

Redundancy of Markov Family with Unbounded Memory

Changlong Wu Maryam Hosseini Narayana Santhanam
 University of Hawaii at Manoa
 Honolulu
 96822 Hawaii, USA
 Email: {wuchangl, hosseini, nsanthan}@hawaii.edu

June 21, 2018

Abstract

We study the redundancy of universally compressing strings X_1, \dots, X_n generated by a binary Markov source p without any bound on the memory. To better understand the connection between compression and estimation in the Markov regime, we consider a class of Markov sources restricted by a *continuity* condition. In the absence of an upper bound on memory, the continuity condition implies that $p(X_0|X_{-m}^{-1})$ gets closer to the true probability $p(X_0|X_{-\infty}^{-1})$ as m increases, rather than vary around arbitrarily. For such sources, we prove asymptotically matching upper and lower bounds on the redundancy. In the process, we identify what sources in the class matter the most from a redundancy perspective.

1 Introduction

Estimation and compression are virtually synonymous in the *i.i.d.* regime. Indeed, in the *i.i.d.* case, the add-half (and other add-constant) estimators that provide reasonable estimates of probabilities of various symbols are described naturally using a universal compression setup. These estimators simply correspond to the conditional probabilities assigned by a Bayesian mixture when Dirichlet priors on the parameter space—and indeed encoding the probabilities given by these Bayesian mixtures yields good universal compression schemes for these classes of distributions.

In the Markov setup, there is an additional complication not seen in the *i.i.d.* setup—*mixing*—that complicates the relation between compression and estimation. Without going into the technicalities of mixing, slow mixing sources do not explore the state space efficiently.

For example, consider a memory-1 binary Markov source that assigns conditional probability of $1 - \epsilon$ for a 1 given 1, and ϵ for the conditional probability of a 1 given 0. If we start the source from the state 1, for sample length $n \ll \frac{1}{\epsilon}$, we will see a sequence of all 1s with high probability in length- n samples. This sequence of 1s is, of course, easy to compress—but clearly precludes the estimation of the conditional probabilities associated with 0.

Previous work in the Markov regime, however, has typically considered classes Markov sources with bounded memory (say, all memory-3 Markov sources) as a natural hierarchy of classes. As the prior example shows, these classes are definitely not natural from an estimation point of view. Small memory sources—even with memory one can be arbitrarily slow mixing—and hence hard to estimate. On the other hand, sources with longer memory may be easier to estimate if they are fast mixing and satisfy certain other conditions, as we will see below.

As a consequence, we look for a different way to resolve the class of all finite memory Markov sources into a nested hierarchy of classes. Therefore, in [1], a new class of Markov sources was introduced, one that was more amenable to estimation. These classes of sources were not bounded in memory, rather they can have arbitrarily long memory. However, these sources satisfy a continuity condition [1, 2], as described technically in Section 2.1 in the paper.

Roughly speaking, the continuity condition imposes two intuitive constraints closely related to each other. Let p be a binary Markov source with finite but unknown memory, and consider $p(X_0|X_{-\infty}^{-1})$. Because the source has finite memory, there is a suffix of the past, $X_{-\infty}^{-1}$ that determines the conditional probability above. Since we do not have an a-priori bound on the memory, we cannot say how much of the past we need. Yet, the conditional probabilities $p(X_0|X_{-m}^{-1})$ are well defined for all m . It is now possible to construct Markov sources where, unless we have suffixes long enough to encapsulate the true state (or equivalently, m larger than the true memory), $p(X_0|X_{-m}^{-1})$ is simply not a reflection of the true probability $p(X_0|X_{-\infty}^{-1})$.

The continuity condition prohibits this pathological property—imposing on the other hand that the more of the context (X_{-m}^{-1}) we see, the better $p(X_0|X_{-m}^{-1})$ reflects $p(X_0|X_{-\infty}^{-1})$. Second, given a history X_{-m}^{-1} , the continuity condition implies that the conditional information one more bit in the past, X_{-m-1}^{-1} , provides on X_0 diminishes with m .

The continuity condition may be imagined as a soft constraint on the memory, but it does not control mixing. Suppose we consider the collection of all Markov sources that satisfy a given continuity condition. It was shown in [1] that it is possible to use length- n samples to estimate the conditional probabilities $p(1|\mathbf{w})$ of all strings \mathbf{w} of length $\log n$ that appear in the sample, as well as provide deviation bounds on the estimates.

This hints that Markov sources nested by the continuity condition (as opposed to memory) are a natural way to break down the collection of all Markov sources. In order to better understand these model classes, in this paper we study compressing Markov sources constrained by the continuity condition, and obtain asymptotically tight bounds on their redundancy.

Part of the reason to study this is to understand what portions of the model classes are more important (namely, contribute primarily) to the redundancy. Indeed, it turns out that the primary contribution to the redundancy comes from essentially fast mixing sources whose state probabilities that are not towards the extremes (not near 0 or 1), while, of course, these sources do not complicate estimation at all. On the other hand, slow mixing sources hit estimation, but do not matter for compression at all—a dichotomy which was hinted at with our very first memory-1 example.

Our main results are matching lower and upper bounds, in Sections 3 and 4 respectively. In Section 2, we set up notations, and briefly describe the continuity condition in 2.1.

1.1 Prior Work

Davisson formulated the average and worst case redundancy in his seminal paper in 1973 [3]. A long sequence of work has characterized the worst case redundancy for memoryless sources [4] [5] [6] [7]. Average redundancy of Markov sources with fixed memory has been studied in [8] [9] [10]. In [11], the authors obtain the worst case redundancy of such Markov sources and later [12] derived the exact worst case redundancy of such Markov sources. The estimation and compression of finite memory context tree models was studied in [13] and [14]. [2] and [15] studied the estimation of context tree with unbounded memory.

For a different comparison of estimation and compression in Markov settings, see [16]. Here the authors obtain the redundancy of conditionally describing one additional symbol after obtaining a sample. Finally, [17] considers compression of patterns of Hidden Markov models.

2 Setup and Notation

We denote length- n strings in bold face \mathbf{x} or as x_1^n , and their random counterparts are \mathbf{X} and X_1^n respectively. For any $\mathbf{x} \in \{0,1\}^n$ and $\mathbf{w} \in \{0,1\}^\ell$ with $\ell \leq n$, let $n_{\mathbf{w}}$ be the number of appearances of \mathbf{w} in \mathbf{x} as a *substring*. For $\mathbf{x}, \mathbf{y} \in \{0,1\}^*$, $\mathbf{x} \prec \mathbf{y}$ denotes that \mathbf{x} is a suffix of \mathbf{y} (e.g., $010 \prec 11010$), \mathbf{xy} is the concatenation of \mathbf{x} followed by \mathbf{y} (e.g., if $\mathbf{x} = 010$ and $\mathbf{y} = 1$ then $\mathbf{xy} = \mathbf{x}1 = 0101$), and $|\mathbf{x}|$ as the length of \mathbf{x} .

We consider a two-sided infinite binary Markov random process p generating a sample $\dots, X_{-1}, X_0, X_1, \dots$. The memory of p is finite (though not bounded a-priori), and let S_p to be *context tree* [5] of source p . S_p is a set of leaves of a complete binary tree and p is completely described by the conditional (transition) probabilities $p(1 | \mathbf{s})$ for $\mathbf{s} \in S_p$.

Let \mathcal{M}^ℓ to be all the Markov chains with memory at most ℓ , and $\mathcal{M} = \cup_\ell \mathcal{M}^\ell$ be the family of all the finite memory Markov chains. As mentioned before, while \mathcal{M}^ℓ is a natural class, we are looking to break \mathcal{M} down into a more natural hierarchy of classes.

2.1 Markov chain with continuity condition

Let $\delta := \mathbb{N} \rightarrow \mathbb{R}^+$ be a function such that $\delta(n) \rightarrow 0$ as $n \rightarrow \infty$. A Markov chain p satisfies the *continuity condition* subject to δ if for all $\mathbf{s}_1, \mathbf{s}_2 \in S_p$, and $a \in \{0, 1\}$, we have

$$\left| \frac{p(a | \mathbf{s}_1)}{p(a | \mathbf{s}_2)} - 1 \right| \leq \delta(|\mathbf{w}|)$$

for all $\mathbf{w} \in \{0, 1\}^*$ such that $\mathbf{w} \prec \mathbf{s}_1$ and $\mathbf{w} \prec \mathbf{s}_2$ (namely \mathbf{w} is a common suffix of \mathbf{s}_1 and \mathbf{s}_2). For technical reasons, we will assume that $\delta(n) < \frac{1}{4n}$, see [1].

Denote \mathcal{M}_δ to be the family of all the Markov chains with continuity condition subject to δ and $\mathcal{M}_\delta^\ell = \mathcal{M}_\delta \cap \mathcal{M}^\ell$. Roughly speaking, the continuity condition constraints the transition probabilities of states with long common suffix to be close.

Let $S_p(\mathbf{w}) = \{\mathbf{s} \in S_p : \mathbf{w} \prec \mathbf{s}\}$. Clearly we have that the stationary probability $p(\mathbf{w}) = \sum_{\mathbf{s} \in S_p(\mathbf{w})} p(\mathbf{s})$ and that $p(1 | \mathbf{w}) = (\sum_{\mathbf{s} \in S_p(\mathbf{w})} p(1 | \mathbf{s})p(\mathbf{s})) / p(\mathbf{w})$. Let p have memory ℓ . In analogy to the true conditional probability $p(1 | \mathbf{w})$, for a given $\mathbf{x} \in \{0, 1\}^n$ and past $x_{-\ell+1}^0$, let

$$\tilde{p}(1 | \mathbf{w}) = \frac{\sum_{\mathbf{s} \in S_p(\mathbf{w})} n_{\mathbf{s}} p(1 | \mathbf{s})}{n_{\mathbf{w}}}, \quad (1)$$

be the *empirical aggregated distribution* of p , write it as $\tilde{p}_{\mathbf{w}}$ for simplicity. In a slight abuse of notation here, $n_{\mathbf{s}}$ (respectively $n_{\mathbf{w}}$) is the number of bits in \mathbf{x} with context \mathbf{s} (respectively \mathbf{w}), *i.e.*, number of bits in \mathbf{x} immediately preceded by \mathbf{s} (respectively \mathbf{w}) when taking the past $x_{-\ell+1}^0$ into account).

2.2 The redundancy

For any distribution family \mathcal{P} on $\{0, 1\}^n$, the worst case minimax redundancy of \mathcal{P} is defined as

$$R(\mathcal{P}) = \inf_q \sup_{p \in \mathcal{P}} \max_{\mathbf{x} \in \{0, 1\}^n} \log \frac{p(\mathbf{x})}{q(\mathbf{x})},$$

similarly, the average minimax redundancy is defined as

$$\tilde{R}(\mathcal{P}) = \inf_q \sup_{p \in \mathcal{P}} E_{\mathbf{X} \sim p} \log \frac{p(\mathbf{X})}{q(\mathbf{X})},$$

where q is choosing from all the possible distributions on $\{0, 1\}^n$. For any given (*fixed*) past $x_{-\infty}^0$, we know that for any $p \in \mathcal{M}$ we will have a well defined distribution over $\{0, 1\}^n$, given by

$$\bar{p}(x_1^n) = p(x_1^n | x_{-\infty}^0).$$

The main result of this paper is a lower and upper bound on the average redundancy of \mathcal{M}_δ over $\{0, 1\}^n$ for any past, *i.e.*

$$\tilde{R}(\mathcal{M}_\delta) \triangleq \inf_q \sup_{p \in \mathcal{M}_\delta} \sup_{x_{-\infty}^0} E_{\mathbf{X} \sim \bar{p}} \log \frac{\bar{p}(\mathbf{X})}{q(\mathbf{X})}.$$

3 The lower bound

The *Redundancy-Capacity* theorem [18] is a common approach to lower bound the minimax redundancy. This approach gets complicated in our case since there is no universal bound on the memory of sources in \mathcal{M}_δ , rendering the parameter space to be infinite dimensional. We therefore first consider \mathcal{M}_δ^ℓ (see Section II), the subset of sources in \mathcal{M}_δ which also have finite memory $\leq \ell$, and obtain a lower bound on $\tilde{R}(\mathcal{M}_\delta^\ell)$. Since $\tilde{R}(\mathcal{M}_\delta) \geq \tilde{R}(\mathcal{M}_\delta^\ell)$ for all ℓ , we optimize over ℓ to obtain the best possible lower bound on $\tilde{R}(\mathcal{M}_\delta^\ell)$. Recall that $\delta(\ell) \leq 1/(4\ell)$.

Theorem 1 (Lower Bound). For any ℓ , we have

$$\tilde{R}(\mathcal{M}_\delta^\ell) \geq 2^{\ell-1} \log n - 2^\ell (\log \frac{1}{\delta(\ell)} + \ell/2) - 2^{\ell-1} (\log 4\pi e\ell + 1),$$

and $\tilde{R}(\mathcal{M}_\delta) \geq \max_\ell \tilde{R}(\mathcal{M}_\delta^\ell)$. \square

Before proving this theorem, we consider specific forms for δ to get an idea of the order of magnitude of the redundancy in Theorem 1.

Corollary 2. If $\delta(\ell) = \frac{1}{\ell^c}$ with $c > 1$, then we have

$$\tilde{R}(\mathcal{M}_\delta) = \Omega(n/\log^{2c-1} n),$$

for $\ell = \log n - 2c \log \log n + o(1)$. If $\delta(\ell) = 2^{-c\ell}$, then

$$\tilde{R}(\mathcal{M}_\delta) = \Omega(n^{1/(2c+1)} \log n),$$

for $\ell = \frac{1}{2c+1} \log n + o(1)$. If $\delta(\ell) = 2^{-2^\ell}$, then

$$\tilde{R}(\mathcal{M}_\delta) = \Omega(\log^{1+1/c} n),$$

for $\ell = \frac{1}{c} \log \log n + o(1)$. \square

For any $p \in \mathcal{M}_\delta^\ell$, we know that the distribution of p on $\{0, 1\}^n$ can be uniquely determined by at most 2^ℓ parameters, i.e. the transition probabilities $p(1 | \mathbf{s})$. Let

$$\hat{\mathcal{M}}_\delta^\ell = \{p \in \mathcal{M}_\delta^\ell \mid \forall \mathbf{s} \in S_p, |p(1 | \mathbf{s}) - 1/2| \leq \delta(\ell)\},$$

be a sub-family of \mathcal{M}_δ^ℓ with all the transition probabilities close to 1/2. The following lemma is directly from *Redundancy-Capacity* theorem [18].

Lemma 3. Let $\epsilon_{\mathbf{s}}$ be the maximum mean square error of estimating parameters $p \in \hat{\mathcal{M}}_\delta^\ell$ from their length n sample. Then

$$\tilde{R}(\hat{\mathcal{M}}_\delta^\ell) \geq 2^\ell \log \delta(\ell) - 2^{\ell-1} \log (2\pi e \epsilon_{\mathbf{s}}). \quad (2)$$

Proof

Consider the following Markov chain

$$\hat{\mathcal{M}}_\delta^\ell \xrightarrow{(a)} P \xrightarrow{(b)} X_{-\ell+1}^n \xrightarrow{(c)} \hat{P},$$

where (a) P is a random Markov process chosen from a uniform prior over $\hat{\mathcal{M}}_\delta^\ell$, (b) X_1^n is a length n sample from distribution P , (c) \hat{P} is an estimate of P from the sample X_1^n that uses the empirical probabilities $\frac{n_{\mathbf{s}1}}{n_{\mathbf{s}}}$ to estimate $P(1 | \mathbf{s})$ for any $\mathbf{s} \in \{0, 1\}^\ell$.

By capacity-redundancy theorem one knows that

$$\tilde{R}(\hat{\mathcal{M}}_\delta^\ell) \geq I(P; X_1^n) \geq I(P; \hat{P}),$$

where the second inequality is by data processing inequality. Note that

$$\begin{aligned}
I(P; \hat{P}) &= h(P) - h(P|\hat{P}) \\
&= h(P) - h(P - \hat{P}|\hat{P}) \\
&\geq h(P) - h(P - \hat{P}),
\end{aligned} \tag{3}$$

where the last inequality follows since conditioning reduces entropy. To bound first term in (3) let $P \in \hat{\mathcal{M}}_\delta^\ell$ be uniform on the hypercube A with edge lengths $\delta(\ell)$. Then

$$h(P) = 2^\ell \log \delta(\ell).$$

since $h(P) = \log \text{Vol}(A)$.

To bound $h(P - \hat{P})$, let K be the covariance matrix of any estimator of parameter space condition on \mathbf{x}^n . Then using Theorem 8.6.5 in [19]

$$h(P - \hat{P}) \leq \frac{1}{2} \log(2\pi e)^{2^\ell} |K|.$$

Let $|K|$ and λ_i show determinant and eigenvalues of matrix K , respectively. Let ϵ_i be the element diagonal elements of covariance matrix. Then by the definition of trace of a matrix

$$\sum_i \epsilon_i = \text{tr}(K) = \sum_i \lambda_i.$$

Using arithmetic-geometric inequality, we get

$$\left(\frac{\sum_i \lambda_i}{2^\ell}\right)^{2^\ell} \geq \prod_i \lambda_i. \tag{4}$$

Also $\sum_i \epsilon_i \leq 2^\ell \epsilon_s$. Then

$$|K| = \prod_i \lambda_i \leq \left(\frac{\sum_i \epsilon_i}{2^\ell}\right)^{2^\ell} \leq \epsilon_s^{2^\ell}. \tag{5}$$

Applying (5) in $\frac{1}{2} \log(2\pi e)^{2^\ell} |K|$, we have

$$h(P - \hat{P}) \leq 2^{\ell-1} \log(2\pi e \epsilon_s).$$

and lemma follows. \square

To bound ϵ_s one needs to find an estimator that makes it as small as possible. We will show that the empirical estimation $\hat{P}(\mathbf{s}) = \frac{n_{s1}}{n_s}$, is sufficient to establish our lower bound. We now find an upper bound on the estimation error of state \mathbf{s} using naive estimator.

Lemma 4. Consider the naive estimator $\hat{P}(\mathbf{s}) = \frac{n_{s1}}{n_s}$. Then,

$$E[(\hat{P}(\mathbf{s}) - P(\mathbf{s}))^2] \leq \min\left\{E\left[\frac{1}{n_s}\right], 1\right\}.$$

Proof Note that

$$E[(\hat{P}(\mathbf{s}) - P(\mathbf{s}))^2] = E[E[(\hat{P}(\mathbf{s}) - P(\mathbf{s}))^2 | n_s]].$$

Condition on n_s , the symbols follow string \mathbf{s} can be considered as outputs of an *i.i.d.* Bernoulli with parameter $P(\mathbf{s})$. For a sequence of zeros and ones with length n drawn *i.i.d.* from $B(p)$ with k one, it is easy to see that $E[(\frac{k}{n} - p)^2] \leq \frac{1}{n}$, so

$$E[E[(\hat{P}(\mathbf{s}) - P(\mathbf{s}))^2 | n_s]] \leq \min\left\{E\left[\frac{1}{n_s}\right], 1\right\}.$$

\square

So finding the lower bound on redundancy reduces to find an upper bound on $E[\frac{1}{n_s}]$. We need two following technical lemmas to bound $E[\frac{1}{n_s}]$.

Lemma 5. Let X_1, X_2, \dots, X_n be binary random variables, such that for any $1 \leq t \leq n$

$$\Pr(X_t = 1 | X_1, \dots, X_{t-1}) \geq q,$$

for some $q \in [0, 1]$. Then, for any $1 \leq k \leq n$

$$\Pr\left(\sum_{i=1}^n X_i \leq k\right) \leq \sum_{i=0}^k \binom{n}{i} q^i (1-q)^{n-i}.$$

Proof We use double induction on k and n to prove this theorem.

Consider the base case for $k = 0$, and an arbitrary n , then we need to bound $\Pr(\sum_{i=1}^n X_i \leq 0)$, which is equal to say that $\Pr(X_1 = 0, X_2 = 0, \dots, X_n = 0)$. But

$$\begin{aligned} \Pr(X_1 = 0, \dots, X_n = 0) &= \Pr(X_n = 0 | X_1 = 0, \dots, X_{n-1} = 0) \dots \Pr(X_1 = 0) \\ &\leq (1-q)^n \end{aligned}$$

where the first equation is using chain rule and the inequality follows by the assumption that

$$\Pr(X_n = 1 | X_1, \dots, X_{n-1}) \geq q.$$

If $k = n$ then

$$\sum_{i=0}^n \binom{n}{i} q^i (1-q)^{n-i} = 1$$

so we need to have $\Pr(\sum_{i=1}^n X_i \leq k) < 1$ which holds always.

For the induction step we just need to show that if $(n', k') \leq (n, k)$ and

$$\Pr\left(\sum_{i=1}^{n'} X_i \leq k'\right) \leq \sum_{i=0}^{k'} \binom{n'}{i} q^i (1-q)^{n'-i},$$

holds, then

$$\Pr\left(\sum_{i=1}^n X_i \leq k\right) \leq \sum_{i=0}^k \binom{n}{i} q^i (1-q)^{n-i}$$

holds. To see it, define

$$\begin{aligned} A_k^n &= \left\{ \sum_{i=1}^n X_i \leq k \right\}, \\ B_k^n &= \left\{ X_1 = 0 \wedge \sum_{i=1}^n X_i \leq k \right\}, \quad \text{and} \\ C_k^n &= \left\{ X_1 = 1 \wedge \sum_{i=1}^n X_i \leq k-1 \right\}. \end{aligned}$$

Define

$$T_k^n = \sum_{i=0}^k \binom{n}{i} q^i (1-q)^{n-i}.$$

Using chain rule we have

$$\Pr\{B_k^n\} = \Pr\{X_1 = 0\} \Pr\left\{\sum_{i=1}^n X_i \leq k | X_1 = 0\right\}$$

$$\Pr\{C_k^n\} = \Pr\{X_1 = 1\} \Pr\{\sum_{i=1}^n X_i \leq k-1 | X_1 = 1\}.$$

Let $P(X_1 = 1) = \tilde{q} > q$, then $P(X_1 = 0) = 1 - \tilde{q} < 1 - q$.

Note that $A_k^n = B_k^n \cup C_k^n$. Using union bound and since B_k^n and C_k^n are disjoint,

$$\Pr\{A_k^n\} = \Pr\{B_k^n\} + \Pr\{C_k^n\},$$

Then

$$\begin{aligned} \Pr\{A_k^n\} &= (1 - \tilde{q}) \Pr\{\sum_{i=1}^n X_i \leq k | X_1 = 0\} + \tilde{q} \Pr\{\sum_{i=1}^n X_i \leq k-1 | X_1 = 1\} \\ &\stackrel{(a)}{\leq} (1 - q) T_k^{n-1} + q T_{k-1}^{n-1} \\ &= (1 - q) \sum_{i=0}^k \binom{n-1}{i} q^i (1 - q)^{n-1-i} + q \sum_{i=0}^{k-1} \binom{n}{i} q^i (1 - q)^{n-i} \\ &= \sum_{i=0}^k \binom{n}{i} q^i (1 - q)^{n-i} \end{aligned}$$

and (a) follows by using induction assumption. \square

Lemma 6. For all $p \in \hat{\mathcal{M}}_\delta^\ell$, we have

$$p\left(n_{\mathbf{s}} \leq \frac{n}{2\ell 2^\ell} - \sqrt{\frac{n \log n}{\ell 2^\ell}}\right) \leq \frac{1}{n},$$

Proof Divide length n sequence to subsequence of length l and let $m = \frac{n}{\ell}$. Let 1_i for $i \in \{1, 2, \dots, m\}$ as

$$1_i = 1\{\mathbf{s} \prec X_{(i-1)\ell}^{i\ell}\}$$

Note that

$$\begin{aligned} p(1_i = 1 | 1_0, 1_1, \dots, 1_{i-1}) &\stackrel{(a)}{\geq} \left(\frac{1}{2} - \delta(\ell)\right)^\ell \\ &\geq \frac{1}{2^\ell} (1 - 2\delta(\ell))^\ell \\ &\stackrel{(b)}{\geq} \frac{1}{2^\ell} (1 - 2\ell\delta(\ell)) \\ &\geq \frac{1}{2^{\ell+1}}, \end{aligned}$$

where (a) follows since $p(1 | \mathbf{s}') \in [1/2 - \delta(\ell), 1/2 + \delta(\ell)]$ for all $\mathbf{s}' \in S_p$ and (b) is by union bound. Let $q = \frac{1}{2^{\ell+1}}$ in Lemma 5, then

$$\Pr(n_s \leq k) \leq \sum_{i=0}^k \binom{m}{i} \left(\frac{1}{2^{\ell+1}}\right)^i \left(1 - \frac{1}{2^{\ell+1}}\right)^{m-i}.$$

Right hand side in last inequality is the probability that sum of some *i.i.d.* random variables (we denote it by S) drawn from $\mathcal{B}(\frac{1}{2^{\ell+1}})$ with mean $\mu = \frac{m}{2^{\ell+1}}$ is less than k . Let $k = (1 - \gamma)\mu$ where $0 \leq \gamma \leq 1$ is arbitrary. Then

$$\sum_{i=0}^k \binom{m}{i} \left(\frac{1}{2^{\ell+1}}\right)^i \left(1 - \frac{1}{2^{\ell+1}}\right)^{m-i} = \Pr(S \leq (1 - \gamma)\mu).$$

Using Chernoff bound, we get

$$\Pr(S \leq (1 - \gamma)\mu) \leq e^{-\frac{\gamma^2 \mu}{2}}.$$

Let $\gamma = 2\sqrt{\frac{2^\ell \ell \log n}{n}}$, then

$$e^{-\frac{\gamma^2 \mu}{2}} = e^{-2\frac{2^{\ell+1} \ell \log n}{n} \frac{n}{2^{\ell+1} \ell}} \leq \frac{1}{n}.$$

So $k = (1 - \gamma)\mu = \frac{n}{2^{\ell+1} \ell} - \sqrt{\frac{n \log n}{2^\ell \ell}}$ and lemma follows. \square

We now combine Lemma 5 and Lemma 6 to bound $E\left[\frac{1}{n_s}\right]$.

Lemma 7. For n large enough,

$$E\left[\frac{1}{n_s}\right] \leq \frac{\ell 2^{\ell+1}}{n}.$$

Proof Let $k = \frac{n}{2^{\ell+1} \ell} - \sqrt{\frac{n \log n}{2^\ell \ell}}$, then

$$\begin{aligned} E\left[\frac{1}{n_s}\right] &= \sum_{\frac{1}{n_s} \geq \frac{1}{k}} \frac{1}{n_s} p\left(\frac{1}{n_s} \geq \frac{1}{k}\right) + \sum_{\frac{1}{n_s} < \frac{1}{k}} \frac{1}{n_s} p\left(\frac{1}{n_s} < \frac{1}{k}\right) \\ &\leq \sum_{\frac{1}{n_s} \geq \frac{1}{k}} p\left(\frac{1}{n_s} \geq \frac{1}{k}\right) + \frac{1}{k} \sum_{\frac{1}{n_s} < \frac{1}{k}} p\left(\frac{1}{n_s} < \frac{1}{k}\right), \end{aligned}$$

where the inequality is by the fact that $\frac{1}{n_s} < 1$. Using Lemma 6

$$\sum_{\frac{1}{n_s} \geq \frac{1}{k}} p\left(\frac{1}{n_s} \geq \frac{1}{k}\right) < \frac{1}{n}.$$

Also $\sum_{\frac{1}{n_s} < \frac{1}{k}} p\left(\frac{1}{n_s} < \frac{1}{k}\right) < 1$. So

$$E\left[\frac{1}{n_s}\right] \leq \frac{1}{n} + \frac{1}{k}.$$

But $k = \frac{n}{2^{\ell+1} \ell} - \sqrt{\frac{n \log n}{2^\ell \ell}} = \frac{n}{2^{\ell+1} \ell} (1 - 2\sqrt{\frac{2^\ell \ell \log n}{n}})$, then

$$E\left[\frac{1}{n_s}\right] \leq \frac{1}{n} + \frac{\ell 2^{\ell+1}}{n} \left(\frac{1}{1 - 2\sqrt{\frac{2^\ell \ell \log n}{n}}} \right).$$

But we can choose n large enough so that $\sqrt{\frac{2^\ell \ell \log n}{n}} < \frac{1}{16}$, then

$$\frac{1}{1 - 2\sqrt{\frac{2^\ell \ell \log n}{n}}} < 2.$$

So

$$E\left[\frac{1}{n_s}\right] \leq \frac{1 + 2\ell 2^{\ell+1}}{n} \simeq \frac{2\ell 2^{\ell+1}}{n}.$$

\square

We are now ready to give proof of theorem 1.

Proof of Theorem 1

$$\begin{aligned}
\tilde{R}(\hat{\mathcal{M}}_\delta^\ell) &\stackrel{(a)}{\geq} 2^\ell \log \delta(\ell) - 2^{\ell-1} \log (2\pi e \epsilon_s) \\
&\stackrel{(b)}{\geq} 2^\ell \log \delta(\ell) - 2^{\ell-1} \log (2\pi e E[\frac{1}{n_s}]) \\
&\stackrel{(c)}{\geq} 2^\ell \log \delta(\ell) - 2^{\ell-1} \log (2\pi e \frac{2\ell 2^{\ell+1}}{n}) \\
&= 2^{\ell-1} \log n - 2^{\ell-1} \log 4\pi e \ell \\
&\quad - 2^{\ell-1} \log 2^{\ell+1} - 2^\ell \log \frac{1}{\delta(\ell)} \\
&= 2^{\ell-1} \log n - 2^\ell (\log \frac{1}{\delta(\ell)} + \ell/2) - 2^{\ell-1} (\log 4\pi e \ell + 1)
\end{aligned}$$

where (a) is using Lemma 3, (b) is by Lemma 4 and (c) follows by Lemma 7.

4 The upper bound

To obtain the upper bound, we first show that for any given $\mathbf{x} \in \{0, 1\}^n$, the maximum probability from any distribution in \mathcal{M}_δ will not much greater than that from \mathcal{M}_δ^ℓ for an appropriate choice of ℓ . This allows us a simple upper bound based on truncating the memory of sources. Unfortunately (or fortunately), this simple argument does not allow for tight matching bounds—hence we will need to refine our argument further.

Lemma 8. Fix any past $x_{-\infty}^0$. For any $\mathbf{x} \in \{0, 1\}^n$, let $\hat{p}(\mathbf{x}) = \max_{p \in \mathcal{M}_\delta} p(\mathbf{x} | x_{-\infty}^0)$ and $\hat{p}_\ell(\mathbf{x}) = \max_{p \in \mathcal{M}_\delta^\ell} p(\mathbf{x} | x_{-\infty}^0)$. Then

$$\hat{p}(\mathbf{x}) \leq 2^{2n\delta(\ell)} \hat{p}_\ell(\mathbf{x}).$$

Proof The continuity condition implies that for any $p \in \mathcal{M}_\delta$, we can find $p_\ell \in \mathcal{M}_\delta^\ell$ such that for all $a \in \{0, 1\}$, $\mathbf{w} \in \{0, 1\}^\ell$ and $\mathbf{s} \in S_p(\mathbf{w})$, we have

$$p(a | \mathbf{s}) \leq (1 + \delta(\ell)) p_\ell(a | \mathbf{w}).$$

Thus we have $p(\mathbf{x}) \leq (1 + \delta(\ell))^n p_\ell(\mathbf{x}) \leq 2^{2n\delta(\ell)} \hat{p}_\ell(\mathbf{x})$, where the last inequality follows since $(1 + \delta(\ell))^n \approx e^{n\delta(\ell)}$. \square

Proposition 1.

$$R(\mathcal{M}_\delta) \leq \min_\ell 2^{\ell-1} \log n + 2n\delta(\ell).$$

Proof Shtarkov's sum [20] gives

$$2^{R(\mathcal{M}_\delta)} = \sum_{\mathbf{x} \in \{0, 1\}^n} \hat{p}(\mathbf{x}).$$

Thus by Lemma 8, we have

$$2^{R(\mathcal{M}_\delta)} \leq 2^{2n\delta(\ell)} \sum_{\mathbf{x} \in \{0, 1\}^n} \hat{p}_\ell(\mathbf{x}) = 2^{2n\delta(\ell)} 2^{R(\mathcal{M}_\delta^\ell)}.$$

Therefore,

$$R(\mathcal{M}_\delta) \leq 2n\delta(\ell) + R(\mathcal{M}_\delta^\ell).$$

Observe that,

$$R(\mathcal{M}_\delta^\ell) \leq R(\mathcal{M}^\ell) \leq 2^{\ell-1} \log n,$$

where the last inequality holds whenever $\ell > 1$, see *e.g.*, [5]. \square

As before, we work out the above bounds for specific δ .

Corollary 9. If $\delta(\ell) = \frac{1}{\ell^c}$ with $c > 1$, then

$$R(\mathcal{M}_\delta) = O(n / \log^{c-1} n),$$

for $\ell = \log n - c \log \log n + o(1)$. For $\delta(\ell) = 2^{-c\ell}$, we have

$$R(\mathcal{M}_\delta) = O(n^{1/(c+1)} \log n),$$

for $\ell = \frac{1}{c+1} \log n + o(1)$. For $\delta(\ell) = 2^{-2^{c\ell}}$, we have

$$R(\mathcal{M}_\delta) = O(\log^{1+1/c} n),$$

, for $\ell = \frac{1}{c} \log \log n$. □

Comparing Corollary 2 and Corollary 8, the upper and lower bounds on the redundancy have asymptotically tight order when δ diminishes doubly exponentially. For polynomial δ the lower bound and upper bound orders differ by $\log n$ factors. However, for δ to be exponential, we have a polynomial gap between the lower and upper bound.

This suggests that either the lower or upper or both bound are too loose. For the lower bound, recall that the main contribution came from the fast mixing sources in M , while the other sources—the ones that are problematic to estimate, were summarily ignored.

Yet we will show in what follows that our lower bound given in Theorem 1 is actually tight. We need the following technical lemmas to refine our upper bound

Lemma 10 (Extended Azuma inequality). Let X_1, \dots, X_k, \dots be martingale differences with $|X_i| \leq 1$, τ is a stopping time (i.e. event $\{\tau = k\}$ only depends on $\sigma(X_1, \dots, X_k)$). If $\tau \leq n$, then we have

$$\Pr \left(\left| \sum_{i=1}^{\tau} X_i \right| \geq \gamma \sqrt{\tau} \right) \leq n e^{-\gamma^2/2}.$$

Proof Define $A_k = \{|X_1 + \dots + X_k| \geq \gamma \sqrt{k}\}$, $B_k = \{\tau = k\}$, let $C_k = A_k \cap B_k$. In fact, C_n is the event that we stop at n while it is a wrong time to stop. Note that $|X_i - X_{i-1}| < 1$, using Azuma inequality we have

$$\Pr[A_n] = \Pr \left\{ \left| \sum_{i=1}^n X_i \right| \geq \gamma \sqrt{n} \right\} \leq e^{-\gamma^2/2}.$$

Then

$$\Pr[\bigcup_{k=1}^n C_k] \leq \sum_{k=1}^n \Pr[A_k \cap B_k] \leq \sum_{k=1}^n \Pr[A_k] \leq n e^{-\gamma^2/2}.$$

and the Lemma follows. □

Lemma 11. For any $p \in \mathcal{M}_\delta$, we have

$$p \left(\sum_{\mathbf{w} \in \{0,1\}^\ell} |n_{\mathbf{w}1} - n_{\mathbf{w}} \tilde{p}_{\mathbf{w}}| \leq \log n \sqrt{n 2^\ell} \right) \geq 1 - \frac{2^\ell}{n^3},$$

for n large enough that $\log n \geq 6$, where $\tilde{p}_{\mathbf{w}}$ is defined in Section II.

Proof Define

$$1_i(\mathbf{s}) = \begin{cases} 1, & \text{the } i\text{-th appearance of } \mathbf{w} \text{ in } \mathbf{w} \prec \mathbf{s} \\ 0, & \text{otherwise.} \end{cases}$$

Let $W_i = \sum_{w \leq s} 1_i(\mathbf{s})p(1|\mathbf{s})$, and define

$$Y_i(\mathbf{w}) = \begin{cases} 1, & \text{the } i\text{-th appearance of } \mathbf{w} \text{ happens follows by one} \\ 0, & \text{otherwise.} \end{cases}$$

Let $Z_i = Y_i - W_i$, then by Lemma 2 in [21], Z_i are martingale differences and $|Z_i| < 1$.

Note that $n_{\mathbf{w}1} = \sum_i Y_i$ and $n_{\mathbf{w}}\tilde{p}_{\mathbf{w}} = \sum_i W_i$ and

$$|n_{\mathbf{w}1} - \tilde{p}_{\mathbf{w}}n_{\mathbf{w}}| = \sum_{\mathbf{s} \in S_{\mathbf{w}}(p)} |n_{\mathbf{s}1} - p(1|\mathbf{s})n_{\mathbf{s}}| = \left| \sum_{i=1}^{n_{\mathbf{w}}} Z_i \right|$$

Define $z_{\mathbf{w}} = |n_{\mathbf{w}1} - \tilde{p}_{\mathbf{w}}n_{\mathbf{w}}|$. Then using Lemma 10,

$$\begin{aligned} p\left(z_{\mathbf{w}} \geq \log n\sqrt{n_{\mathbf{w}}}\right) &\leq ne^{-\log n^2/2} \\ &= ne^{\log n - \log n/2} \\ &= \frac{n}{n^{\log n/2}} \\ &\leq \frac{1}{n^3}. \end{aligned}$$

Let $A_{\mathbf{w}} = \{z_{\mathbf{w}} < \log n\sqrt{n_{\mathbf{w}}}\}$. Then

$$\begin{aligned} p(\cup_{\mathbf{w}} A_{\mathbf{w}}^c) &= p(\cup_{\mathbf{w}} \{z_{\mathbf{w}} \geq \log n\sqrt{n_{\mathbf{w}}}\}) \\ &\leq \sum_{\mathbf{w} \in \{0,1\}^{\ell}} p(z_{\mathbf{w}} \geq \log n\sqrt{n_{\mathbf{w}}}) \\ &\leq \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \frac{1}{n^3} \\ &= \frac{2^{\ell}}{n^3}. \end{aligned}$$

Therefore,

$$p(\cap_{\mathbf{w}} A_{\mathbf{w}}) = 1 - p(\cup_{\mathbf{w}} A_{\mathbf{w}}^c) \geq 1 - \frac{2^{\ell}}{n^3}.$$

Note that event $\{\cap_{\mathbf{w}} A_{\mathbf{w}}\}$ implies event $\{\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}} < \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\}$, so

$$p\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}} < \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\right) \geq p(\cap_{\mathbf{w}} A_{\mathbf{w}}) \geq 1 - \frac{2^{\ell}}{n^3}.$$

Also,

$$p\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}} < \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\right) = p\left(\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}}\right)^2 < \left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\right)^2\right).$$

Using Cauchy-Schwartz inequality we have,

$$\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\right)^2 \leq \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log^2 n \sum_{\mathbf{w} \in \{0,1\}^{\ell}} n_{\mathbf{w}} = n2^{\ell} \log^2 n.$$

So,

$$p\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}} < \sum_{\mathbf{w} \in \{0,1\}^{\ell}} \log n\sqrt{n_{\mathbf{w}}}\right) = p\left(\sum_{\mathbf{w} \in \{0,1\}^{\ell}} z_{\mathbf{w}} < \sqrt{n2^{\ell}} \log n\right),$$

and the Lemma follows. \square

Consider a sample \mathbf{x} from $p \in \mathcal{M}_\delta$, past $x_{-\infty}^0$ and consider the empirical aggregated probabilities in (1) for $\mathbf{w} \in \{0, 1\}^\ell$. We now consider a memory ℓ Markov source that has its conditional probability of 1 given $\mathbf{w} \in \{0, 1\}^\ell$ equal to the empirically aggregated probabilities in (1), call the source \tilde{p}_ℓ . Note that \tilde{p}_ℓ need not be in the class \mathcal{M}_δ^ℓ or \mathcal{M}_δ , and while we do not explicitly say so in notation for ease of readability, \tilde{p}_ℓ depends on the sample \mathbf{x} .

For any \mathbf{x} and $p \in \mathcal{M}_\delta$, and for $\mathbf{w} \in \{0, 1\}^\ell$, let $z_\ell = \sum_{\mathbf{w}} z_{\mathbf{w}}$.

Lemma 12. For any $p \in \mathcal{M}_\delta$ and $\mathbf{x} \in \{0, 1\}^n$, we have

$$\tilde{p}_{\ell+1}(\mathbf{x}) \leq 2^{2n\delta^2(\ell)+2z_{\ell+1}\delta(\ell)} \tilde{p}_\ell(\mathbf{x}).$$

Moreover, we have

$$p\left(\left\{\mathbf{x} : z_{\ell+1} \leq \log n \sqrt{n2^{\ell+1}}\right\}\right) \geq 1 - \frac{2^\ell}{n^3}$$

Proof

Note that

$$\tilde{p}_{\ell+1}(\mathbf{x}) = \prod_{\mathbf{w} \in \{0, 1\}^\ell} \tilde{p}_{1\mathbf{w}}^{n_{1\mathbf{w}}} (1 - \tilde{p}_{1\mathbf{w}})^{n_{1\mathbf{w}} - n_{1\mathbf{w}1}} \tilde{p}_{0\mathbf{w}}^{n_{0\mathbf{w}}} (1 - \tilde{p}_{0\mathbf{w}})^{n_{0\mathbf{w}} - n_{0\mathbf{w}1}}$$

, and

$$\tilde{p}_\ell(\mathbf{x}) = \prod_{\mathbf{w} \in \{0, 1\}^\ell} \tilde{p}_{\mathbf{w}}^{n_{\mathbf{w}}} (1 - \tilde{p}_{\mathbf{w}})^{n_{\mathbf{w}} - n_{\mathbf{w}1}}$$

So we just need to show that

$$\begin{aligned} & \prod_{\mathbf{w} \in \{0, 1\}^\ell} \tilde{p}_{1\mathbf{w}}^{n_{1\mathbf{w}}} (1 - \tilde{p}_{1\mathbf{w}})^{n_{1\mathbf{w}} - n_{1\mathbf{w}1}} \tilde{p}_{0\mathbf{w}}^{n_{0\mathbf{w}}} (1 - \tilde{p}_{0\mathbf{w}})^{n_{0\mathbf{w}} - n_{0\mathbf{w}1}} \\ & \leq 2^{2n\delta^2(\ell)+2z_{\ell+1}\delta(\ell)} \prod_{\mathbf{w} \in \{0, 1\}^\ell} \tilde{p}_{\mathbf{w}}^{n_{\mathbf{w}}} (1 - \tilde{p}_{\mathbf{w}})^{n_{\mathbf{w}} - n_{\mathbf{w}1}} \end{aligned}$$

To see it, note

$$\tilde{p}_{\mathbf{w}} n_{\mathbf{w}} = \tilde{p}_{1\mathbf{w}} n_{1\mathbf{w}} + \tilde{p}_{0\mathbf{w}} n_{0\mathbf{w}}.$$

Let

$$\begin{aligned} \tilde{p}_{1\mathbf{w}} &= \tilde{p}_{\mathbf{w}} + \tilde{p}_{\mathbf{w}} \delta_1, \\ \tilde{p}_{0\mathbf{w}} &= \tilde{p}_{\mathbf{w}} + \tilde{p}_{\mathbf{w}} \delta_0, \end{aligned}$$

for some $|\delta_1| < l$ and $|\delta_0| < l$. Then

$$n_{0\mathbf{w}} \delta_0 + n_{1\mathbf{w}} \delta_1 = 0.$$

Let

$$\begin{aligned} n_{1\mathbf{w}1} &= \tilde{p}_{1\mathbf{w}} n_{1\mathbf{w}} + z'_1 = \tilde{p}_{\mathbf{w}} n_{1\mathbf{w}} + \tilde{p}_{\mathbf{w}} n_{1\mathbf{w}} \delta_1 + z'_1, \\ n_{0\mathbf{w}1} &= \tilde{p}_{0\mathbf{w}} n_{0\mathbf{w}} + z'_0 = \tilde{p}_{\mathbf{w}} n_{0\mathbf{w}} + \tilde{p}_{\mathbf{w}} n_{0\mathbf{w}} \delta_0 + z'_0, \end{aligned}$$

for some z'_0 and z'_1 . Also

$$\begin{aligned}
\log \tilde{p}_{1\mathbf{w}}^{n_{1\mathbf{w}1}} (1 - \tilde{p}_{1\mathbf{w}})^{n_{1\mathbf{w}} - n_{1\mathbf{w}1}} &= n_{1\mathbf{w}1} \log \tilde{p}_{\mathbf{w}} \\
&\quad + n_{1\mathbf{w}0} \log(1 - \tilde{p}_{1\mathbf{w}}) \\
&\leq n_{1\mathbf{w}1} (\log \tilde{p}_{1\mathbf{w}} + \delta_1) \\
&\quad + n_{1\mathbf{w}0} (\log(1 - \tilde{p}_{\mathbf{w}}) - \frac{\tilde{p}_{\mathbf{w}}}{1 - \tilde{p}_{\mathbf{w}}} \delta_1) \\
&= A_{1\mathbf{w}} + (\tilde{p}_{\mathbf{w}} + \frac{\tilde{p}_{\mathbf{w}}^2}{1 - \tilde{p}_{\mathbf{w}}}) n_{1\mathbf{w}} \delta_1^2 \\
&\quad + (\frac{\tilde{p}_{\mathbf{w}}}{1 - \tilde{p}_{\mathbf{w}}}) \delta_1 z'_1 \\
&\leq A_{1\mathbf{w}} + 2n_{1\mathbf{w}} \delta(\ell)^2 + 2\delta(\ell) |z'_1|.
\end{aligned}$$

where

$$A_{1\mathbf{w}} = n_{1\mathbf{w}1} \log \tilde{p}_{\mathbf{w}} + n_{1\mathbf{w}0} \log(1 - \tilde{p}_{\mathbf{w}}).$$

Similarly,

$$\begin{aligned}
\log \tilde{p}_{0\mathbf{w}}^{n_{0\mathbf{w}1}} (1 - \tilde{p}_{0\mathbf{w}})^{n_{0\mathbf{w}} - n_{0\mathbf{w}1}} &= n_{0\mathbf{w}1} \log \tilde{p}_{0\mathbf{w}} + n_{0\mathbf{w}0} (1 - \log \tilde{p}_{0\mathbf{w}}) \\
&\leq n_{0\mathbf{w}1} (\log \tilde{p}_{\mathbf{w}} + \delta_0) \\
&\quad + n_{0\mathbf{w}0} (\log(1 - \tilde{p}_{\mathbf{w}}) - \frac{\tilde{p}_{\mathbf{w}}}{1 - \tilde{p}_{\mathbf{w}}} \delta_0) \\
&= A_{0\mathbf{w}} + (\tilde{p}_{\mathbf{w}} + \frac{\tilde{p}_{\mathbf{w}}^2}{1 - \tilde{p}_{\mathbf{w}}}) n_{0\mathbf{w}} \delta_0^2 \\
&\quad + (\frac{\tilde{p}_{\mathbf{w}}}{1 - \tilde{p}_{\mathbf{w}}}) \delta_0 z'_0 \\
&\leq A_{0\mathbf{w}} + 2n_{0\mathbf{w}} \delta(\ell)^2 + 2\delta(\ell) |z'_0|,
\end{aligned}$$

and $A_{0\mathbf{w}} = n_{0\mathbf{w}1} \log \tilde{p}_{\mathbf{w}} + n_{0\mathbf{w}0} \log(1 - \tilde{p}_{\mathbf{w}})$. Summing over all \mathbf{w} , we have

$$\begin{aligned}
&\sum_{\mathbf{w}} \log \tilde{p}_{1\mathbf{w}}^{n_{1\mathbf{w}1}} (1 - \tilde{p}_{1\mathbf{w}})^{n_{1\mathbf{w}} - n_{1\mathbf{w}1}} \tilde{p}_{0\mathbf{w}}^{n_{0\mathbf{w}1}} (1 - \tilde{p}_{0\mathbf{w}})^{n_{0\mathbf{w}} - n_{0\mathbf{w}1}} \\
&\leq \sum_{\mathbf{w}} (A_{1\mathbf{w}} + A_{0\mathbf{w}}) + 2n\delta(\ell)^2 + 2\delta(\ell) z_{\ell+1} \\
&= \log \tilde{p}_{\mathbf{w}}^{n_{\mathbf{w}1}} (1 - \tilde{p}_{\mathbf{w}})^{n_{\mathbf{w}} - n_{\mathbf{w}1}} + 2n\delta(\ell)^2 + 2\delta(\ell) z_{\ell+1},
\end{aligned}$$

where we use the fact that $z_{\ell+1} = \sum_{\mathbf{w}} |z'_0| + |z'_1|$. Also using Lemma 11 one can see that

$$p \left(\left\{ \mathbf{x} : z_{\ell+1} \leq \log n \sqrt{n 2^{\ell+1}} \right\} \right) \geq 1 - \frac{2^\ell}{n^3}$$

□

Lemma 13. For any $p \in \mathcal{M}_\delta$, we have

$$p(\{\mathbf{x} : p(\mathbf{x}) \leq 2^{r_\ell} \tilde{p}_\ell(\mathbf{x})\}) \geq 1 - \frac{\sum_{k=\ell}^{2\ell} 2^k}{n^3},$$

where

$$r_\ell = n\delta(2\ell) + \sum_{k=\ell}^{2\ell} 2n\delta^2(k) + 2\log n \sqrt{n 2^{k+1}} \delta(k).$$

Proof Using Lemma 8 we have

$$p(\mathbf{x}) \leq \tilde{p}_{2\ell}(\mathbf{x}) 2^{2n\delta(2\ell)}. \quad (6)$$

Also,

$$\begin{aligned}\tilde{p}_{\ell+1}(\mathbf{x}) &= \prod_{\mathbf{w} \in \{0,1\}^\ell} \tilde{p}_{1\mathbf{w}}^{n_{1\mathbf{w}1}} (1 - \tilde{p}_{1\mathbf{w}})^{n_{1\mathbf{w}} - n_{1\mathbf{w}1}} \tilde{p}_{0\mathbf{w}}^{n_{0\mathbf{w}1}} (1 - \tilde{p}_{0\mathbf{w}})^{n_{0\mathbf{w}} - n_{0\mathbf{w}1}} \\ &\leq 2^{2n\delta^2(\ell) + 2z_\ell\delta(\ell)} \prod_{\mathbf{w} \in \{0,1\}^\ell} \tilde{p}_{\mathbf{w}}^{n_{\mathbf{w}1}} (1 - \tilde{p}_{\mathbf{w}})^{n_{\mathbf{w}} - n_{\mathbf{w}1}}.\end{aligned}$$

Similarly,

$$\begin{aligned}\tilde{p}_{2\ell+1}(\mathbf{x}) &\leq \sum_{k=\ell}^{2\ell} 2^{2n\delta^2(k) + 2z_{k+1}\delta(k)} \prod_{\mathbf{w} \in \{0,1\}^k} \tilde{p}_{\mathbf{w}}^{n_{\mathbf{w}1}} (1 - \tilde{p}_{\mathbf{w}})^{n_{\mathbf{w}} - n_{\mathbf{w}1}} \\ &= 2^{\sum_{k=\ell}^{2\ell} 2n\delta^2(k) + 2z_{k+1}\delta(k)} \tilde{p}_\ell(\mathbf{x}).\end{aligned}\tag{7}$$

and using equation (6) and equation (7), we have

$$p(\mathbf{x}) \leq \tilde{p}_\ell(\mathbf{x}) 2^{n\delta(2\ell) + \sum_{k=\ell}^{2\ell} 2n\delta^2(k) + 2z_{k+1}\delta(k)}.$$

but note that for all k

$$p\left(\left\{\mathbf{x} : z_{k+1} \leq \log n \sqrt{n2^{k+1}}\right\}\right) \geq 1 - \frac{2^k}{n^3},$$

Using union bound

$$p\left(\left\{\mathbf{x} : \sum_{k=l}^{2l} z_{k+1} \leq \sum_{k=l}^{2l} \log n \sqrt{n2^{k+1}}\right\}\right) \geq 1 - \sum_{k=l}^{2l} \frac{2^k}{n^3},$$

and the lemma follows. \square

Theorem 14 (Improved Upper Bound). Redundancy of \mathcal{M}_δ is upper bounded by

$$\tilde{R}(\mathcal{M}_\delta) \leq 2^{\ell-1} \log n + n\delta(2\ell) + \sum_{k=\ell}^{2\ell} (n\delta^2(k) + \log n \sqrt{n2^k} \delta(k)) + (2^{2\ell+1} - 2^\ell) \frac{n}{n^3}$$

for any integer $\ell \in \mathbb{N}$.

Proof Let $\mathcal{T}_p = \{\mathbf{x} : p(\mathbf{x}) \leq 2^{r_\ell} \tilde{p}_\ell(\mathbf{x})\}$ be the set of good sequences and $\mathcal{T}_p^c = \{\mathbf{x} : \mathbf{x} \notin \mathcal{T}_p\}$. Let $c(\mathbf{x})$ be the best code for memory ℓ sources. Let $|c(\mathbf{x})|$ denote the length of $c(\mathbf{x})$. Let $q(\mathbf{x}) = \frac{2^{-c(\mathbf{x})} + 2^{-n}}{2}$. We can choose $c(\mathbf{x})$ tight enough so that $\sum 2^{-c(\mathbf{x})} = 1$. Then

$$\begin{aligned}\tilde{R}(\mathcal{M}_\delta) &= \max_{p \in \mathcal{M}_\delta} \sum_{\mathbf{x}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} \\ &= \max_{p \in \mathcal{M}_\delta} \sum_{\mathbf{x} \in \mathcal{T}_p} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} + \sum_{\mathbf{x} \in \mathcal{T}_p^c} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} \\ &\leq \max_{p \in \mathcal{M}_\delta} \sum_{\mathbf{x} \in \mathcal{T}_p} p(\mathbf{x}) \max_{\mathbf{x} \in \{0,1\}^n} \log \frac{p(\mathbf{x})}{q(\mathbf{x})} + \sum_{\mathbf{x} \in \mathcal{T}_p^c} p(\mathbf{x}) \max_{\mathbf{x} \in \{0,1\}^n} \log \frac{p(\mathbf{x})}{q(\mathbf{x})} \\ &\leq \max_{p \in \mathcal{M}_\delta} \sum_{\mathbf{x} \in \mathcal{T}_p} p(\mathbf{x}) \max_{\mathbf{x} \in \{0,1\}^n} \log \frac{p_\ell 2^{r_\ell}}{q(\mathbf{x})} + \sum_{\mathbf{x} \in \mathcal{T}_p^c} p(\mathbf{x}).n\end{aligned}$$

$$\begin{aligned}
&\leq \max_{p \in \mathcal{M}_\delta} \max_{\mathbf{x} \in \{0,1\}^n} \log \frac{p_\ell 2^{r_\ell}}{q(\mathbf{x})} + n \frac{\sum_{k=\ell}^{2\ell} 2^k}{n^3} \\
&= \max_{p \in \mathcal{M}_\delta} \max_{\mathbf{x} \in \{0,1\}^n} [\log p_\ell(\mathbf{x}) + c(\mathbf{x}) + 1] + r_\ell + n \frac{\sum_{k=\ell}^{2\ell} 2^k}{n^3} \\
&= 2^{\ell-1} \log n + r_\ell + (2^{2\ell+1} - 2^\ell) \frac{n}{n^3}.
\end{aligned}$$

Where $r_\ell = n\delta(2\ell) + \sum_{k=\ell}^{2\ell} n\delta^2(k) + \log n\sqrt{n2^\ell}\delta(k)$. Note that first term in the last equation follows since the worst case redundancy of Markov sources with memory ℓ and is bounded by $2^{\ell-1} \log n$. So

$$\begin{aligned}
\tilde{R}(\mathcal{M}_\delta) &\leq 2^{\ell-1} \log n + n\delta(2\ell) \\
&+ \sum_{k=\ell}^{2\ell} (n\delta^2(k) + \log n\sqrt{n2^k}\delta(k)) + (2^{2\ell+1} - 2^\ell) \frac{n}{n^3}.
\end{aligned}$$

□

Corollary 15. For $\delta(\ell) = 2^{-c\ell}$, we have

$$\tilde{R}(\mathcal{M}_\delta) = O(n^{1/(2c+1)} \log n).$$

□

Proof Note that,

$$\begin{aligned}
\tilde{R}(\mathcal{M}_\delta) &\leq 2^{\ell-1} \log n + n\delta(2\ell) \\
&+ \sum_{k=\ell}^{2\ell} (n\delta^2(k) + \log n\sqrt{n2^k}\delta(k)) + (2^{2\ell+1} - 2^\ell) \frac{n}{n^3},
\end{aligned}$$

let $\delta(k) = 2^{-ck}$ then

$$\begin{aligned}
\sum_{k=\ell}^{2\ell} \delta^2(k) &= \sum_{k=\ell}^{2\ell} 2^{-2ck} = 2^{-2cl} + 2^{-2c(l+1)} + \dots + 2^{-2c(2\ell)} \\
&= 2^{-2cl} (1 + 2^{-2c} + \dots + 2^{-2cl}) \\
&= 2^{-2cl} \frac{1 - 2^{-2c(\ell+1)}}{1 - 2^{-2c}} \\
&= \frac{2^{-2cl} - 2^{-2c(2\ell+1)}}{1 - 2^{-2c}},
\end{aligned}$$

and

$$\begin{aligned}
\sum_{k=\ell}^{2\ell} 2^{k/2} \delta(k) &= \sum_{k=\ell}^{2\ell} 2^{(-c+1/2)k} \\
&= 2^{(-c+1/2)\ell} + 2^{(-c+1/2)(\ell+1)} + \dots + 2^{(-c+1/2)2\ell} \\
&= 2^{(-c+1/2)\ell} (1 + 2^{-c} + \dots + 2^{-2cl}) \\
&= 2^{(-c+1/2)\ell} \frac{1 - 2^{-c(\ell+1)}}{1 - 2^{-c}} \\
&= \frac{2^{(-c+1/2)\ell} - 2^{(-2c+1/2)\ell-c}}{1 - 2^{-c}}.
\end{aligned}$$

Let $\ell = c' \log n$, then

$$\begin{aligned}\tilde{R}(\mathcal{M}_\delta) &\leq \frac{n^{c'}}{2} \log n + \frac{n}{n^{-2cc'}} \\ &+ n \left(\frac{n^{-2cc'} - \frac{n^{-4cc'}}{2^{-2c}}}{1 - 2^{-2c}} \right) \\ &+ \log n \sqrt{n} \left(\frac{n^{c'(-c+1/2)} - \frac{n^{(-2c+1/2)c'}}{2^c}}{1 - 2^{-c}} \right) \\ &+ (2n^{2c'} - n^{c'}) \frac{n}{n^3}.\end{aligned}$$

Let $c' = \frac{1}{2c+1}$

$$\begin{aligned}\tilde{R}(\mathcal{M}_\delta) &\leq \frac{1}{2} n^{\frac{1}{2c+1}} \log n + n^{\frac{1}{2c+1}} \\ &+ \frac{n^{\frac{1}{2c+1}} - \frac{n^{\frac{-2c}{2c+1}}}{2^{-2c}}}{1 - 2^{-2c}} \\ &+ \log n \left(\frac{n^{\frac{1}{2c+1}} - \frac{1}{2^c} n^{\frac{-c+1}{2c+1}}}{1 - 2^{-c}} \right) \\ &+ (2n^{\frac{2c}{2c+1}} - n^{\frac{1}{2c+1}}) \frac{n}{n^3} \\ &= \mathcal{O}(n^{\frac{1}{2c+1}} \log n).\end{aligned}$$

□

5 Conclusion

We proved matching (in order) upper and lower bounds on the redundancy of universally compressing length- n strings from \mathcal{M}_δ . In the process, we examined which sources contribute predominantly to the redundancy—discovering that fast mixing sources, or where estimation is uncomplicated by mixing biases, are the biggest contributors. This reveals an interesting dichotomy in the Markov setup—the sources that make estimation complicated are different from the sources that complicate compression.

In our future work, we hope to use these results to better understand what sort of Bayesian priors work in the Markov setups for various tasks, resampling techniques for Markov sampling—the initial results of which are available in [21], as well as extend our understanding of data-derived estimation [22] into the Markov regime.

Acknowledgment

The authors were supported by NSF grants CCF 1619452 “Slow mixing Markov processes” and by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

References

- [1] Meysam Asadi, Ramezan Paravi Torghabeh, and Narayana P Santhanam. Stationary and transition probabilities in slow mixing, long memory markov processes. *IEEE Transactions on Information Theory*, 60(9):5682–5701, 2014.

- [2] Imre Csiszár and Zsolt Talata. Context tree estimation for not necessarily finite memory processes, via bic and mdl. *IEEE Transactions on Information theory*, 52(3):1007–1016, 2006.
- [3] L.D. Davisson. Universal noiseless coding. *IEEE Transactions on Information Theory*, 19(6):783—795, November 1973.
- [4] R.E. Krichevsky and V.K. Trofimov. The performance of universal coding. *IEEE Transactions on Information Theory*, 27(2):199—207, March 1981.
- [5] Y.M. Shtarkov, T.J. Tjalkens, and F.M.J. Willems. Multialphabet universal coding of memoryless sources. *Problems of Information Transmission*, 31(2):114—127, 1995.
- [6] Q. Xie and A. Barron. Minimax redundancy for the class of memoryless sources. *IEEE Transactions on Information Theory*, 43(2):647—657, March 1997.
- [7] W. Szpankowski. On asymptotics of certain recurrences arising in universal coding. *Problems of Information Transmission*, 34(2):142—146, 1998.
- [8] Viktor Kupriyanovich Trofimov. Redundancy of universal coding of arbitrary markov sources. *Problemy Peredachi Informatsii*, 10(4):16—24, 1974.
- [9] L Davisson. Minimax noiseless universal coding for markov sources. *IEEE Transactions on Information Theory*, 29(2):211—215, 1983.
- [10] Kevin Atteson. The asymptotic redundancy of bayes rules for markov chains. *IEEE Transactions on Information Theory*, 45(6):2104—2109, 1999.
- [11] J. Rissanen. Fisher information and stochastic complexity. *IEEE Transactions on Information Theory*, 42(1):40—47, January 1996.
- [12] Philippe Jacquet and Wojciech Szpankowski. Markov types and minimax redundancy for markov sources. *IEEE Transactions on Information Theory*, 50(7):1393—1402, 2004.
- [13] Frans MJ Willems, Yuri M Shtarkov, and Tjalling J Tjalkens. The context-tree weighting method: basic properties. *IEEE Transactions on Information Theory*, 41(3):653—664, 1995.
- [14] Jorma Rissanen. A universal data compression system. *IEEE Transactions on information theory*, 29(5):656—664, 1983.
- [15] Frans MJ Willems. The context-tree weighting method: Extensions. *IEEE Transactions on Information Theory*, 44(2):792—798, 1998.
- [16] Moein Falahatgar, Alon Orlitsky, Venkatadheeraj Pichapati, and Ananda Theertha Suresh. Learning markov distributions: Does estimation trump compression? In *Information Theory (ISIT), 2016 IEEE International Symposium on*, pages 2689—2693. IEEE, 2016.
- [17] Anand K Dhulipala and Alon Orlitsky. Universal compression of markov and related sources over arbitrary alphabets. *IEEE Transactions on Information Theory*, 52(9):4182—4190, 2006.
- [18] Imre Csiszár, Paul C Shields, et al. Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory*, 1(4):417—528, 2004.
- [19] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley and sons., 1991.
- [20] Yurii Mikhailovich Shtarkov. Universal sequential coding of single messages. *Problemy Peredachi Informatsii*, 23(3):3—17, 1987.

- [21] K. Oshiro, C. Wu, and N. Santhanam. Jackknife estimation for markov processes with no mixing constraints. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pages 3020–3024. IEEE, 2017.
- [22] N. Santhanam and V. Anantharam. Agnostic insurance of model classes. *Journal of Machine Learning Research*, November 2015. Full version available from arXiv doc id: 1212:3866.