \\ &= P(X_{l+\lambda_{l,k,0}^++1} = z_{-i}, \dots, X_{l+\lambda_{l,k,i-1}^++1} = z_{-1}, X_{l+\lambda_{l,k,i+1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,i+j}^++1} = u_j, \\ & \quad X_{l+\lambda_{l,k,i}^++1-k+1} = w_{-k+1}^0, X_{l+\lambda_{l,k,i}^++1} = x). \end{aligned} \tag{3}$$

Now telescoping the right-hand side we get

$$\begin{aligned} & P(X_{l-\lambda_{l,k,i}^-+1} = z_{-i}, \dots, X_{l-\lambda_{l,k,1}^-+1} = z_{-1}, X_{l+\lambda_{l,k,1}^++1} = u_1, \dots, X_{l+\lambda_{l,k,j}^++1} = u_j | X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x) \\ &= P(X_{-k+1}^0 = w_{-k+1}^0) \prod_{h=1}^i P(X_1 = z_{-h} | X_{-k+1}^0 = w_{-k+1}^0) P(X_1 = x | X_{-k+1}^0 = w_{-k+1}^0) \\ & \quad \times \frac{\prod_{h=1}^j P(X_1 = u_h | X_{-k+1}^0 = w_{-k+1}^0)}{P(X_{-k+1}^0 = w_{-k+1}^0) P(X_1 = x | X_{-k+1}^0 = w_{-k+1}^0)} \\ &= \prod_{h=1}^i P(X_1 = z_{-h} | X_{-k+1}^0 = w_{-k+1}^0) \prod_{h=1}^j P(X_1 = u_h | X_{-k+1}^0 = w_{-k+1}^0). \end{aligned}$$

We have to prove that

$$P(X_1 = z_{-h} | X_{-k+1}^0 = w_{-k+1}^0) = P(X_{l-\lambda_{l,k,h}^-} = z_{-h} | X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x).$$

Indeed, by (3) and stationarity,

$$\begin{aligned} & P(X_{l-\lambda_{l,k,h}^-} = z_{-h} | X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x) \\ &= \frac{P(X_{l-\lambda_{l,k,h}^-} = z_{-h}, X_{l-\lambda_{l,k,h}^-}^{l-\lambda_{l,k,h}^-} = w_{-k+1}^0, X_{l+1} = x)}{P(X_{l-k+1}^l = w_{-k+1}^0, X_{l+1} = x)} \\ &= P(X_{-k+1}^0 = w_{-k+1}^0) P(X_1 = z_{-h} | X_{-k+1}^0 = w_{-k+1}^0) \frac{P(X_{l+1} = x | X_{l-k+1}^l = w_{-k+1}^0, X_{l-\lambda_{l,k,h}^-} = z_{-h})}{P(X_{-k+1}^0 = w_{-k+1}^0) P(X_{l+1} = x | X_{l-k+1}^l = w_{-k+1}^0)} \\ &= P(X_1 = z_{-h} | X_{-k+1}^0 = w_{-k+1}^0). \end{aligned}$$

The proof of Lemma 1 is complete. \square

Lemma 2.

$$\begin{aligned} & P(\text{For some } 0 \leq k < n, -n+k-1 \leq l \leq -1: X_{l-k+1}^{l+1} \in \mathcal{L}_{k+1}^n, K(X_{-\infty}^l) \leq k, \\ & \quad |\hat{p}_n(X_{l+1} | X_{l-k+1}^l) - p(X_{l+1} | X_{l-k+1}^l)| > n^{-\beta}) \\ & \leq n^2 \sum_{h=\lfloor n^{1-\gamma} \rfloor}^{\infty} h 2 e^{-2n^{-2\beta}h}. \end{aligned}$$

Proof. For a given $0 \leq k < n$, $-n+k-1 \leq l \leq -1$ assume that $X_{l-k+1}^{l+1} = w_{-k+1}^0 x$ and w_{-k+1}^0 is a memory word. Since w_{-k+1}^0 is a memory word, by Lemma 1 and by Hoeffding's inequality (cf. Hoeffding [9] or Theorem 8.1 of Devroye et al. [7]) for sums of bounded independent random variables implies

$$\begin{aligned} & P\left(\left| \frac{\sum_{h=1}^i \mathbf{1}_{\{X_{l-\lambda_{l,k,h}^-} = x\}} + \sum_{h=1}^j \mathbf{1}_{\{X_{l+\lambda_{l,k,h}^+} = x\}}}{i+j} - p(x | w_{l-k+1}^0) \right| \right. \\ & \quad \left. \geq n^{-\beta} | X_{l-k+1}^{l+1} = w_{-k+1}^0 x \right) \leq 2 e^{-2n^{-2\beta}(i+j)}. \end{aligned}$$

Multiplying both sides by $P(X_{l-k+1}^{l+1} = w_{-k+1}^0 x)$ and summing over all possible memory words w_{-k+1}^0 and x we get that

$$\begin{aligned} & P\left(K(X_{-\infty}^l) \leq k, X_{l-k+1}^{l+1} \in \mathcal{L}_{k+1}^n, \right. \\ & \quad \left| \frac{\sum_{h=1}^i \mathbf{1}_{\{X_{l-\lambda_{l,k,h}^-} = X_{l+1}\}} + \sum_{h=1}^j \mathbf{1}_{\{X_{l+\lambda_{l,k,h}^+} = X_{l+1}\}}}{i+j} - p(X_{l+1} | X_{l-k+1}^l) \right| > n^{-\beta} \Big) \\ & \leq 2 e^{-2n^{-2\beta}(i+j)}. \end{aligned}$$

Summing over all pairs (k, l) such that $0 \leq k < n$ and all $-n+k-1 \leq l \leq -1$ and over all pairs (i, j) such that $i \geq 0$, $j \geq 0$, $i+j \geq \lfloor n^{1-\gamma} \rfloor$ we complete the proof of Lemma 2. \square

Lemma 3.

$$P\left(\max_{w_{-k+1}^0 \in \mathcal{W}_k} \hat{\Delta}_k^n(w_{-k+1}^0) > n^{-\beta}\right) \leq n^3 \sum_{h=\lfloor n^{1-\gamma} \rfloor}^{\infty} h 4 e^{-\frac{n^{-2\beta}h}{2}}.$$

Proof.

$$\begin{aligned}
 & P\left(\max_{w_{-k+1}^0 \in \mathcal{W}_k} \hat{\Delta}_k^n(w_{-k+1}^0) > n^{-\beta}\right) \\
 & \leq \sum_{i=1}^n P\left(\max_{w_{-k+1}^0 \in \mathcal{W}_k} \max_{(z_{-k-i+1}^{-k}, w_{-k+1}^0, x) \in \mathcal{L}_{k+i}^n} |\hat{p}_n(x|w_{-k+1}^0) - \hat{p}_n(x|z_{-k-i+1}^{-k}, w_{-k+1}^0)| > n^{-\beta}\right) \\
 & \leq \sum_{i=1}^n P\left(\max_{w_{-k+1}^0 \in \mathcal{W}_k} \max_{(z_{-k-i+1}^{-k}, w_{-k+1}^0, x) \in \mathcal{L}_{k+i}^n} |\hat{p}_n(x|w_{-k+1}^0) - p(x|w_{-k+1}^0)| > n^{-\beta}/2\right) \\
 & \quad + \sum_{i=1}^n P\left(\max_{w_{-k+1}^0 \in \mathcal{W}_k} \max_{(z_{-k-i+1}^{-k}, w_{-k+1}^0, x) \in \mathcal{L}_{k+i}^n} |p(x|z_{-k-i+1}^{-k}, w_{-k+1}^0) - \hat{p}_n(x|z_{-k-i+1}^{-k}, w_{-k+1}^0)| > n^{-\beta}/2\right).
 \end{aligned}$$

By Lemma 2, both terms inside the sum can be upperbounded by an exponential, and summing over i we get the statement and so the proof of Lemma 3 is complete. \square

Proof of Theorem 1. If w_{-k+1}^0 is not a memory word, then there are z_{-k-i+1}^{-k} and x such that $p(x|w_{-k+1}^0) \neq p(x|z_{-k-i+1}^{-k} w_{-k+1}^0)$ and $p(z_{-k-i+1}^{-k} w_{-k+1}^0, x) > 0$. By ergodicity, $NTEST_n(w_{-k+1}^0) = NO$ eventually almost surely.

Assume w_{-k+1}^0 is a memory word. We will estimate the probability of the undesirable event as follows: by Lemma 3,

$$P(\hat{\Delta}_k^n(w_{-k+1}^0) > n^{-\beta}) \leq n^3 \sum_{h=\lfloor n^{1-\gamma} \rfloor}^{\infty} h 4 e^{-\frac{n^{-2\beta} h}{2}}.$$

The right-hand side is summable provided $2\beta + \gamma < 1$ and the Borel–Cantelli lemma yields that

$$P(\hat{\Delta}_k^n(w_{-k+1}^0) \leq n^{-\beta} \text{ eventually}) = 1$$

and so $NTEST_n(w_{-k+1}^0) = YES$ eventually almost surely. The proof of Theorem 1 is complete. \square

Proof of Theorem 2. Since $X_{-K(X_{-\infty}^0)+1}^0$ is a memory word and none of its suffixes has this property, $\chi_n = K(X_{-\infty}^0)$ eventually almost surely, by Theorem 1. The proof of Theorem 2 is complete. \square

3. Forward estimation of the memory length for finitarily Markovian processes

Define $PTEST_n(w_{-k+1}^0)(X_0^n) = NTEST_n(w_{-k+1}^0)(T^n X_0^n)$ where T is the left shift operator.

Theorem 3. *Eventually almost surely, $PTEST_n(w_{-k+1}^0) = YES$ if and only if w_{-k+1}^0 is a memory word.*

Define a list of words $\{w(0), w(1), w(2), \dots, w(n), \dots\}$ such that all words of all lengths are listed and a word cannot precede its suffix. Note that $w(0)$ is the empty word.

Now define sets of indices A_n^i as follows. Let $A_n^0 = \{0, 1, \dots, n\}$ and for $i > 0$ define

$$A_n^i = \{|w(i)| - 1 \leq j \leq n: X_{j-|w(i)|+1}^j = w(i)\}. \tag{4}$$

Let $\epsilon > 0$ be fixed. Define $\theta_n(\epsilon) < n$ to be the minimal j such that

$$\frac{|\bigcup_{i \leq j: PTEST_n(w(i))=YES} A_n^i|}{n+1} \geq 1 - \frac{\epsilon}{2} \tag{5}$$

and n if no such j exists. We estimate for the length of the memory of $X_{-\infty}^n$ looking backwards if $n \in \bigcup_{i \leq \theta_n(\epsilon), PTEST_n(w(i))=YES} A_n^i$. The set of n 's for which this holds will be the set for which we estimate the memory and we denote this set by \mathcal{N} . Note that the event $n \in \mathcal{N}$ depends only on X_0^n , and thus \mathcal{N} can be thought of as a sequence of stopping times.

We define for $n \in \mathcal{N}$,

$$\kappa_n = \min\{i \geq 0: X_{n-|w(i)|+1}^n = w(i), PTEST_n(w(i)) = YES\}.$$

For $n \in \mathcal{N}$ define

$$\rho_n(X_0^n) = |w(\kappa_n)|.$$

Note that $\rho_n, \theta_n, \kappa_n$ and \mathcal{N} depend on ϵ , however, we will not denote this dependence on ϵ explicitly.

Theorem 4. *Let $\epsilon > 0$ be fixed. Then for $n \in \mathcal{N}$,*

$$\rho_n = K(X_{-\infty}^n) \text{ eventually almost surely,} \tag{6}$$

and

$$\liminf_{n \rightarrow \infty} \frac{|\mathcal{N} \cap \{0, 1, \dots, n-1\}|}{n} \geq 1 - \epsilon. \tag{7}$$

For $n \in \mathcal{N}$, $X_{n-\rho_n+1}^n$ appears at least $n^{1-\gamma}$ times eventually almost surely.

Proof of Theorem 3. Since the proof of Theorem 1 was based on a Borel–Cantelli lemma, the time shift in defining $PTEST_n$ makes no difference and we literally copy the proof of Theorem 1. The proof of Theorem 3 is complete. \square

Proof of Theorem 4. There is a N large enough such that

$$P(K(X_{-\infty}^0) < N) \geq 1 - \epsilon/4.$$

The sequence θ_n is bounded along individual sequences of the process with probability one. (This may be seen by first choosing a sufficiently large finite set $\{w(0), \dots, w(M)\}$ of memory words so that the probability of seeing at least one of them in position zero is greater than $1 - \epsilon/4$ and then applying Theorem 3 and the ergodic theorem we see that almost surely for all sufficiently large n , $\theta_n \leq M$. This implies of course that θ_n is bounded pointwise as claimed.) Thus by Theorem 3,

$$\rho_n = K(X_{-\infty}^n) \text{ provided } X_{n-K(X_{-\infty}^n)+1}^n \in \mathcal{W}_{K(X_{-\infty}^n)} \cap \{w(0), \dots, w(\theta_n)\}$$

eventually almost surely. We have proved the consistency. Let J denote the smallest j such that

$$\sum_{k=0}^j P(X_{-k+1}^0 \in \mathcal{W}_k \cap \{w(0), \dots, w(j)\}) \geq 1 - \frac{\epsilon}{2}.$$

It is obvious from the definition above that

$$\sum_{k=0}^{J-1} P(X_{-k+1}^0 \in \mathcal{W}_k \cap \{w(0), \dots, w(J-1)\}) < 1 - \frac{\epsilon}{2}.$$

Thus $\theta_n \geq J$ eventually almost surely. Thus

$$\liminf_{n \rightarrow \infty} \frac{|\mathcal{N} \cap \{0, 1, \dots, n\}|}{n+1} \geq \liminf_{n \rightarrow \infty} \frac{|\bigcup_{i=0}^J A_n^i|}{n+1} \geq 1 - \frac{\epsilon}{2} \text{ almost surely.}$$

We have proved that \mathcal{N} has density at least $1 - \epsilon/2$. Since θ_n is bounded, for $n \in \mathcal{N}$ eventually, $w(\kappa_n)$ appears at least $n^{1-\gamma}$ times. The proof of Theorem 4 is complete. \square

4. Another approach to estimating the memory length for finitarily Markovian processes

In the preceding section we made use of the fact that the proof that we gave for the backward memory estimator was via a rough probability estimate and the Borel–Cantelli lemma. This enabled us to copy it directly for the forward estimation. In this section we shall show that any successful backward memory estimator can be used to get the same kind of result. We will denote by χ_n some fixed consistent backward estimator for the memory length such as the χ_n

of Section 2. To this end, based on the successive forward samples we construct many infinite sample points of the $X_{-\infty}^0$ process.

To construct a sample of the $X_{-\infty}^0$ process from the forward data segment X_0^n , we use the procedure that we used in Morvai and Weiss [15]. Begin with X_0 , then look for its first recurrence, i.e. the minimum $t_0 > 0$ such that $X_{t_0} = X_0$ and then extend X_0 to the left by adding X_{t_0-1} . Next look for the first recurrence of $X_{t_0-1}X_{t_0}$, in a position $t_1 > t_0$, i.e. $X_{t_1-1}X_{t_1} = X_{t_0-1}X_{t_0}$ and then again extend to the left by adding X_{t_1-2} obtaining $X_{t_1-2}X_{t_1-1}X_{t_1}$ as the first three symbols of our sample for the backward process. We will denote this by $\tilde{X}_{-2}^0 = X_{t_1-2}^{t_1}$. Continuing in this way, we can develop from X_0^n a point $\tilde{X}_{-\infty}^0$ which we shall show has the same distribution as $X_{-\infty}^0$. We need to do this starting at each $i \geq 0$. Here are the formulas that accomplish this end.

For $i = 0, 1, \dots$ define auxiliary stopping times. Set $\zeta_{-1}(i) = -i$ and $\zeta_0(i) = 0$. For $n = 1, 2, \dots$, let

$$\zeta_n(i) = \zeta_{n-1}(i) + \min\{t > 0: X_{i+\zeta_{n-1}(i)-(n-1)+t}^{i+\zeta_{n-1}(i)+t} = X_{i+\zeta_{n-1}(i)-(n-1)}^{i+\zeta_{n-1}(i)}\}. \quad (8)$$

Among other things, using $\zeta_n(i)$ we can define very useful processes $\{\tilde{X}_n(i)\}_{n=-\infty}^0$ as a function of X_0^∞ as follows.

Define

$$\tilde{X}_{-n}(i) = X_{i+\zeta_n(i)-n}. \quad (9)$$

It is clear that in this way we defined processes $\{\tilde{X}_n(i)\}_{n=-\infty}^0$. We will see that the $\{\tilde{X}_n(i)\}_{n=-\infty}^0$ has the same distribution as the original process, and for now assume that this is so.

Let

$$\eta_n(i) = \max\{j \geq -1: i + \zeta_j(i) \leq n\}. \quad (10)$$

Note that $(\tilde{X}_{-\eta_n(i)}(i), \dots, \tilde{X}_0(i))$ is measurable with respect to X_0^n .

Define $\rho_n^i = \chi_{\eta_n(i)}(\tilde{X}_{-\infty}^0(i))$ if $\eta_n(i) \geq 0$ and $\rho_n^i = 0$ otherwise. Note that ρ_n^i is also measurable with respect to X_0^n .

Define sets of indices A_n^i as follows.

$$A_n^i = \{\rho_n^i \leq j \leq n: X_{j-\rho_n^i+1}^j = \tilde{X}_{-\rho_n^i+1}^0(i)\}. \quad (11)$$

For any fixed i , eventually, $\tilde{X}_{-\rho_n^i+1}^0(i)$ is a memory word, so the sets A_n^i are simply the places where this fixed word occurs. Let $\epsilon > 0$ be fixed. Define $\theta_n(\epsilon)$ to be the minimal j such that

$$\frac{|\bigcup_{i \leq j} A_n^i|}{n} \geq 1 - \frac{\epsilon}{2}. \quad (12)$$

We estimate for the order of $X_{-\infty}^n$ looking backwards if $n \in \bigcup_{i \leq \theta_n(\epsilon)} A_n^i$. The set of n 's for which this holds will be the set for which we estimate the memory and we denote this set by \mathcal{N} . Note that the event $n \in \mathcal{N}$ depends only on X_0^n , and thus \mathcal{N} can be thought of as a sequence of stopping times.

In case $n \in \bigcup_{i \leq \theta_n(\epsilon)} A_n^i$ we define

$$\kappa_n = \min\{i \geq 0: \tilde{X}_{-\rho_n^i+1}^0(i) = X_{n-\rho_n^i+1}^n\}.$$

Note that θ_n, κ_n and \mathcal{N} depend on ϵ , however, we will not denote this dependence on ϵ explicitly.

Theorem 5. *Let $\epsilon > 0$ be fixed. Then for $n \in \mathcal{N}$,*

$$\rho_n^{\kappa_n} = K(X_{-\infty}^n) \quad \text{eventually almost surely,} \quad (13)$$

and

$$\liminf_{n \rightarrow \infty} \frac{|\mathcal{N} \cap \{0, 1, \dots, n-1\}|}{n} \geq 1 - \epsilon. \quad (14)$$

For $n \in \mathcal{N}$, $X_{n-\rho_n^{\kappa_n}}^n$ appears at least $n^{1-\gamma}$ times eventually almost surely.

Lemma 4. *For all i the time series $\{\tilde{X}_n(i)\}_{n=-\infty}^0$ and $\{X_n\}_{n=-\infty}^0$ have identical distribution.*

Proof. For all $k \geq 1$ and $1 \leq i \leq k$ define $\hat{\zeta}_0^k = 0$ and

$$\hat{\zeta}_i^k = \hat{\zeta}_{i-1}^k - \min\{t > 0: X_{\hat{\zeta}_{i-1}^k - (k-i) - t}^{\hat{\zeta}_{i-1}^k - t} = X_{\hat{\zeta}_{i-1}^k - (k-i)}^{\hat{\zeta}_{i-1}^k}\}.$$

Let T denote the left shift operator, that is, $(Tx_{-\infty}^\infty)_i = x_{i+1}$. It is easy to see that if and only if $\zeta_k(i)(x_{-\infty}^\infty) = l$ then $\hat{\zeta}_k^k(T^{(i+l)}x_{-\infty}^\infty) = -l$.

Now the statement follows from stationarity and the fact that for $k \geq 0$, $x_{-k}^0 \in \mathcal{X}^{k+1}$, $l \geq 0$,

$$T^{i+l}\{X_{i+\hat{\zeta}_k(i)-k}^{i+\hat{\zeta}_k(i)} = x_{-k}^0, \zeta_k(i) = l\} = \{X_{-k}^0 = x_{-k}^0, \hat{\zeta}_k^k(X_{-\infty}^0) = -l\}. \tag{15}$$

The proof of Lemma 4 is complete. \square

Lemma 5. If $P(X_0^n = w_0^n) > 0$ for the string w_0^n then almost surely,

$$\tilde{X}_{-n}^0(i) = w_0^n \text{ for some } i. \tag{16}$$

Proof. Let t denote the $(n + 1)$ th occurrence of the string w_0^n in X_0^∞ . It is easy to see that there must be a $0 \leq i \leq t$ such that

$$X_{i+\zeta_k(i)-k}^{i+\zeta_k(i)} = X_{t-k}^t \text{ for } k = 0, 1, \dots, n$$

and so

$$\tilde{X}_{-n}^0(i) = w_0^n.$$

The proof of Lemma 5 is complete. \square

Proof of Theorem 5. There is a N large enough such that

$$P(K(X_{-\infty}^0) < N) \geq 1 - \epsilon/4.$$

Then by Lemma 5 and ergodicity, θ_n is a bounded sequence (cf. the proof of Theorem 4). By Lemma 4 and Theorem 2

$$\rho_n^i = K(\tilde{X}_{-\infty}^0(i)) \text{ for all } i = 1, \dots, \theta_n$$

eventually almost surely. We have proved (13). We have to prove (14). Let J denote the smallest j such that

$$\sum_{i=0}^j P(\tilde{X}_{-K(\tilde{X}_{-\infty}^0(i))}^0(i)) \geq 1 - \frac{\epsilon}{2}.$$

It is obvious from the definition above that

$$\sum_{i=0}^{J-1} P(\tilde{X}_{-K(\tilde{X}_{-\infty}^0(i))}^0(i)) < 1 - \frac{\epsilon}{2}.$$

Thus $\theta_n \geq J$ eventually almost surely. Thus

$$\liminf_{n \rightarrow \infty} \frac{|\mathcal{N} \cap \{0, 1, \dots, n\}|}{n+1} \geq \liminf_{n \rightarrow \infty} \frac{|\bigcup_{i=0}^J A_n^i|}{n+1} \geq 1 - \frac{\epsilon}{2} \text{ almost surely.}$$

The proof of Theorem 5 is complete. \square

5. Memory estimation for Markov processes

In this section we shall examine how well can one estimate the local memory length for finite order Markov chains. In the case of finite alphabets this can be done with stopping times that eventually cover all time epochs. (Indeed, assume $\{X_n\}$ is a Markov chain taking values from a finite set. Assume $ORDEST_n$ estimates the order in a pointwise sense from data X_0^n e.g. as in Csiszár and Shields [4] or in Morvai and Weiss [16]. Then let

$$\rho_n = \min\{0 \leq t \leq ORDEST_n: PTEST_n(X_{n-t+1}^n) = YES\}$$

if there is such t and 0 otherwise. Since $ORDEST_n$ eventually gives the right order and there are finitely many possible strings with length not greater than the order thus $\rho_n = K(X_{-\infty}^n)$ eventually almost surely by Theorem 3.)

However, as soon as one goes to a countable alphabet, even if the order is known to be two and we are just trying to decide whether the X_n alone is a memory word or not, there is no sequence of stopping times which is guaranteed to succeed eventually and whose density is one. This shows that the ϵ in the preceding sections cannot be eliminated.

Theorem 6. *There are no strictly increasing sequence of stopping times $\{\lambda_n\}$ and estimators $\{h_n(X_0, \dots, X_{\lambda_n})\}$ taking the values one and two, such that for all countable alphabet Markov chains of order two:*

$$\lim_{n \rightarrow \infty} \frac{\lambda_n}{n} = 1$$

and

$$\lim_{n \rightarrow \infty} |h_n(X_0, \dots, X_{\lambda_n}) - K(X_0^{\lambda_n})| = 0 \quad \text{with probability one.}$$

To prove the theorem we will assume that such a pair of stopping times and estimators exist and construct a Markov chain $\{X_n\}$ of order two for which they fail. The Markov chain of order two that we construct will have for its state space the nonnegative integers \mathbf{N} , and it will be a perturbation of the 1-step Markov chain Z_n defined by the following formulae:

$$P_{s,s+r} = 2^{-r-1} \quad \text{for all } r \geq 1,$$

$$P_{s,s} = 2^{-s-1}$$

and

$$P_{s,j} = 2^{-j-2} \quad \text{for all } 0 \leq j < s.$$

Notice that from any state $s > 0$ there is a fixed probability of $\frac{1}{4}$ of going to 0. Also there is a strictly positive probability of going from any state to any other. These properties ensure that there is a finite stationary measure. The ultimate chain will preserve most of these conditional probabilities, with the difference depending on a sequence of integers $t_k \gg k$ which will be defined later. The perturbed chain $\{X_n\}$ will have the same transition probabilities as the original chain $\{Z_n\}$ for X_n given (X_{n-2}, X_{n-1}) when the latter, (X_{n-2}, X_{n-1}) , equals any pair (t, s) with the exception of (t_k, k) for $k \geq 0$. In that case we will modify the probability of the transitions to k and $k+1$ by interchanging the values of $P_{k,k}$ and $P_{k,k+1}$. As soon as the first change is made 0 ceases to be a memory word and therefore the order of the new chain is two. Eventually all singletons cease to be memory words.

The t_k 's will be chosen inductively in a fashion depending on the purported sequence of stopping times and estimators that we are trying to show cannot exist. At the k th stage we will have only made these changes up to k . Let us denote the Markov process of order 2 that this defines by $\{Y_n^{(k)}\}$. More explicitly the process $\{Y_n^{(k)}\}$ is defined as follows. It has transition probabilities given by: for all $j \leq k$

$$P(Y_{n+2}^{(k)} = j | Y_n^{(k)} = t_j, Y_{n+1}^{(k)} = j) = P(Z_{n+2} = j+1 | Z_{n+1} = j),$$

$$P(Y_{n+2}^{(k)} = j+1 | Y_n^{(k)} = t_j, Y_{n+1}^{(k)} = j) = P(Z_{n+2} = j | Z_{n+1} = j),$$

and for all other values of (u, t, s) we have

$$P(Y_{n+2}^{(k)} = u | Y_n^{(k)} = t, Y_{n+1}^{(k)} = s) = P(Z_{n+2} = u | Z_{n+1} = s).$$

Thus for this process, all singletons j for $j > k$ are still memory words of length one, but none of the j with $0 \leq j \leq k$ are. The main technical lemma that we will need is that the distribution of finite blocks up to some preassigned length N of the $\{Y_n^{(k)}\}$ and $\{Y_n^{(k+1)}\}$ processes are arbitrarily close if the t_{k+1} is chosen sufficiently large. This is independent of the other properties of t_k that are needed, and so we begin by establishing this fact.

Our proof will be via a coupling argument. We will calculate the finite distributions of these two processes by calculating time averages of a pair of typical sequences generated by transition matrices starting from the pair 00. The coupling is especially easy since for both processes, from any state there is a fixed probability of moving to 00 of at least $\frac{1}{8}$ and therefore no matter how the two sequences diverge if we continue their evolution independently there is at every moment a fixed probability, namely $\frac{1}{64}$, of the processes returning simultaneously to 00.

Lemma 6. *With the definitions above for $\{Y_n^{(k)}\}$ and $\{Y_n^{(k+1)}\}$ if N and $\delta > 0$ are arbitrary, for any choice of t_{k+1} that is sufficiently large we will have that the variational distance between the distributions of $(Y_0^{(k)}, \dots, Y_N^{(k)})$ and $(Y_0^{(k+1)}, \dots, Y_N^{(k+1)})$ is at most δ .*

Proof. Since the $\{Y_n^{(k)}\}$ is fixed at the start, given N and δ we can choose T sufficiently large so that the stationary probability $\pi^{(k)}(t) < \gamma$ for any $t > T$, where $\gamma = \delta / (64 + 2N)$.

Suppose that we choose $t_{k+1} > T$. We begin the coupling by starting each of the processes at the pair 00. Denote by u_j and v_j the random sequences constructed by applying the transition functions for the two processes $\{Y_n^{(k)}\}, \{Y_n^{(k+1)}\}$ respectively. Until $u_j = t_{k+1}$ for the first time the sequences can be taken to be identical, since they have the same transition probabilities for pairs that do not include this state. We denote by σ_1 this moment, and continue the coupling now independently waiting for the first moment $j > \sigma_1$ that the equality $(u_j, u_{j+1}) = (v_j, v_{j+1}) = (0, 0)$ holds. Call this moment τ_1 . Notice that σ_1 is a function of the u'_j 's while τ_1 is a function of both processes. Beginning with τ_1 we can once again continue the evolution in an identical fashion until the first moment $j > \tau_1$ that $u_j = t_{k+1}$. Call that stopping time σ_2 . Note that this stopping time also depends on both processes. As before, as soon as this happens continue the processes independently until the first moment $j > \sigma_2$ that the equality $(u_j, u_{j+1}) = (v_j, v_{j+1}) = (0, 0)$ holds. It should now be clear how this is continued to build (with probability one) typical sequences for the two processes. In order to compare the stationary distributions of words up to length N in the two processes we need to know what is the relative frequency of the periods when we are coupling independently compared to the periods when we are producing the same symbols.

The asymptotic frequency of the occurrence of t_{k+1} in the u_j sequence is known to be at most γ and at the stopping times σ_i $u_{\sigma_i} = t_{k+1}$. The gaps $\tau_i - \sigma_i$ are independent for different i 's and have a length which has a geometric distribution with fixed parameter $\frac{1}{64}$ as we remarked earlier. Thus the average fraction of the time that the N -strings in the u and v sequences do not match exactly is at most $(64 + 2N)\gamma$. It follows that the variational distance between $(Y_0^{(k)}, \dots, Y_N^{(k)})$ and $(Y_0^{(k+1)}, \dots, Y_N^{(k+1)})$ is at most $(64 + 2N)\gamma = \delta$ and thus the lemma has been established. \square

We can now give the

Proof of Theorem 6. Suppose that there does exist a sequence of stopping times and estimators as in the statement of the theorem. We begin with the one step Markov chain Z_n described above and observe that the state 0 has a positive stationary probability. Since the λ_n 's have density one we can find an N_0 so that with probability at least $1 - \frac{1}{10}$ in the string $Z_0^{N_0}$ there will be some $\lambda_n < N_0$ with $Z_{\lambda_n} = 0$ and

$$h_n(Z_0, \dots, Z_{\lambda_n}) = 1.$$

We can apply the lemma with $N = N_0$ and $\delta = \frac{1}{10}$ to find a suitable t_0 with which we can define a $\{Y_n^{(0)}\}$ process in which now 0 is not a memory word, so that for those strings where $Z_{\lambda_n} = 0$ and $h_n(Z_0, \dots, Z_{\lambda_n}) = 1$ a definite mistake is being made. Such strings with length N_0 still have probability at least $1 - \frac{2}{10}$. Having defined $\{Y_n^{(0)}\}$ we notice now that the state 1 is still a memory word of length one with positive stationary probability, and therefore we can find an N_1 sufficiently large so that with probability at least $1 - \frac{1}{10^2}$ in the string $(Y_0^{(0)}, \dots, Y_{N_1}^{(0)})$ there will be some $\lambda_n < N_1$ with $Y_{\lambda_n}^{(0)} = 1$ and

$$h_n(Y_0^{(0)}, \dots, Y_{\lambda_n}^{(0)}) = 1.$$

As before we apply the lemma with $N = N_1$ and $\delta = 10^{-2}$ to find a suitable t_1 with which we can define the next process $\{Y_n^{(1)}\}$. For this process although 1 fails to be a memory word we still can estimate the probability that in a $(Y_0^{(1)}, \dots, Y_{N_1}^{(1)})$ string there will be some $\lambda_n < N_1$ with $Y_{\lambda_n}^{(1)} = 1$ and

$$h_n(Y_0^{(1)}, \dots, Y_{\lambda_n}^{(1)}) = 1$$

as being at least $1 - 2 \times 10^{-2}$. In addition, the previous estimate on strings of length N_0 is degraded only by 10^{-2} , since the estimate on the variational distance descends to strings of shorter length. By now it should be clear how to continue the inductive construction of the t_k 's. The ultimate process that we obtain, which we may denote simply by

$\{X_n\}$ is of course a Markov chain of order two, and it has no memory words of length one at all. However, for every k , the probability that there will be some $\lambda_n < N_k$ with $X_{\lambda_n} = k$ and

$$h_n(X_0, \dots, X_{\lambda_n}) = 1$$

will be at least $1 - \frac{2}{9 \times 10^k}$. The Borel–Cantelli lemma implies that with probability one there will be infinitely many mistakes being made by our estimator contrary to the assumption. This concludes the proof of Theorem 6. \square

6. Limitations for binary finitarily Markovian processes

In the preceding section we showed that we cannot achieve density one in the forward memory length estimation problem even in the class of Markov chains on a countable alphabet. In this section we shall show something similar in the class of binary (i.e. 0, 1) valued finitarily Markov processes. To prove this we will assume that there is given a sequence of estimators and stopping times, (h_n, λ_n) that do succeed to estimate successfully the memory length for binary Markov chains of finite order and construct a finitarily Markovian binary process on which the scheme fails infinitely often. This differs from the proof outline of the previous section. There a contradiction was reached showing that the purported estimators do not exist. In the present case, as we remarked in the opening paragraph of Section 5 there does exist a sequence of estimators h_n which eventually succeed in giving the memory length almost surely for all binary Markov chains of finite order. Here is a precise statement:

Theorem 7. *For any strictly increasing sequence of stopping times $\{\lambda_n\}$ and sequence of estimators $\{h_n(X_0, \dots, X_{\lambda_n})\}$, such that for all stationary and ergodic binary Markov chains with arbitrary finite order, $\lim_{n \rightarrow \infty} \frac{\lambda_n}{n} = 1$, and*

$$\lim_{n \rightarrow \infty} |h_n(X_0, \dots, X_{\lambda_n}) - K(X_0^{\lambda_n})| = 0 \quad \text{almost surely}$$

there is a stationary, ergodic finitarily Markovian binary time series such that on a set of positive measure of process realizations

$$h_n(X_0, \dots, X_{\lambda_n}) \neq K(X_{-\infty}^{\lambda_n})$$

infinitely often.

Proof. First we define the same Markov-chain as in Ryabko [23] (cf. also Györfi, Morvai, Yakowitz [8], Morvai and Weiss [17]) which serves as the technical tool for construction of our counterexample. Let the state space S be the nonnegative integers and define the transition probabilities $p_{i,j}$ as follows: $p_{0,1} = p_{1,2} = 1$, and for all $s > 1$: $p_{s,0} = p_{s,s+1} = \frac{1}{2}$.

This construction yields a stationary and ergodic Markov chain $\{M_i\}$ with stationary distribution

$$P(M = 0) = P(M = 1) = \frac{1}{4}$$

and

$$P(M = i) = \frac{1}{2^i} \quad \text{for } i \geq 2.$$

We shall construct a finitarily Markovian process X_n by defining a certain function f from the state space S to $\{0, 1\}$ and setting $X_n = f(M_n)$. We will ensure that it is finitarily Markovian by taking care that $f(0) = f(1) = 0$, $f(2) = 1$ and for all $s > 2$ if $f(s) = 0$ then $f(s+1) = 1$. Thus, in the X_n process whenever one observes two successive zeroes and a one, it is known that the underlying states in the Markov chain were 012. Note that if there is an integer K such that $f(i) = 1$ for all $i \geq K - 1$ then the process $\{X_n\}$ is a binary Markov-chain with order not greater than K . (Indeed, the probabilities $P(X_n = 1 | X_0, \dots, X_{n-1})$ are determined by the last K bits $(X_{n-K}, \dots, X_{n-1})$.)

We will define f in stages using the stopping times and estimators that the hypotheses of the theorem give us. At stage j there will be an $f^{(j)}$ and we will define a binary-valued process, $\{X_i^{(j)}\}$ by the formula: $X_i^{(j)} = f^{(j)}(M_i)$ where $f^{(j)}$ will be a $\{0, 1\}$ valued function of the state space S which is eventually one. As remarked, this ensures that all these processes are actually finite order Markov chains. The desired f will be the limit of these $f^{(j)}$'s and it will take the value 0 infinitely often. Now for the definition.

For all $0 \leq j \leq \infty$, set $f^{(j)}(0) = 0$, $f^{(j)}(1) = 0$, and $f^{(j)}(2) = 1$.

Define $f^{(0)}(k) = 1$ for all $k \geq 3$, hence since $f^{(0)}(i)$ is eventually 1, the process $\{X_i^{(0)} = f^{(0)}(M_i)\}$ is a stationary ergodic binary Markov chain with order $k_0 \leq 3$.

Recalling the stopping times and estimators define the event

$$A_1(t_1, s_1) = \left\{ \text{For some } n: h_n(f^{(0)}(M_0), \dots, f^{(0)}(M_{\lambda_n})) \leq k_0, \right. \\ \left. f^{(0)}(M_i) = 1 \text{ for } \lambda_n - k_0 + 1 \leq i \leq \lambda_n, t_1 \leq \lambda_n \leq s_1 \right\}.$$

Notice that this is a well defined event in the sample space of the Markov chain M_n . All of the events that we are about to define are in that one fixed sample space, only the function $f^{(i)}$ will be changing. By the hypotheses of the theorem there are sufficiently large $s_1 > t_1 > 3$ such that the probability

$$P(A_1(t_1, s_1) | M_0 = 0, M_1 = 1, M_2 = 2) > 1 - 2^{-1}.$$

Let $f^{(1)}(i) = f^{(0)}(i)$ for $i = 0, 1, \dots, s_1$ and let $f^{(1)}(s_1 + 1) = 0$, $f^{(1)}(i) = 1$ for $i \geq s_1 + 2$. It is clear that the memory of a sequence with prefix $1_{-k_0-1}^0$ in the process $X_n^{(1)} = f^{(1)}(M_n)$ is greater than k_0 but if the event A_1 occurs then the estimator will commit an error at least once in the interval $[t_1, s_1]$. The new process has an order $k_1 \leq s_1 + 3$. We will continue in this manner inductively. Assuming that we have already defined k_j, t_j, s_j, A_j we will now go to stage $j + 1$ and show how to update these parameters.

Let $s_{j+1} > t_{j+1} > s_j + 3$ be chosen such that for the event

$$A_{j+1}(t_{j+1}, s_{j+1}) = \left\{ \text{For some } n: h_n(f^{(j)}(M_0), \dots, f^{(j)}(M_{\lambda_n})) \leq k_j, \right. \\ \left. f^{(j)}(M_i) = 1 \text{ for } \lambda_n - k_j + 1 \leq i \leq \lambda_n, t_{j+1} \leq \lambda_n \leq s_{j+1} \right\}$$

we have:

$$P(A_{j+1}(t_{j+1}, s_{j+1}) | M_0 = M_1 = 0, M_2 = 1) > 1 - 2^{-(j+1)}.$$

Set now $f^{(j+1)}(i) = f^{(j)}(i)$ for $i = 0, 1, \dots, s_{j+1}$ and let $f^{(j+1)}(s_{j+1} + 1) = 0$ and $f^{(j+1)}(i) = 1$ for $i \geq s_{j+1} + 2$. It is clear that the memory of a sequence with suffix a string of 1's of length k_j in the process $X_n^{(j+1)} = f^{(j+1)}(M_n)$ is greater than k_j and if the event A_{j+1} happens then the estimator will commit error at least once in the interval $[t_{j+1}, s_{j+1}]$. The new process has an order $k_{j+1} \leq s_{j+1} + 3$. By induction, we have defined all the functions $f^{(j)}$ for $0 \leq j < \infty$. To complete the definition of f simply put $f = \lim f^{(j)}$. By the construction this is certainly well defined.

By the Borel–Cantelli lemma, conditioned on the positive probability event $M_0 M_1 M_2 = 012$, the events A_j occur infinitely often almost surely and this completes the proof of Theorem 7. \square

Remark 1. In the final process X_n that we constructed $P(K(X_{-\infty}^0) = k)$ decays to zero exponentially fast and in particular is summable. It follows that with probability one eventually $K(X_0^n) \leq n$ so that the reason for our failure to estimate the memory length correctly is not coming about because we do not even see the memory word.

It is also worth pointing out the sequence of moments on which the estimator is failing is of density zero. It follows fairly easily from the ergodic theorem that if one is willing to tolerate such failures then a straightforward application of any backward estimation scheme will converge outside a set of density zero. The effort that we expended in Sections 3, 4 to achieve density $1 - \epsilon$ for the stopping was because eventually we wanted to guarantee that there would be no failures at all.

7. Forward estimation of the conditional probability for finitarily Markovian processes

Let the alphabet be finite or countably infinite. Now our goal is to estimate the conditional probability $P(X_{n+1} = x | X_0^n)$ on stopping times in a pointwise sense.

Let \mathcal{N} be a sequence of stopping times such that eventually almost surely $X_{n-K(X_{-\infty}^n)+1}^n$ appears at least $n^{1-\gamma}$ times in X_0^n .

Let ρ_n be any estimate of the length of the memory from samples X_0^n such that $\rho_n - K(X_{-\infty}^n) \rightarrow 0$ on \mathcal{N} .

Define our estimate $\hat{q}_n(x)$ of the conditional probability $P(X_{n+1} = x | X_0^n)$ on \mathcal{N} as

$$\hat{q}_n(x) = \frac{\#\{\rho_n - 1 \leq i < n: X_{i-\rho_n+1}^i = X_{n-\rho_n+1}^n, X_{i+1} = x\}}{\#\{\rho_n - 1 \leq i < n: X_{i-\rho_n+1}^i = X_{n-\rho_n+1}^n\}}.$$

Theorem 8. *On $n \in \mathcal{N}$,*

$$|\hat{q}_n(x) - P(X_{n+1} = x | X_0^n)| \rightarrow 0 \quad \text{almost surely.}$$

Corollary 1. *For the stopping times \mathcal{N} and estimator ρ_n in Theorem 4, Theorem 8 holds and the density of \mathcal{N} is at least $1 - \epsilon$.*

To prove the above theorem we define Markov estimators of the conditional probabilities. Use $\lambda_{n,K(X_{-\infty}^n),i}^-$ defined in (2). Define the Markov estimator using j samples as

$$q_n^j(x) = \frac{1}{j} \sum_{i=1}^j 1_{\{X_{n-\lambda_{n,K(X_{-\infty}^n),i}^-}+1} = x\}}. \quad (17)$$

Lemma 7. *Almost surely,*

$$\max_{j \geq \lfloor n^{1-\gamma} \rfloor} |q_n^j(x) - P(X_{n+1} = x | X_{-\infty}^n)| \rightarrow 0.$$

Proof. Since by Lemma 1, $q_n^j(x)$ is an average of independent and identically distributed bounded random variables so one may apply Hoeffding's inequality (cf. Hoeffding [9] or Theorem 8.1 of Devroye et al. [7]):

$$\sum_{j=\lfloor n^{1-\gamma} \rfloor}^{\infty} P(|q_n^j(x) - P(X_{n+1} = x | X_{-\infty}^n)| > \epsilon) \leq \sum_{j=\lfloor n^{1-\gamma} \rfloor}^{\infty} 2e^{-2\epsilon^2 j}.$$

The right-hand side is summable in n and the Borel–Cantelli lemma yields Lemma 7. The proof of Lemma 7 is complete. \square

Proof of Theorem 8. Since eventually, for $n \in \mathcal{N}$, $X_{n-K(X_{-\infty}^n)+1}^n$ appears $n^{1-\gamma}$ times in X_0^n so $\hat{q}_n(x) = q_n^j(x)$ for some $j \geq \lfloor n^{1-\gamma} \rfloor$, and $P(X_{n+1} = x | X_0^n) = P(X_{n+1} = x | X_{-\infty}^n)$, the result follows from Lemma 7. The proof of Theorem 8 is complete. \square

8. Forward estimation of the conditional probability for Markov processes

Let $\{X_n\}$ be a stationary and ergodic finite or countably infinite alphabet Markov chain with order K . Let $ORDEST_n$ be an estimator of the order from samples X_0^n such that $ORDEST_n \rightarrow K$ almost surely. Such an estimator can be found e.g. in Morvai and Weiss [16]. Let $n \in \mathcal{N}$ if $X_{n-ORDEST_{n+1}}^n$ appears at least $n^{1-\gamma}$ times in X_0^n . \mathcal{N} is a sequence of stopping times. Let

$$\hat{q}_n(x) = \frac{\#\{ORDEST_n - 1 \leq i < n: X_{i-ORDEST_{n+1}}^i = X_{n-ORDEST_{n+1}}^n, X_{i+1} = x\}}{\#\{ORDEST_n - 1 \leq i < n: X_{i-ORDEST_{n+1}}^i = X_{n-ORDEST_{n+1}}^n\}}.$$

Theorem 9. *Assume $ORDEST_n$ equals the order eventually almost surely. Then on $n \in \mathcal{N}$,*

$$|\hat{q}_n(x) - P(X_{n+1} = x | X_{n-K}^n)| \rightarrow 0 \quad \text{almost surely}$$

and

$$\liminf_{n \rightarrow \infty} \frac{|\mathcal{N} \cap \{0, 1, \dots, n-1\}|}{n} = 1.$$

If the Markov chain turns out to take values from a finite set, then \mathcal{N} takes as values all but finitely many positive integers.

To prove the above theorem we define Markov estimators of the conditional probabilities. Use $\lambda_{n,K,i}^-$ defined in (2). Define the Markov estimator using j samples as

$$q_n^j(x) = \frac{1}{j} \sum_{i=1}^j 1_{\{X_{n-\lambda_{n,K,i}^-} = x\}}. \quad (18)$$

Lemma 8. *Almost surely,*

$$\max_{j \geq \lfloor n^{1-\gamma} \rfloor} |q_n^j(x) - P(X_{n+1} = x | X_{n-K}^n)| \rightarrow 0.$$

Proof. The proof goes along the lines of the proof of Lemma 7. The proof of Lemma 8 is complete. \square

Proof of Theorem 9. Since $ORDEST_n = K$ eventually, and so for $n \in \mathcal{N}$: X_{n-K+1}^n appears at least $n^{1-\gamma}$ times thus $\hat{q}_n(x) = q_n^j(x)$ for some $j \geq \lfloor n^{1-\gamma} \rfloor$ and the result follows from Lemma 8. Since any word of length K with positive probability appears eventually almost surely $n^{1-\gamma}$ times in X_0^n thus \mathcal{N} has density one. If the alphabet is finite, then the number of words with length K is finite and by ergodicity, eventually almost surely all words with length K which has positive probability appears at least $n^{1-\gamma}$ times. The proof of Theorem 9 is complete. \square

References

- [1] D.H. Bailey, Sequential schemes for classifying and predicting ergodic processes, Ph.D. thesis, Stanford University, 1976.
- [2] P. Bühlmann, A.J. Wyner, Variable-length Markov chains, *Ann. Statist.* 27 (1999) 480–513.
- [3] I. Csiszár, Large-scale typicality of Markov sample paths and consistency of MDL order estimators, *IEEE Trans. Inform. Theory* 48 (2002) 1616–1628.
- [4] I. Csiszár, P. Shields, The consistency of the BIC Markov order estimator, *Ann. Statist.* 28 (2000) 1601–1619.
- [5] I. Csiszár, Zs. Talata, Context tree estimation for not necessarily finite memory processes via BIC and MDL, *IEEE Trans. Inform. Theory*, in press.
- [6] A. Dembo, Y. Peres, A topological criterion for hypothesis testing, *Ann. Statist.* 22 (1994) 106–117.
- [7] L. Devroye, L. Györfi, G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer-Verlag, New York, 1996.
- [8] L. Györfi, G. Morvai, S. Yakowitz, Limits to consistent on-line forecasting for ergodic time series, *IEEE Trans. Inform. Theory* 44 (1998) 886–892.
- [9] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Amer. Statist. Assoc.* 58 (1963) 13–30.
- [10] S. Kalikow, Y. Katznelson, B. Weiss, Finitarily deterministic generators for zero entropy systems, *Israel J. Math.* 79 (1992) 33–45.
- [11] G. Morvai, Guessing the output of a stationary binary time series, in: Y. Haitovsky, H.R. Lerche, Y. Ritov (Eds.), *Foundations of Statistical Inference*, Physika-Verlag, 2003, pp. 207–215.
- [12] G. Morvai, S. Yakowitz, L. Györfi, Nonparametric inference for ergodic, stationary time series, *Ann. Statist.* 24 (1996) 370–379.
- [13] G. Morvai, B. Weiss, Forecasting for stationary binary time series, *Acta Appl. Math.* 79 (2003) 25–34.
- [14] G. Morvai, B. Weiss, Intermittent estimation of stationary time series, *Test* 13 (2004) 525–542.
- [15] G. Morvai, B. Weiss, Prediction for discrete time series, *Probab. Theory Related Fields* 132 (2005) 1–12.
- [16] G. Morvai, B. Weiss, Order estimation of Markov chains, *IEEE Trans. Inform. Theory* 51 (2005) 1496–1497.
- [17] G. Morvai, B. Weiss, Limitations on intermittent forecasting, *Statist. Probab. Lett.* 72 (2005) 285–290.
- [18] G. Morvai, B. Weiss, On classifying processes, *Bernoulli* 11 (2005) 523–532.
- [19] G. Morvai, B. Weiss, Inferring the conditional mean, *Theory Stochastic Process.* 11 (1–2) (2005) 112–120.
- [20] A. Nobel, Limits to classification and regression estimation from ergodic processes, *Ann. Statist.* 27 (1999) 262–273.
- [21] D.S. Ornstein, Guessing the next output of a stationary process, *Israel J. Math.* 30 (1978) 292–296.
- [22] D.S. Ornstein, B. Weiss, How sampling reveals a process, *Ann. Probab.* 18 (1990) 905–930.
- [23] B.Ya. Ryabko, Prediction of random sequences and universal coding, *Problems Inform. Trans.* 24 (April–June 1988) 87–96.
- [24] P.C. Shields, *The Ergodic Theory of Discrete Sample Paths*, Grad. Stud. Math., vol. 13, American Mathematical Society, Providence, RI, 1996.