

An Implementable Scheme for Universal Lossy Compression of Discrete Markov Sources

Shirin Jalali,^{*} Andrea Montanari^{*} and Tsachy Weissman^{*†},

^{*}Department of Electrical Engineering, Stanford University, Stanford, CA 94305,

[†] Department of Electrical Engineering, Technion, Haifa 32000, Israel

{shjalali, montanar, tsachy}@stanford.edu

Abstract

We present a new lossy compressor for discrete sources. For coding a source sequence x^n , the encoder starts by assigning a certain cost to each reconstruction sequence. It then finds the reconstruction that minimizes this cost and describes it losslessly to the decoder via a universal lossless compressor. The cost of a sequence is given by a linear combination of its empirical probabilities of some order $k + 1$ and its distortion relative to the source sequence. The linear structure of the cost in the empirical count matrix allows the encoder to employ a Viterbi-like algorithm for obtaining the minimizing reconstruction sequence simply. We identify a choice of coefficients for the linear combination in the cost function which ensures that the algorithm universally achieves the optimum rate-distortion performance of any Markov source in the limit of large n , provided k is increased as $o(\log n)$.

I. INTRODUCTION

Let $\mathbf{X} = \{X_i : i \geq 1\}$ represent a discrete-valued stationary ergodic process with unknown statistics, and consider the problem of compressing \mathbf{X} at rate R such that the incurred distortion is minimized. Let \mathcal{X} and $\hat{\mathcal{X}}$ denote finite source and reconstruction alphabets respectively. The performance of the described coding scheme is measured by its average expected distortion between source and reconstruction blocks, i.e.

$$D = \mathbb{E} d_n(X^n, \hat{X}^n) \triangleq \frac{1}{n} \sum_{i=1}^n \mathbb{E} d(X_i, \hat{X}_i), \quad (1)$$

where $d : \mathcal{X} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}^+$ is a single-letter distortion measure. For any $R \geq 0$, the minimum achievable distortion (cf. [4] for exact definition of achievability) is characterized as [1], [2], [3]

$$D(\mathbf{X}, R) = \lim_{n \rightarrow \infty} \min_{p(\hat{X}^n | X^n) : I(X^n; \hat{X}^n) \leq R} \mathbb{E} d_n(X^n, \hat{X}^n). \quad (2)$$

A sequence of codes at rate R is called universal if for every stationary ergodic source \mathbf{X} its asymptotic performance converges to $D(\mathbf{X}, R)$, i.e.,

$$\limsup_{n \rightarrow \infty} \mathbb{E} d_n(X^n, \hat{X}^n) \leq D(\mathbf{X}, R). \quad (3)$$

For lossless compression where the source is to be recovered without any errors, there already exist well-known implementable universal schemes such as Lempel-Ziv coding [5] or arithmetic coding [6]. In contrast to the situation of lossless compression, for $D > 0$, there are no well-known practical schemes that universally achieve the rate-distortion curve. In recent years, there has been progress towards designing universal lossy compressor especially in trying to tune some of the existing universal lossless coders to work in the lossy case as well [7], [8], [9]. All of these algorithms are either provably suboptimal, or optimal but with exponential complexity.

Another approach for lossy compression, which is very well-studied in the literature and even implemented in JPEG 2000 image compression standard, is Trellis coded quantization, i.e. Trellis structured code plus Viterbi encoding (c.f. [10], [11] and references therein). This method is in general suboptimal for coding sources that have memory [11]. In [12], an algorithm for fixed-slope Trellis source coding is

proposed, and is shown to be able to get arbitrary close to the rate-distortion curve for continuous-valued stationary ergodic sources. The proposed method is efficient in low rate region.

In a recent work [13], a new implementable algorithm for lossy compression of discrete-valued stationary ergodic sources was proposed. Instead of fixing rate (or distortion) and minimizing distortion (or rate), the new algorithm fixes Lagrangian coefficient α , and minimizes $R + \alpha D$. This is done by assigning energy $\mathcal{E}(y^n)$ representing $R + \alpha D$ to each possible reconstruction sequence and finding the sequence that minimizes the cost by simulated annealing. The algorithm starts by letting $y^n = x^n$, and at each iteration chooses an index $i \in \{1, \dots, n\}$ uniformly at random, and probabilistically changes y_i to some $y \in \hat{\mathcal{X}}$ such that there is a positive probability (which goes to zero as the number of iterations increases) that the resulting sequence has higher energy than the original sequence. Allowing the energy to increase especially at initial steps prevents the algorithm from being entrapped in a local minimum. It was shown that using a universal lossless compressor to describe the reconstruction sequence resulting from this process to the decoder results in a scheme which is universal in the limit of many iterations and large block length. The drawback of the proposed scheme is that although its computational complexity per iteration is independent of the block length n and linear in a parameter $k_n = o(\log n)$, there is no useful bound on the number of iterations required for convergence. In this paper, inspired by the previous method, we propose yet another approach for lossy compression of discrete Markov sources which universally achieves optimum rate-distortion performance for any discrete Markov source. We start by assigning the same cost that was defined for each possible reconstruction sequence in [13]. The cost of each sequence is a linear combination of two terms: its empirical conditional entropy and its distance to the source sequence to be coded. We show that there exists proper linear approximation of the first term such that minimizing the linearized cost results in the same performance as minimizing the original cost. But the advantage is that minimizing the modified cost can be done via Viterbi algorithm in lieu of simulated annealing which was used for minimizing the original cost.

The organization of the paper is as follows. In Section II, we set up the notation, and define the count matrix and empirical conditional entropy of a sequence. Section III describes a new coding scheme for fixed-slope lossy compression which universally achieves the rate-distortion curve for any discrete Markov source and IV describes how to compute the coefficients required by the algorithm outlined in the previous section. Section V explains how Viterbi algorithm can be used for implementing the coding scheme described in Section III. Section VI presents some simulations results, and finally, Section VII concludes the paper with a discussion of some future directions.

Proofs that are not presented in the paper will appear in the full version.

II. NOTATIONS AND REQUIRED DEFINITIONS

Let \mathcal{X} and $\hat{\mathcal{X}}$ denote the source and reconstruction alphabets respectively. Let matrix $\mathbf{m}(y^n) \in \mathbb{R}^{|\hat{\mathcal{X}}| \times \mathbb{R}^{|\hat{\mathcal{X}}|^k}}$ represent $(k+1)^{\text{th}}$ order empirical count of y^n defined as

$$m_{\beta, \mathbf{b}}(y^n) = \frac{1}{n} |\{1 \leq i \leq n : y_{i-k}^{i-1} = \mathbf{b}, y_i = \beta\}|. \quad (4)$$

In (4), and throughout we assume a cyclic convention whereby $y_i \triangleq y_{n+i}$ for $i \leq 0$. Let $H_k(y^n)$ denote the conditional empirical entropy of order k induced by y^n , i.e.

$$H_k(y^n) = H(Y_{k+1}|Y^k), \quad (5)$$

where Y^{k+1} on the right hand side of (5) is distributed according to

$$P(Y^{k+1} = [\mathbf{b}, \beta]) = m_{\beta, \mathbf{b}}(y^n), \quad (6)$$

where $\beta \in \hat{\mathcal{X}}$, and $\mathbf{b} \in \hat{\mathcal{X}}^k$, and $[\mathbf{b}, \beta]$ represents the vector made by concatenation of \mathbf{b} and β . We will use the same notation throughout the paper, namely, $\beta, \beta', \dots \in \hat{\mathcal{X}}$, and $\mathbf{b}, \mathbf{b}', \dots \in \hat{\mathcal{X}}^k$. The conditional empirical entropy in (5) can be expressed as a function of $\mathbf{m}(y^n)$ as follows

$$H_k(y^n) = H_k(\mathbf{m}(y^n)) := \frac{1}{n} \sum_{\beta, \mathbf{b}} \mathcal{H}(\mathbf{m}_{\cdot, \mathbf{b}}(y^n)) \mathbf{1}^T \mathbf{m}_{\cdot, \mathbf{b}}(y^n), \quad (7)$$

where $\mathbf{1}$ and $\mathbf{m}_{\cdot, \mathbf{b}}(y^n)$ denote the all-ones column vector of length $|\hat{\mathcal{X}}|$, and the column in $\mathbf{m}(y^n)$ corresponding to \mathbf{b} respectively. For a vector $\mathbf{v} = (v_1, \dots, v_\ell)^T$ with non-negative components, we let $\mathcal{H}(\mathbf{v})$ denote the entropy of the random variable whose probability mass function (pmf) is proportional to \mathbf{v} . Formally,

$$\mathcal{H}(\mathbf{v}) = \begin{cases} \sum_{i=1}^{\ell} \frac{v_i}{\|\mathbf{v}\|_1} \log \frac{\|\mathbf{v}\|_1}{v_i} & \text{if } \mathbf{v} \neq (0, \dots, 0)^T \\ 0 & \text{if } \mathbf{v} = (0, \dots, 0)^T. \end{cases} \quad (8)$$

III. LINEARIZED COST FUNCTION

Consider the following scheme for lossy source coding at fixed slope $\alpha > 0$. For each source sequence x^n let the reconstruction block \hat{x}^n be

$$\hat{x}^n = \arg \min_{y^n \in \hat{\mathcal{X}}^n} [H_k(y^n) + \alpha d_n(x^n, y^n)]. \quad (9)$$

The encoder, after computing \hat{x}^n , losslessly conveys it to the decoder using LZ compression. Let k grow slowly enough with n so that

$$\limsup_{n \rightarrow \infty} \max_{y^n} \left[\frac{1}{n} \ell_{\text{LZ}}(y^n) - H_k(y^n) \right] \leq 0, \quad (10)$$

where $\ell_{\text{LZ}}(y^n)$ denotes the length of the LZ representation of y^n . Note that Ziv's inequality guarantees that if $k = k_n = o(\log n)$ then (10) holds.

Theorem 1: [13] Let \mathbf{X} be a stationary and ergodic source, let $R(\mathbf{X}, D)$ denote its rate distortion function, and let \hat{X}^n denote the reconstruction using the above scheme for coding X^n . Then

$$\mathbb{E} \left[\frac{1}{n} \ell_{\text{LZ}}(\hat{X}^n) + \alpha d_n(X^n, \hat{X}^n) \right] \xrightarrow{n \rightarrow \infty} \min_{D \geq 0} [R(\mathbf{X}, D) + \alpha D]. \quad (11)$$

In other words, conveying the reconstruction sequence to the decoder via universal lossless compression (selection of LZ algorithm here is for concreteness, but other universal lossless methods can be used as well) achieves optimum fixed-slope rate-distortion performance universally.

As proposed in [13], the exhaustive search required by this algorithm can be tackled through simulated annealing Gibbs sampling. Here assuming the source is a discrete Markov source, we propose another method for finding a sequence achieving the minimum in (9). The advantage of the new method is that its computational complexity is linear in n for fixed k .

Before describing the new scheme, consider the problems (P1) and (P2) described below.

$$(P1) : \quad \min_{y^n} [H_k(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n)], \quad (12)$$

and

$$(P2) : \quad \min_{y^n} \left[\sum_{\beta} \sum_{\mathbf{b}} \lambda_{\beta, \mathbf{b}} m_{\beta, \mathbf{b}}(y^n) + \alpha d_n(x^n, y^n) \right]. \quad (13)$$

Comparing (P1) with (9) reveals that it is the optimization required by the exhaustive search coding scheme described before. The question is whether it is possible to choose a set of coefficients $\{\lambda_{\beta,\mathbf{b}}\}_{\beta,\mathbf{b}}$, $\beta \in \hat{\mathcal{X}}$ and $\mathbf{b} \in \hat{\mathcal{X}}^k$, such that (P1) and (P2) have the same set of minimizers or at least, the set of minimizers of (P2) is a subset of minimizers of (P1). If the answer to this question is affirmative, then instead of solving (P1) one can solve (P2), which, as we describe in Section V, can be done simply via the Viterbi algorithm.

Let \mathcal{S}_1 and \mathcal{S}_2 denote the set of minimizers of (P1) and (P2). Consider some $z^n \in \mathcal{S}_1$, and let $\mathbf{m}_n^* = \mathbf{m}(z^n)$. Since $H(\mathbf{m})$ is concave in \mathbf{m} ,¹ for any empirical count matrix \mathbf{m} , we have

$$H(\mathbf{m}) \leq H(\mathbf{m}_n^*) + \sum_{\beta,\mathbf{b}} \frac{\partial}{\partial m_{\beta,\mathbf{b}}} H(\mathbf{m}) \Big|_{\mathbf{m}_n^*} (m_{\beta,\mathbf{b}} - m_{\beta,\mathbf{b}}^*) \quad (14)$$

$$\triangleq \hat{H}(\mathbf{m}). \quad (15)$$

Now assume that in (P2), the coefficients are chosen as follows

$$\lambda_{\beta,\mathbf{b}} = \frac{\partial}{\partial m_{\beta,\mathbf{b}}} H(\mathbf{m}) \Big|_{\mathbf{m}_n^*}. \quad (16)$$

Lemma 1: (P1) and (P2) have the same minimum value, if the coefficients are chosen according to (16). Moreover, if all the sequences in \mathcal{S}_1 have the same type, then $\mathcal{S}_1 = \mathcal{S}_2$.

Proof: For any $y^n \in \hat{\mathcal{X}}^n$,

$$H(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n) \leq \hat{H}(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n). \quad (17)$$

Therefore,

$$\min_{y^n} [H(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n)] \leq \min_{y^n} [\hat{H}(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n)] \quad (18)$$

$$\leq \hat{H}(\mathbf{m}(z^n)) + \alpha d_n(x^n, z^n) \quad (19)$$

$$= \min_{y^n} [H(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n)]. \quad (20)$$

This shows that (P1) and (P2) have the same minimum values. For any sequence y^n with $\mathbf{m}(y^n) \neq \mathbf{m}_n^*$, by strict concavity of $H(\mathbf{m})$,

$$\hat{H}(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n) > H(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n), \quad (21)$$

$$\geq \min_{y^n} [H(\mathbf{m}(y^n)) + \alpha d_n(x^n, y^n)]. \quad (22)$$

As a result all the sequences in \mathcal{S}_2 should have the empirical count matrix equal to \mathbf{m}_n^* . Since for these sequences $H(\mathbf{m}) = \hat{H}(\mathbf{m})$, we also conclude that $\mathcal{S}_2 \subset \mathcal{S}_1$. If there is a unique minimizing type \mathbf{m}_n^* , then $\mathcal{S}_1 = \mathcal{S}_2$. ■

This shows that if we knew the optimal type \mathbf{m}_n^* , then we could compute the optimal coefficients via (16), and solve (P2) instead of (P1). The problem is that \mathbf{m}_n^* is not known to the encoder (since knowledge of \mathbf{m}_n^* requires solving (P1) which is the problem we are trying to avoid). In the next section, we describe a method for approximating \mathbf{m}_n^* , and hence the coefficients $\{\lambda_{\beta,\mathbf{b}}\}$.

¹As proved in Appendix B.

IV. HOW TO CHOOSE THE COEFFICIENTS?

For a given stationary ergodic source \mathbf{X} , and for any given count matrix \mathbf{m} define $D(\mathbf{m})$ to be the minimum average expected distortion among all processes \mathbf{Y} that are jointly stationary ergodic with \mathbf{X} and their $(k+1)^{\text{th}}$ order stationary distribution is according to \mathbf{m} .² $D(\mathbf{m})$ can equivalently be defined as

$$D(\mathbf{m}) = \lim_{k_1 \rightarrow \infty} \min_{p(x^{k_1}, y^{k_1}) \in \mathcal{M}^{(k_1)}} \mathbb{E}_p d(x^{k_1}, y^{k_1}), \quad (23)$$

where $\mathcal{M}^{(k_1)}$ is the set of all jointly stationary distributions $p(x^{k_1}, y^{k_1})$ of (X^{k_1}, Y^{k_1}) with marginal distributions with respect to x coinciding with the k_1^{th} order distribution of \mathbf{X} process, and with marginal distributions with respect to y coinciding with \mathbf{m} , i.e., having the $(k+1)^{\text{th}}$ order marginal distribution described by \mathbf{m} .

Lemma 2: If the source is ℓ^{th} order Markov, then

$$D(\mathbf{m}) = \min_{p(x^{k_1}, y^{k_1}) \in \mathcal{M}^{(k_1)}} \mathbb{E}_p d(x^{k_1}, y^{k_1}), \quad (24)$$

where $k_1 = \max(\ell, k+1)$.

Proof: [outline] Using the technique described in Appendix A, for any legitimate given joint distribution $p(x^{k_1}, y^{k_1})$ with the marginal distribution with respect to x coinciding with the source distribution and with marginal distribution with respect to y coinciding with some given distribution \mathbf{m} , it is possible to construct a process which is jointly stationary and ergodic with our source process and also has the $(k+1)^{\text{th}}$ order joint distribution as $p(x^{k_1}, y^{k_1})$. Using this gives us an achievable distortion, i.e., an upper bound on $D(\mathbf{m})$. On the other hand, the limit given in (23) is approaching $D(\mathbf{m})$ from below. Combining the upper and lower bounds yields the desired equality. ■

Since by assumption the encoder does not know ℓ , therefore it can not compute $\max(\ell, k+1)$. But letting $k_1 = k+1$, where $k = o(\log n)$, for any fixed order ℓ , k_1 will eventually for n large enough, exceed ℓ , and hence be equal to $\max(\ell, k+1)$. Having this observation in mind, consider the following optimization problem,

$$\begin{aligned} \min \quad & H(\mathbf{m}) + \alpha D(\mathbf{m}) \\ \text{s.t.} \quad & \mathbf{m} \in \mathcal{M}^{(k_1)}. \end{aligned} \quad (25)$$

By Lemma 2, an equivalent representation of (25) is

$$\begin{aligned} \min \quad & H(\mathbf{m}) + \alpha \sum_{\beta, \beta', \mathbf{b}, \mathbf{b}'} d_{k_1}(\beta' \mathbf{b}', \beta \mathbf{b}) p_x(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') \\ \text{s.t.} \quad & m_{\beta, \mathbf{b}} = \sum_{\beta' \mathbf{b}'} p_x(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}'), \quad \forall \beta, \mathbf{b}, \\ & 0 \leq q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') \leq 1, \quad \forall \beta, \beta', \mathbf{b}, \mathbf{b}', \\ & \sum_{\beta, \mathbf{b}} q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') = 1, \quad \forall \beta', \mathbf{b}', \\ & \sum_{\beta, \beta'} p_x(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') = \sum_{\beta, \beta'} p_x(\mathbf{b}' \beta') q_{y|x}(\mathbf{b} \beta | \mathbf{b}' \beta') \quad \forall \mathbf{b}, \mathbf{b}'. \end{aligned} \quad (26)$$

The last constraint in (26) is the *stationarity condition* defined in (A-1), and ensures that the joint distribution defined by $p_x(\beta \mathbf{b}) q_{y|x}(\beta' \mathbf{b}' | \beta, \mathbf{b}')$ over (x^{k+1}, y^{k+1}) corresponds to $(k+1)^{\text{th}}$ order marginal distribution of some jointly stationary processes (\mathbf{X}, \mathbf{Y}) . Note that the variables in (26) are conditional distributions $q_{y|x}(y^{k_1} | x^{k_1})$, but we are only interested in the \mathbf{m} that they induce.

²As discussed in Appendix A, the set of such processes is non-empty for any legitimate \mathbf{m} .

Lemma 3: If for each n , (P1) has a unique minimizing type \mathbf{m}_n^* , then

$$\|\mathbf{m}_n^* - \hat{\mathbf{m}}_n^*\|_{\text{TV}} \rightarrow 0, \quad \text{a.s.}, \quad (27)$$

where $\hat{\mathbf{m}}_n^*$ is the solution of (26).

Remark: In (26), the only dependence on n is through k_1 .

Therefore, if the encoder knew the distribution of the source, it could solve (26), find a good approximation of \mathbf{m}_n^* , and then use (16) to compute the coefficients required by (P2). The problem is that the encoder does not have this information, and only knows that the source is Markov (but does not know its order). To overcome its lack of information, a reasonable step is to use empirical distribution of the source instead of the true unknown distribution in (26). For $a^{k_1} \in \mathcal{X}^{k_1}$, define the k_1^{th} order empirical distribution of the source as

$$\hat{p}_x^{(k_1)}(a^{k_1}) \triangleq \frac{|\{i : (x_{i-k_1}, \dots, x_{i-1}) = a^{k_1}\}|}{n}. \quad (28)$$

The following lemma shows that for $k_1 = o(\log n)$, $\hat{p}^{(k_1)}$ converges to the actual k_1^{th} order distribution of the source, and therefore can be considered as a good approximation for it.

Lemma 4: For $k_1 = o(\log n)$, and any stationary ergodic Markov source,

$$\|\hat{p}^{(k_1)} - p^{(k_1)}\|_{\text{TV}} \rightarrow 0 \quad \text{a.s.}, \quad (29)$$

where $p^{(k_1)}$ is the true k_1^{th} order distribution of the Markov source.

Assume x^n is generated by a discrete Markov source, and let $\hat{p}_x^{(k_1)}$ be its empirical distribution defined in (28). Consider the following optimization problem

$$\begin{aligned} \min \quad & H(\mathbf{m}) + \alpha \sum_{\beta, \beta', \mathbf{b}, \mathbf{b}'} d_{k_1}(\beta' \mathbf{b}', \beta \mathbf{b}) \hat{p}_x^{(k_1)}(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') \\ \text{s.t.} \quad & m_{\beta, \mathbf{b}} = \sum_{\beta' \mathbf{b}'} \hat{p}_x^{(k_1)}(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}'), \quad \forall \beta, \mathbf{b}, \\ & 0 \leq q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') \leq 1, \quad \forall \beta, \beta', \mathbf{b}, \mathbf{b}' \\ & \sum_{\beta, \mathbf{b}} q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') = 1, \quad \forall \beta', \mathbf{b}', \\ & \sum_{\beta, \beta'} \hat{p}_x^{(k_1)}(\beta' \mathbf{b}') q_{y|x}(\beta \mathbf{b} | \beta' \mathbf{b}') = \sum_{\beta, \beta'} \hat{p}_x^{(k_1)}(\mathbf{b}' \beta') q_{y|x}(\mathbf{b} \beta | \mathbf{b}' \beta'), \quad \forall \mathbf{b}, \mathbf{b}'. \end{aligned} \quad (30)$$

and let $\tilde{\mathbf{m}}_n^*$ denote the output of the above optimization problem.

Lemma 5: For $k_1 = k_1(n) = o(\log n)$,

$$\|\tilde{\mathbf{m}}_n^* - \hat{\mathbf{m}}_n^*\|_{\text{TV}} \rightarrow 0, \quad \text{a.s.}$$

Proof: [outline] The input parameters of the optimization problem (30) are $\{\hat{p}^{(k_1)}(a^{k_1})\}_{a^{k_1} \in \mathcal{X}^{k_1}}$, therefore $\hat{\mathbf{m}}_n^* = \hat{\mathbf{m}}_n^*(\{\hat{p}^{(k_1)}(a^{k_1})\}_{a^{k_1} \in \mathcal{X}^{k_1}})$. On the other hand, both the cost function and the constraints of (30) are continuous both in input parameters and optimization variables. This means that $\hat{\mathbf{m}}_n^*$ in turn is a continuous function of $\{\hat{p}^{(k_1)}(a^{k_1})\}_{a^{k_1} \in \mathcal{X}^{k_1}}$. ■

Let $\{\lambda_{\beta, \mathbf{b}}(n)\}_{\beta, \mathbf{b}}$ denote the optimal values of the coefficients defined at \mathbf{m}_n^* (as given in (16)), and let $\{\hat{\lambda}_{\beta, \mathbf{b}}(n)\}_{\beta, \mathbf{b}}$ be coefficients computed at $\tilde{\mathbf{m}}_n^*$, then

Lemma 6:

$$\max_{\beta, \mathbf{b}} |\lambda_{\beta, \mathbf{b}}(n) - \hat{\lambda}_{\beta, \mathbf{b}}(n)| \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (31)$$

These results suggest that for computing the coefficients we can solve the optimization problem given in (30) (whose complexity can be controlled with the rate of increase of k_1), and then substitute the result

in (16) to obtain the approximate coefficients. After that (P2) defined by these coefficients can be solved using the Viterbi algorithm in a way that will be detailed in the next section. The succession of lemmas detailed in the previous sections then allow us to prove the following theorem.

Theorem 2: Let \mathbf{X} let a stationary and ergodic Markov source, and $R(\mathbf{X}, D)$ denote its rate distortion function. Let \hat{X}^n be the reconstruction sequence obtained using the above scheme for coding X^n choosing $k_1 = k + 1$, where $k = o(\log n)$. Then

$$\mathbb{E} \left[\frac{1}{n} H_k(\mathbf{m}(\hat{X}^n)) + \alpha d_n(X^n, \hat{X}^n) \right] \xrightarrow{n \rightarrow \infty} \min_{D \geq 0} [R(\mathbf{X}, D) + \alpha D]. \quad (32)$$

Remark: Theorem 2 implies the fixed-slope universality of the scheme which does the lossless compression of the reconstruction by first describing its count matrix (costing a number of bits which is negligible for large n) and then doing the conditional entropy coding.

V. VITERBI CODER

As proved in Section III, instead of solving (P1), one can solve (P2) for proper choices of coefficients $\{\lambda_{\mathbf{b}, \beta}\}$. Note that

$$\sum_{\beta, \mathbf{b}} [\lambda_{\beta, \mathbf{b}} m_{\beta, \mathbf{b}}(y^n) + \alpha d_n(x^n, y^n)] = \frac{1}{n} \sum_{i=1}^n [\lambda_{y_i, y_{i-k}^{i-1}} + \alpha d(x_i, y_i)]. \quad (33)$$

This alternative representation of the cost function suggests that instead of using simulated annealing, we can find the sequence that minimizes the cost function by the Viterbi algorithm. For $i = k + 1, \dots, n$, let $s_i = y_{i-k}^i$ be the state at time i , \mathcal{S} be the set of all 2^{k+1} possible states, and for $s = b^{k+1}$ define

$$w(s, i) := \lambda_{b_{k+1}, b^k} + \alpha d(x_i, b_{k+1}).$$

From our definition of the states $s_i = g(s_{i-1}, y_i)$, where $g : \mathcal{S} \times \mathcal{X} \rightarrow \mathcal{S}$. This representation leads to a Trellis diagram corresponding to the evolution of the states $\{s_i\}_{i=k+1}^m$ in which each state has two states leading to it and two states branching from it. Assume that weight $w(s_i)$ is assigned to the edge connecting states s_{i-1} and s_i , i.e., the cost of each edge only depends on the tail state.

It is clear that in our representation, there is a 1-to-1 correspondence between binary sequences y^n and sequences of states $\{s_i\}_{i=k+1}^m$, and minimizing (33) is equivalent to finding the path of minimum weight in the corresponding Trellis diagram, i.e., the path $\{s_i\}_{i=k+1}^n$ that minimizes $\sum_{i=k+1}^n w(s_i, i)$. Solving this minimization can readily be done by Viterbi algorithm which can be described as follows. For each state s , let $\mathcal{L}(s)$ be the two states leading to it, and for any $i > 1$,

$$C(s, i) := \min_{s' \in \mathcal{L}(s)} [w(s') + C(s', i-1)]. \quad (34)$$

For $i = 1$ and $s = b^{k+1}$, let $C(s, 1) := \lambda_{b_{k+1}, b^k} + \alpha d_{k+1}(x^{k+1}, b^{k+1})$. Using this procedure, each state s at each time j has a path of length $j - k - 1$ which is the minimum path among all the possible paths between $i = k + 1$ and $i = j$ such that $s_j = s$. After computing $\{C(s, i)\}_{i \in \{k+1, \dots, n\}}$, at time $i = n$, let

$$s^* = \arg \min_{s \in \mathcal{S}} C(s, n). \quad (35)$$

It is not hard to see that the path leading to s^* is the path of minimum weight among all possible paths.

Note that the computational complexity of this procedure is linear in n but exponential in k because the number of states increases exponentially with k .




Fig. 1. $(d_n(x^n, \hat{x}^n), H_k(\hat{x}^n))$ of output points of Viterbi encoder when the coefficients are computed at $\mathbf{m}[x^n]$. For each value of α , the algorithm is run $L = 20$ times. Here $n = 5000$, $k = 7$, and the source is binary Markov with $q = 0.2$

VI. SIMULATION RESULTS

In this section, some preliminary simulation results of the application of Viterbi encoder described in the previous section is presented. In our simulations, instead of computing the coefficients $\{\lambda_{\beta,b}\}$ from (16) at the optimal point \mathbf{m}_n^* , we compute them at the count matrix of the input sequence x^n , $\mathbf{m}(x^n)$. Fig. 1 demonstrates $(d_n(x^n, y^n), H_k(\mathbf{m}(y^n)))$ of output points of the described algorithm. The block length is $n = 5000$, $k = 7$ and the source is 1st order binary symmetric Markov with transition probability $q = 0.2$. For each value of α the algorithm is applied to $L = 20$ different randomly generated sequences. The reason of getting some points below the rate-distortion curve is that the actual number of bits required for describing \hat{x}^n losslessly to the decoder is larger than $nH_k(\hat{x}^n)$, but converges to it as n grows. For example, for the simple scheme of separately describing the subsequences corresponding to different preceding contexts, this surplus is of order $2^k \log n/n$. The effect of this excess rate is not reflected in the figure, which explains why some points appear below the rate-distortion curve.

It can be observed that for larger values of α the output points are closer to the curve. The reason is that large values of α correspond to small values of distortion, and if the distortion is small then $\mathbf{m}(x^n)$ is a good approximation of $\mathbf{m}(y^n)$.

Finally, Fig. 2 compares the performance of the new Viterbi encoder and the MCMC encoder described in [13]. Here the source is again binary symmetric Markov with $q = 0.2$, and the other parameters are: $k = 7$, $n = 5,000$, $\beta_t = n \log t$, $r = 10n$, where β_t determines the cooling schedule of the MCMC coder and r is its number of iterations. Each point in the figure corresponds to the average performance of $L = 10$ random realizations of the source. It can be observed that even for this simplistic choice of the coefficients the performance of the algorithms are comparable, while the Viterbi encoder for example in this example runs at least 40 times faster.

VII. CONCLUSIONS AND CURRENT DIRECTIONS

In this paper, a new method for universal fixed-slope lossy compression of discrete Markov sources was proposed. The new method achieves the rate-distortion curve for any discrete Markov source. Extending




Fig. 2. Comparing the performances of Viterbi encoder and MCMC encoder proposed in [13]

the algorithm to work on any stationary ergodic source is under current investigation. We believe that in fact the same algorithm works for the general class of stationary ergodic sources, and only the proof should be extended to work in this case as well. Another direction for future work is finding a simple method for approximating the optimal coefficients that would alleviate the need for solving the optimization problem (30).

APPENDIX A: STATIONARITY CONDITION

Assume that we are given a $|\hat{\mathcal{X}}| \times |\hat{\mathcal{X}}|^k$ matrix \mathbf{m} with all elements positive and summing up to one. The question is under what condition(s) this matrix can be $(k + 1)^{\text{th}}$ order stationary distribution of a stationary process. For the ease of notations, instead of matrix \mathbf{m} consider $p(x^{k+1})$ as a distribution defined on $\hat{\mathcal{X}}^{k+1}$. We show that a necessary and sufficient condition is the so-called *stationarity condition* which is

$$\sum_{\beta \in \hat{\mathcal{X}}} p(\beta x^k) = \sum_{\beta \in \hat{\mathcal{X}}} p(x^k \beta). \quad (\text{A-1})$$

- Necessity: The necessity of (A-1) is just a direct result of the definition of stationarity of a process. If $p(x^{k+1})$ is to represent the $(k + 1)^{\text{th}}$ order marginal distribution of a stationary process, then it should be consistent with the k^{th} order marginal distribution as satisfy (A-1).
- Sufficiency: In order to prove the sufficiency, we assume that (A-1) holds, and build a stationary process with $(k + 1)^{\text{th}}$ order marginal distribution of $p(x^{k+1})$. Consider a k^{th} order Markov chain with transition probabilities of

$$q(x_{k+1}|x^k) = \frac{p(x^{k+1})}{p(x^k)}. \quad (\text{A-2})$$

Note that $p(x^k)$ is well-defined by (A-1). Moreover, again from (A-1), $p(x^{k+1})$ is the stationary

distribution of the defined Markov chain, because

$$\sum_{x_1} q(x_{k+1}|x^k)p(x^k) = \sum_{x_1} p(x^{k+1}) = p(x_2^{k+1}). \quad (\text{A-3})$$

Therefore we have found a stationary process that has the desired marginal distribution.

Finally we show that if \mathbf{m} is the count matrix of a sequence y^n , then there exist a stationary process with the marginal distribution coinciding with \mathbf{m} . From what we just proved, we only need to show that (A-1) holds, i.e.,

$$\sum_{\beta} m_{\beta, \mathbf{b}} = \sum_{\beta} m_{b_k, [\beta, b_1, \dots, b_{k-1}]}. \quad (\text{A-4})$$

But this is true because both sides of (A-4) are equal to $|\{i : y_{i+1}^{i+k} = \mathbf{b}\}|/n$.

APPENDIX B: CONCAVITY OF $H(\mathbf{m})$

For simplicity assume that $\mathcal{X} = \hat{\mathcal{X}} = \{0, 1\}$. By definition

$$H(\mathbf{m}) = \sum_{\mathbf{b} \in \{0,1\}^k} (m_{0,\mathbf{b}} + m_{1,\mathbf{b}}) h\left(\frac{m_{0,\mathbf{b}}}{m_{0,\mathbf{b}} + m_{1,\mathbf{b}}}\right), \quad (\text{B-1})$$

where $h(\alpha) = \alpha \log \alpha + \bar{\alpha} \log \bar{\alpha}$ and $\bar{\alpha} = 1 - \alpha$. We need to show that for any $\theta \in [0, 1]$, and empirical count matrices $\mathbf{m}^{(1)}$ and $\mathbf{m}^{(2)}$,

$$\theta H(\mathbf{m}^{(1)}) + \bar{\theta} H(\mathbf{m}^{(2)}) \leq H(\theta \mathbf{m}^{(1)} + \bar{\theta} \mathbf{m}^{(2)}). \quad (\text{B-2})$$

From the concavity of h , it follows that

$$\begin{aligned} & \theta(m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}) h\left(\frac{m_{0,\mathbf{b}}^{(1)}}{m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}}\right) + \bar{\theta}(m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)}) h\left(\frac{m_{0,\mathbf{b}}^{(2)}}{m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)}}\right) \\ &= (\theta(m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}) + \bar{\theta}(m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)})) \sum_{i \in \{1,2\}} \frac{\theta_i(m_{0,\mathbf{b}}^{(i)} + m_{1,\mathbf{b}}^{(i)})}{(\theta(m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}) + \bar{\theta}(m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)}))} h\left(\frac{m_{0,\mathbf{b}}^{(i)}}{m_{0,\mathbf{b}}^{(i)} + m_{1,\mathbf{b}}^{(i)}}\right) \\ &\leq (\theta(m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}) + \bar{\theta}(m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)})) h\left(\frac{\theta m_{0,\mathbf{b}}^{(1)} + \bar{\theta} m_{0,\mathbf{b}}^{(2)}}{\theta(m_{0,\mathbf{b}}^{(1)} + m_{1,\mathbf{b}}^{(1)}) + \bar{\theta}(m_{0,\mathbf{b}}^{(2)} + m_{1,\mathbf{b}}^{(2)})}\right), \end{aligned} \quad (\text{B-3})$$

where $\theta_1 = 1 - \theta_2 = \theta$. Now summing up both sides of (B-3) over all $\mathbf{b} \in \hat{\mathcal{X}}^k$, yields the desired result.

REFERENCES

- [1] C. Shannon, "Coding theorems for a discrete source with a fidelity criterion," *IRE Nat. Conv. Rec.*, part 4, pp. 142-163, 1959.
- [2] R.G. Gallager, "Information Theory and Reliable Communication," New York, NY: John Wiley & Sons, 1968.
- [3] T. Berger, *Rate-distortion theory: A mathematical basis for data compression*, Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [4] T. M. Cover, and J. A. Thomas, *Elements of Information Theory*, New York: Wiley, 1991.
- [5] J. Ziv and A. Lempel, "Compression of individual sequences via variable-rate coding," *IEEE Trans. on Inf. Theory*, 24(5):530-536, Sep. 1978.
- [6] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic coding for data compression", *Commun. Assoc. Comp. Mach.*, vol. 30, no. 6, pp. 520-540, 1987.
- [7] I. Kontoyiannis, "An implementable lossy version of the Lempel Ziv algorithm-Part I: optimality for memoryless sources," *IEEE Trans. on Inform. Theory*, vol. 45, pp. 2293-2305, Nov. 1999.
- [8] E. Yang, Z. Zhang, and T. Berger, "Fixed-slope universal lossy data compression," *IEEE Trans. on Inform. Theory*, vol. 43, no. 5, pp. 1465-1476, Sep. 1997.
- [9] E. H. Yang and J. Kieffer, "Simple universal lossy data compression schemes derived from the Lempel-Ziv algorithm," *IEEE Trans. on Inform. Theory*, vol. 42, no. 1, pp. 239-245, 1996.
- [10] T. Berger, J.D. Gibson, "Lossy source coding," *IEEE Trans. on Inform. Theory*, vol. 44, no. 6, pp. 2693-2723, 1998.
- [11] A. Gersho, R.M. Gray, *Vector Quantization and Signal Compression* Springer, 1992.
- [12] E. Yang, and Z. Zhang, "Variable-Rate Trellis Source Encoding", *IEEE Trans. on Inform. Theory*, vol. 45, no. 2, pp. 586-608, 1999.
- [13] S. Jalali, T. Weissman, "Lossy coding via Markov chain Monte Carlo," IEEE International Symposium on Information Theory, Toronto, Canada, 2008.