



Exponential bounds for convergence of entropy rate approximations in hidden Markov models satisfying a path-mergeability condition

Nicholas F. Travers*

Department of Mathematics, University of California, Davis, United States

Received 5 February 2014; received in revised form 6 June 2014; accepted 9 July 2014

Available online 17 July 2014

Abstract

A hidden Markov model (HMM) is said to have *path-mergeable states* if for any two states i, j there exist a word w and state k such that it is possible to transition from both i and j to k while emitting w . We show that for a finite HMM with path-mergeable states the block estimates of the entropy rate converge exponentially fast. We also show that the path-mergeability property is asymptotically typical in the space of HMM topologies and easily testable.

© 2014 Elsevier B.V. All rights reserved.

Keywords: Hidden Markov model; Entropy rate; Exponential convergence; Path-mergeable

1. Introduction

Hidden Markov models (HMMs) are generalizations of Markov chains in which the underlying Markov state sequence (S_t) is observed through a noisy or lossy channel, leading to a (typically) non-Markovian output process (X_t) . They were first introduced in the 50s as abstract mathematical models [15,7,6], but have since proved useful in a number of concrete applications, such as speech recognition [20,25,3,19] and bioinformatics [26,13,21,14,4].

* Correspondence to: Department of Mathematics, Technion—Israel Institute of Technology, Israel. Tel.: +972 585055757.

E-mail address: travers@tx.technion.ac.il.

<http://dx.doi.org/10.1016/j.spa.2014.07.011>

0304-4149/© 2014 Elsevier B.V. All rights reserved.

One of the earliest major questions in the study of HMMs [6] was to determine the entropy rate of the output process

$$h = \lim_{t \rightarrow \infty} H(X_t | X_1, \dots, X_{t-1}).$$

Unlike for Markov chains, this actually turns out to be quite difficult for HMMs. Even in the finite case no general closed form expression is known, and it is widely believed that no such formula exists. A nice integral expression was provided in [6], but it is with respect to an invariant density that is not directly computable.

In practice, the entropy rate h is instead often estimated directly by the finite length block estimates

$$h(t) = H(X_t | X_1, \dots, X_{t-1}).$$

Thus, it is important to know about the rate of convergence of these estimates to ensure the quality of the approximation.

Moreover, even in cases where the entropy rate can be calculated exactly, such as for unifilar HMMs, the rate of convergence of the block estimates is still of independent interest. It is important for numerical estimation of various complexity measures, such as the excess entropy and transient information [11], and is critical for an observer wishing to make predictions of future output from a finite observation sequence X_1, \dots, X_t . It is also closely related to the rate of memory loss in the initial condition for HMMs, a problem that has been studied extensively in the field of filtering theory [27,1,2,22,8,12,9]. Though, primarily in the case of continuous real-valued outputs with Gaussian noise or similar, rather than the discrete case we study here.

No general bounds are known for the rate of convergence of the estimates $h(t)$, but exponential bounds have been established for finite HMMs (with both a finite internal state set \mathcal{S} and finite output alphabet \mathcal{X}) under various positivity assumptions. The earliest known, and mostly commonly cited, bound is given in [5], for finite functional HMMs with strictly positive transition probabilities in the underlying Markov chain. A somewhat improved bound under the same hypothesis is also given in [17]. Similarly, exponential convergence for finite HMMs with strictly positive symbol emission probabilities and an aperiodic underlying Markov chain is established in [24] (both for state-emitting and edge-emitting HMMs) using results from [22].

Without any positivity assumptions, though, things become substantially more difficult. In the particular case of unifilar HMMs, we have demonstrated exponential convergence in our recent work on synchronization with Crutchfield [29,28]. Also, in Ref. [16] exponential convergence is established under some fairly technical hypotheses in studying entropy rate analyticity. But, no general bounds on the convergence rate of the block estimates $h(t)$ have been demonstrated for all finite HMMs.

Here we prove exponential convergence for finite HMMs (both state-emitting and edge-emitting) satisfying the following simple path-mergeability condition: *For each pair of distinct states i, j there exist a word w and state k such that it is possible to transition from both i and j to k while emitting w .* We also show that this condition is easily testable (in computational time polynomial in the number of states and symbols) and asymptotically typical in the space of HMM topologies, in that a randomly selected topology will satisfy this condition with probability approaching 1, as the number of states goes to infinity. By contrast, the positivity conditions assumed in [5,17,24], as well as the unifilarity hypothesis we assume in [29,28], are satisfied

only for a vanishingly small fraction of HMM topologies in the limit that the number of states becomes large, and the conditions assumed in [16] are, to our knowledge, not easily testable in an automated fashion like the path-mergeability condition. The conditions assumed in [16] are also intuitively somewhat stronger than path-mergeability in that they require the output process (X_t) to have full support and require uniform exponential convergence of conditional probabilities,¹ neither of which are required for path-mergeable HMMs.

The structure of the remainder of the paper is as follows. In Section 2 we introduce the formal framework for our results, including more complete definitions for hidden Markov models and their various properties, as well as the entropy rate and its finite-block estimates. In Section 3 we provide proofs of our exponential convergence results for edge-emitting HMMs satisfying the path-mergeability property. In Section 4 we use the edge-emitting results to establish analogous results for state-emitting HMMs. Finally, in Section 5 we provide the algorithmic test for path-mergeability and demonstrate that this property is asymptotically typical. The proof methods used in Section 3 are based on the original coupling argument used in [5], but are substantially more involved because of our weaker assumption.

2. Definitions and notation

By an *alphabet* \mathcal{X} we mean simply a set of symbols, and by a *word* w over the alphabet \mathcal{X} we mean a finite string $w = x_1, \dots, x_n$ consisting of symbols $x_i \in \mathcal{X}$. The *length* of a word w is the number of symbols it contains and is denoted by $|w|$. \mathcal{X}^* denotes the set of all (finite) words over an alphabet \mathcal{X} , including the empty word λ .

For a sequence (a_n) (of symbols, random variables, real numbers, \dots , etc.) and integers $n \leq m$, a_n^m denotes the finite subsequence a_n, a_{n+1}, \dots, a_m . This notation is also extended in the natural way to the case $m = \infty$ or $n = -\infty$. In the case $n > m$, a_n^m is interpreted as the null sequence or empty word.

2.1. The entropy rate and finite-block estimates

Throughout this section, and the remainder of the paper, we adopt the following standard information theoretic conventions for logarithms (of any base):

$$0 \cdot \log(0) \equiv \lim_{\xi \rightarrow 0^+} \xi \cdot \log(\xi) = 0.$$

$$0 \cdot \log(1/0) \equiv \lim_{\xi \rightarrow 0^+} \xi \cdot \log(1/\xi) = 0.$$

Note that with these conventions the functions $\xi \log(\xi)$ and $\xi \log(1/\xi)$ are both continuous on $[0, 1]$.

Definition 1. The *entropy* $H(X)$ of a discrete random variable X is

$$H(X) \equiv - \sum_{x \in \mathcal{X}} \mathbb{P}(x) \log_2 \mathbb{P}(x)$$

where \mathcal{X} is the alphabet (i.e. set of possible values) of the random variable X and $\mathbb{P}(x) \equiv \mathbb{P}(X = x)$.

¹ E.g., there exist constants $K > 0$ and $0 < \alpha < 1$ such that for any symbol x and symbol sequence x_{-t-n}, \dots, x_{-1} , $t, n \in \mathbb{N}$, $|\mathbb{P}(X_0 = x | X_{-1} = x_{-1}, \dots, X_{-t} = x_{-t}) - \mathbb{P}(X_0 = x | X_{-1} = x_{-1}, \dots, X_{-t-n} = x_{-t-n})| \leq K\alpha^t$.

Definition 2. For discrete random variables X and Y the *conditional entropy* $H(X|Y)$ is

$$\begin{aligned}
 H(X|Y) &\equiv \sum_{y \in \mathcal{Y}} \mathbb{P}(y) \cdot H(X|Y = y) \\
 &= - \sum_{y \in \mathcal{Y}} \mathbb{P}(y) \sum_{x \in \mathcal{X}} \mathbb{P}(x|y) \log_2 \mathbb{P}(x|y)
 \end{aligned}$$

where \mathcal{X} and \mathcal{Y} are, respectively, the alphabets of X and Y , $\mathbb{P}(y) \equiv \mathbb{P}(Y = y)$, and $\mathbb{P}(x|y) \equiv \mathbb{P}(X = x|Y = y)$.

Intuitively, the entropy $H(X)$ is the amount of uncertainty in predicting X , or equivalently, the amount of information obtained by observing X . The conditional entropy $H(X|Y)$ is the average uncertainty in predicting X given the observation of Y . These quantities satisfy the relations

$$0 \leq H(X|Y) \leq H(X) \leq \log_2 |\mathcal{X}|.$$

Definition 3. Let (X_t) be a discrete time stationary process over a finite alphabet \mathcal{X} . The *entropy rate* h of the process (X_t) is the asymptotic per symbol entropy:

$$h \equiv \lim_{t \rightarrow \infty} H(X_1^t)/t \tag{1}$$

where $X_1^t = X_1, \dots, X_t$ is interpreted as a single discrete random variable taking values in the cross product alphabet \mathcal{X}^t .

Using stationarity it may be shown that this limit h always exists and is approached monotonically from above. Further, it may be shown that the entropy rate may also be expressed as the monotonic limit of the conditional next symbol entropies $h(t)$. That is, $h(t) \searrow h$, where

$$h(t) \equiv H(X_t|X_1^{t-1}). \tag{2}$$

The non-conditioned estimates $H(X_1^t)/t$ can approach no faster than a rate of $1/t$. However, the conditional estimates $h(t) = H(X_t|X_1^{t-1})$ can approach much more quickly, and are therefore generally more useful. Our primary goal is to establish an exponential bound on the rate of convergence of these conditional finite-block estimates $h(t)$ for the output process of a HMM satisfying the path-mergeability property defined in Section 2.2.3.

2.2. Hidden Markov models

We will consider here only finite HMMs, meaning that both the internal state set \mathcal{S} and output alphabet \mathcal{X} are finite. There are two primary types: *state-emitting* and *edge-emitting*. The state-emitting variety is the simpler of the two, and also the more commonly studied, so we introduce them first. However, our primary focus will be on edge-emitting HMMs because the path-mergeability condition we study, as well as the block model presentation of Section 3.2.1 used in the proofs, are both more natural in this context.

Definition 4. A *state-emitting hidden Markov model* is a 4-tuple $(\mathcal{S}, \mathcal{X}, \mathcal{T}, \mathcal{O})$ where:

- \mathcal{S} is a finite set of states.
- \mathcal{X} is a finite alphabet of output symbols.
- \mathcal{T} is an $|\mathcal{S}| \times |\mathcal{S}|$ stochastic state transition matrix: $\mathcal{T}_{ij} = \mathbb{P}(S_{t+1} = j|S_t = i)$.

- \mathcal{O} is an $|\mathcal{S}| \times |\mathcal{X}|$ stochastic observation matrix: $\mathcal{O}_{i,x} = \mathbb{P}(X_t = x | S_t = i)$.

The state sequence (S_t) for a state-emitting HMM is a Markov chain with transition matrix \mathcal{T} , and the observed sequence (X_t) has conditional distribution defined by the observation matrix \mathcal{O} :

$$\mathbb{P}(X_n^m = x_n^m | S_0^\infty = s_0^\infty) = \mathbb{P}(X_n^m = x_n^m | S_n^m = s_n^m) = \prod_{t=n}^m \mathcal{O}_{s_t, x_t}.$$

An important special case is when the observation matrix is deterministic, and the symbol X_t is simply a function of the state S_t . This type of HMMs, known as *functional HMMs* or *functions of Markov chains*, are perhaps the most simple variety conceptually, and also were the first type to be heavily studied. The integral expression for the entropy rate provided in [6] and exponential bound on convergence of the block estimates $h(t)$ established in [5] both dealt with HMMs of this type.

Edge-emitting HMMs are an alternative representation in which the symbol X_t depends not simply on the current state S_t but also the next state S_{t+1} , or rather the transition between them.

Definition 5. An *edge-emitting hidden Markov model* is a 3-tuple $(\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ where:

- \mathcal{S} is a finite set of states.
- \mathcal{X} is a finite alphabet of output symbols.
- $\mathcal{T}^{(x)}, x \in \mathcal{X}$, are $|\mathcal{S}| \times |\mathcal{S}|$ sub-stochastic symbol-labeled transition matrices whose sum \mathcal{T} is stochastic. $\mathcal{T}_{ij}^{(x)}$ is the probability of transitioning from i to j on symbol x .

Visually, one can depict an edge-emitting HMM as a directed graph with labeled edges. The vertices are the states, and for each i, j, x with $\mathcal{T}_{ij}^{(x)} > 0$ there is a directed edge from i to j labeled with the transition probability $p = \mathcal{T}_{ij}^{(x)}$ and symbol x . The sum of the probabilities on all outgoing edges from each state is 1.

The operation of the HMM is as follows: From the current state S_t the HMM picks an outgoing edge E_t according to their probabilities, generates the symbol X_t labeling this edge, and then follows the edge to the next state S_{t+1} . Thus, we have the conditional measure

$$\begin{aligned} \mathbb{P}(S_{t+1} = j, X_t = x | S_t = i, S_0^{t-1} = s_0^{t-1}, X_0^{t-1} = x_0^{t-1}) \\ = \mathbb{P}(S_{t+1} = j, X_t = x | S_t = i) = \mathcal{T}_{ij}^{(x)} \end{aligned}$$

for any $i \in \mathcal{S}, t \geq 0$, and possible length- t joint past (s_0^{t-1}, x_0^{t-1}) which may precede state i . From this it follows, of course, that the state sequence (S_t) is a Markov chain with transition matrix $\mathcal{T} = \sum_x \mathcal{T}^{(x)}$. As for state-emitting HMMs, however, we will be interested primarily in the observable output sequence (X_t) rather than the internal state sequence (S_t) , which is assumed to be hidden from the observer.

Remark. It is assumed that, for a HMM of any type, each symbol $x \in \mathcal{X}$ may actually be generated with positive probability. That is, for each $x \in \mathcal{X}$, there exists $i \in \mathcal{S}$ such that $\mathbb{P}(X_0 = x | S_0 = i) > 0$. Otherwise, the symbol x is useless and the alphabet can be restricted to $\mathcal{X}/\{x\}$.

2.2.1. Irreducibility and stationary measures

A HMM, either state-emitting or edge-emitting, is said to be *irreducible* if the underlying Markov chain over states with transition matrix \mathcal{T} is irreducible. In this case, there exists a

unique stationary distribution π over the states satisfying $\pi = \pi\mathcal{T}$, and the joint state-symbol sequence $(S_t, X_t)_{t \geq 0}$ with initial state S_0 drawn according to π is itself a stationary process.

We henceforth assume all HMMs are irreducible and take our generic probability measure \mathbb{P} on joint sequences $(S_t, X_t)_{t \geq 0}$ to be this stationary process measure:

$$\mathbb{P}(\cdot) \equiv \mathbb{P}(\cdot | S_0 \sim \pi) = \sum_i \pi_i \cdot \mathbb{P}(\cdot | S_0 = i).$$

By stationarity, the measure \mathbb{P} may also be extended to biinfinite sequences $(S_t, X_t)_{t \in \mathbb{Z}}$, and we will do so as necessary. The entropy rate h and block estimate $h(t)$ for a HMM M are defined by (1) and (2) where (X_t) is the stationary output process of M with law \mathbb{P} .

2.2.2. Equivalence of model types

Though they are indeed different objects state-emitting and edge-emitting HMMs are equivalent in the following sense: Given an irreducible HMM M of either type there exists an irreducible HMM M' of the other type such that the stationary output processes (X_t) for the two HMMs M and M' are equal in distribution. We recall below the standard conversions.

1. *State-Emitting to Edge-Emitting*—If $M = (\mathcal{S}, \mathcal{X}, \mathcal{T}, \mathcal{O})$ then $M' = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}'^{(x)}\})$, where $\mathcal{T}'_{ij}{}^{(x)} = \mathcal{T}_{ij}\mathcal{O}_{jx}$.
2. *Edge-Emitting to State-Emitting*—If $M = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ then $M' = (\mathcal{S}', \mathcal{X}, \mathcal{T}', \mathcal{O}')$, where $\mathcal{S}' = \{(i, x) : \sum_j \mathcal{T}_{ij}{}^{(x)} > 0\}$, $\mathcal{T}'_{(i,x)(j,y)} = \left(\mathcal{T}_{ij}{}^{(x)} / \sum_k \mathcal{T}_{ik}{}^{(x)}\right) \cdot \left(\sum_k \mathcal{T}_{jk}{}^{(y)}\right)$, and $\mathcal{O}'_{(i,x)y} = \mathbf{1}\{x = y\}$.

Note that the output M' of the edge-emitting to state-emitting conversion is not just an arbitrary state-emitting HMM, but rather is always a functional HMM. Thus, composition of the two conversion algorithms shows that functional HMMs are also equivalent to either state-emitting or edge-emitting HMMs.

2.2.3. Path-mergeability

For a HMM M , let

$$\mathbb{P}_i(\cdot) \equiv \mathbb{P}(\cdot | S_0 = i)$$

be the measure on joint sequences (S_t, X_t) from initial state $S_0 = i$, and let $\delta_i(w)$ be the set of states j that state i can transition to upon emitting the word w :

$$\delta_i(w) \equiv \{j \in \mathcal{S} : \mathbb{P}_i(X_0^{|w|^{-1}} = w, S_{|w|} = j) > 0\}, \quad \text{for an edge-emitting HMM.}$$

$$\delta_i(w) \equiv \{j \in \mathcal{S} : \mathbb{P}_i(X_1^{|w|} = w, S_{|w|} = j) > 0\}, \quad \text{for a state-emitting HMM.}$$

In either case, if w is the null word λ then $\delta_i(w) \equiv \{i\}$, for each i . The following properties will be of central interest.

Definition 6. A pair of states i, j of a HMM M is said to be *path-mergeable* if there exist some word w and state k such that it is possible to transition from both i and j to k on w . That is,

$$k \in \delta_i(w) \cap \delta_j(w). \tag{3}$$

A HMM M is said to have *path-mergeable states*, or be *path-mergeable*, if each pair of distinct states i, j is path-mergeable.

Definition 7. A symbol x is said to be a *flag* or *flag symbol* for a state k if it is possible to transition to state k upon observing the symbol x , from any state i for which it is possible to generate symbol x as the next output. That is,

$$k \in \delta_i(x), \quad \text{for all } i \text{ with } \delta_i(x) \neq \emptyset.^2 \tag{4}$$

A HMM M is said to be *flag-state* if each state k has some flag symbol x .

Our end goal in Section 3 is to prove exponential bounds on convergence of the entropy rate estimates $h(t)$ for edge-emitting HMMs with path-mergeable states. To do so, however, we will first prove similar bounds for edge-emitting HMMs under the flag-state hypothesis and then bootstrap. As we will show in Section 3.2.1, if an edge-emitting HMM has path-mergeable states then the block model M^n , obtained by considering length- n blocks of outputs as single output symbols, is flag-state, for some n . Thus, exponential convergence bounds for flag-state, edge-emitting HMMs pass to exponential bounds for path-mergeable, edge-emitting HMMs by considering block presentations. In Section 4 we will also consider similar questions for state-emitting HMMs. In this case, analogous convergence results follow easily from the results for edge-emitting HMMs by applying the standard state-emitting to edge-emitting conversion.

Remark. Note that if a state-emitting HMM has strictly positive transition probabilities in the underlying Markov chain, as considered in [5,17], it is always path-mergeable. Similarly, if an edge-emitting or state-emitting HMM has strictly positive symbol emission probabilities and an aperiodic underlying Markov chain, as considered in [24], then it is always path-mergeable. The converse of these statements, of course, do not hold. A concrete example will be given in the next subsection.

2.2.4. An illustrative example

To demonstrate the path-mergeability property and motivate its utility, we provide here a simple example of a 4-state, 2-symbol, edge-emitting HMM M , which is path-mergeable, but such that neither M or the equivalent functional HMM M' given by the conversion algorithm of Section 2.2.2 satisfy any of the previously used conditions for establishing exponential convergence of the entropy rate estimates $h(t)$. Thus, although the entropy rate estimates for the output process of either model converge exponentially, this fact could not be deduced from previous results.

Example 1. Let M be the edge-emitting HMM $(\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ with $\mathcal{S} = \{1, 2, 3, 4\}$, $\mathcal{X} = \{a, b\}$, and transition matrices

$$\mathcal{T}^{(a)} = \begin{pmatrix} 0 & 1/3 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2/3 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathcal{T}^{(b)} = \begin{pmatrix} 0 & 2/3 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 0 & 1/3 \\ 1/2 & 1/2 & 0 & 0 \end{pmatrix}.$$

² Note that if x is a flag symbol for k , then after observing the symbol x it is always possible for the HMM to be in state k at the current time, regardless of its initial state or previous outputs. Thus, we call the symbol x a flag for the state k as it signals or ‘flags’ to the observer that k is now possible as the current state of the HMM.




Fig. 1. Graphical depiction of the HMM M described in Example 1. Edges are labeled $p|x$ for the transition probability $p = T_{ij}^{(x)}$ and symbol x . For visual clarity, parallel edges are omitted and each directed edge between states is labeled with all possible symbols upon which the transition between these two states may occur.

Also, let M' be the equivalent (functional) state-emitting HMM constructed from M by the conversion algorithm of Section 2.2.2. For the reader's convenience, a graphical depiction of the HMM M is given in Fig. 1.

M is, indeed, path-mergeable since any two states i, j can transition to some common state k on the three symbol word bbb . That is, $\delta_i(bbb) \cap \delta_j(bbb) \neq \emptyset$, for all $i, j \in \mathcal{S}$. However, M clearly does not satisfy the unifilarity condition assumed in [29,28] (or exactness condition used in [29]), and neither M or M' satisfy any of the positivity conditions assumed in [5,17,24]. Moreover, M' does not satisfy the conditions assumed in [16] since its output process does not have full support (the two symbol word aa is forbidden).

2.2.5. Additional notation

The following additional notation and terminology for an edge-emitting HMM $M = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ will be used below for our proofs in Section 3.

- $\mathbb{P}(w)$ and $\mathbb{P}_i(w)$ denote, respectively, the probability of generating w according to the measures \mathbb{P} and \mathbb{P}_i :

$$\mathbb{P}(w) \equiv \mathbb{P}(X_0^{|w|-1} = w) \quad \text{and} \quad \mathbb{P}_i(w) \equiv \mathbb{P}_i(X_0^{|w|-1} = w)$$

with the conventions $\mathbb{P}(\lambda) \equiv 1$ and $\mathbb{P}_i(\lambda) \equiv 1, i \in \mathcal{S}$, for the null word λ .

- The process language $\mathcal{L}(M)$ for the HMM M is the set of words w of positive probability in its stationary output process (X_t) , and $\mathcal{L}_n(M)$ is the set of length- n words in the process language.

$$\mathcal{L}(M) \equiv \{w \in \mathcal{X}^* : \mathbb{P}(w) > 0\}.$$

$$\mathcal{L}_n(M) \equiv \{w \in \mathcal{L}(M) : |w| = n\}.$$

- For $w \in \mathcal{L}(M)$, $\mathcal{S}(w)$ is the set of states that can generate w .

$$\mathcal{S}(w) \equiv \{i \in \mathcal{S} : \mathbb{P}_i(w) > 0\}.$$

- Finally, $\phi_i(w)$ is the distribution over the current state, or *filtering distribution*, induced by observing the output w from initial state i , and $\phi(w)$ is the distribution over the current state induced by observing w with the initial state chosen according to π .

$$\phi_i(w) \equiv \mathbb{P}_i(S_{|w|} | X_0^{|w|-1} = w).$$

$$\phi(w) \equiv \mathbb{P}(S_{|w|} | X_0^{|w|-1} = w).$$

That is, $\phi_i(w)$ is the probability vector whose k_{th} component is $\phi_i(w)_k = \mathbb{P}_i(S_{|w|} = k | X_0^{|w|-1} = w)$, and $\phi(w)$ is the probability vector whose k_{th} component is $\phi(w)_k = \mathbb{P}(S_{|w|} = k | X_0^{|w|-1} = w)$. In the case $\mathbb{P}_i(w) = 0$ (respectively $\mathbb{P}(w) = 0$), $\phi_i(w)$ (respectively $\phi(w)$) is, by convention, defined to be the null distribution consisting of all zeros.

All above notation may also be used with time indexed symbol sequences x_n^m in place of the word w as well. In this case, the lower time index n is always ignored, and x_n^m is treated simply as a length- $(m - n + 1)$ word. So, for example, $\mathbb{P}(x_n^m) = \mathbb{P}(X_0^{m-n} = x_n^m)$ and $\mathbb{P}_i(x_n^m) = \mathbb{P}_i(X_0^{m-n} = x_n^m)$.

3. Results for edge-emitting HMMs

In this section we establish exponential convergence of the entropy rate block estimates $h(t)$ for edge-emitting HMMs with path-mergeable states. The basic structure of the arguments is as follows:

1. We establish exponential bounds for flag-state (edge-emitting) HMMs.
2. We extend to path-mergeable (edge-emitting) HMMs by passing to a block model representation (see Section 3.2.1).

Exponential convergence in the flag-state case is established by the following steps:

- (i) By coupling to a sequence of i.i.d. uniform($[0, 1]$) random variables and applying large deviation estimates, we show that the set of “good” length- t symbol sequences G_t defined by (6) has combined probability $1 - O(\text{exponentially small})$.
- (ii) Using a coupling argument similar to that given in [5] we show that $\|\phi_k(x_0^{t-1}) - \widehat{\phi}_k(x_0^{t-1})\|_{\text{TV}}$ is exponentially small, for any symbol sequence $x_0^{t-1} \in G_t$ and states $k, \widehat{k} \in \mathcal{S}(x_0^{t-1})$.
- (iii) Using (ii) we show that the difference $H(X_t | X_0^{t-1} = x_0^{t-1}) - H(X_t | X_0^{t-1} = x_0^{t-1}, S_0)$ is exponentially small for any $x_0^{t-1} \in G_t$. Combining this with the fact that $\mathbb{P}(G_t^c)$ is exponentially small shows that the difference $H(X_t | X_0^{t-1}) - H(X_t | X_0^{t-1}, S_0)$ is exponentially small, from which exponential convergence of the estimates $h(t)$ follows easily by a sandwiching argument.

3.1. Under flag-state assumption

Throughout Section 3.1 we assume $M = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ is a flag-state, edge-emitting HMM, and denote the flag symbol for state j by $y_j: j \in \delta_i(y_j)$, for all i with $\delta_i(y_j) \neq \emptyset$. Also, for notational convenience, we assume the state set is $\mathcal{S} = \{1, \dots, |\mathcal{S}|\}$ and the output alphabet is $\mathcal{X} = \{1, \dots, |\mathcal{X}|\}$. The constants $p_*, q_*, r_*, \eta, \alpha_1, \alpha_2$ for the HMM M are defined as follows:

$$p_j \equiv \mathbb{P}(X_0 = y_j | S_1 = j) \quad \text{and} \quad p_* \equiv \min_j p_j.$$

$$\begin{aligned}
 q_j &\equiv \min_{i \in \mathcal{S}(y_j)} \mathbb{P}_i(S_1 = j | X_0 = y_j) \quad \text{and} \quad q_* \equiv \min_j q_j. \\
 r_* &\equiv \min_{i,j} \pi_i / \pi_j. \\
 \eta &\equiv \frac{p_* r_*}{2|\mathcal{S}|}, \quad \alpha_1 \equiv \exp\left(-\frac{p_*^2 r_*^2}{2|\mathcal{S}|^2}\right), \quad \alpha_2 \equiv (1 - q_*^2)^\eta. \tag{5}
 \end{aligned}$$

Note that, under the flag-state assumption, we always have $p_*, q_*, r_* \in (0, 1]$, $\eta, \alpha_1 \in (0, 1)$, and $\alpha_2 \in [0, 1)$. Our primary objective is to prove the following theorem.

Theorem 1. *The entropy rate block estimates $h(t)$ for the HMM M converge exponentially with*

$$\limsup_{t \rightarrow \infty} \{h(t) - h\}^{1/t} \leq \alpha$$

where $\alpha \equiv \max\{\alpha_1, \alpha_2\}$.

The proof is, however, fairly lengthy. So, we have divided it into subsections following the steps (i)–(iii) outlined above. The motivation for the definitions of the various constants should become clear in the proof.

In the first two Sections 3.1.1 and 3.1.2 we will need to define a number of auxiliary random variables, which are not functions of the HMM random variables X_t and S_t . For notional simplicity, we denote also the probability measure for these additional random variables by the generic \mathbb{P} . Formally, we assume all additional random variables live on a common underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with the HMM random variables X_t and S_t , but are independent from them.

3.1.1. Upper bound for $\mathbb{P}(G_t^c)$

For $x_0^{t-1} \in \mathcal{L}_t(M)$, define

$$N(x_0^{t-1}) \equiv |\{0 \leq \tau \leq t - 1 : x_\tau = y_k, \text{ for some } k \in \mathcal{S}, \text{ and } \mathbb{P}_k(x_{\tau+1}^{t-1}) \geq \mathbb{P}_j(x_{\tau+1}^{t-1}), \forall j \in \mathcal{S}\}|,$$

and, for $t \in \mathbb{N}$, let

$$G_t \equiv \{x_0^{t-1} \in \mathcal{L}_t(M) : N(x_0^{t-1}) \geq \eta t\} \quad \text{and} \quad G_t^c \equiv \mathcal{L}_t(M) / G_t. \tag{6}$$

Recall that, according to our conventions, if $\tau = t - 1$ then $x_{\tau+1}^{t-1} = x_t^{t-1}$ is the empty word λ , and $\mathbb{P}_j(\lambda) = 1$, for each state $j \in \mathcal{S}$. So, $N(x_0^{t-1})$ is indeed well defined.

The purpose of this subsection is to prove the following lemma.

Lemma 1. $\mathbb{P}(G_t^c) \leq \alpha_1^t$, for all $t \in \mathbb{N}$.

The proof of the lemma is based on large deviation estimates for an auxiliary sequence of random variables $(\tilde{X}_t)_{t \in -\mathbb{N}}$, which is equal in law to the stationary output sequence of the HMM M on negative integer times, but is constructed explicitly from a sequence of i.i.d. uniform($[0, 1]$) random variables. This explicit construction will give a coupling that allows large deviation estimates for the i.i.d. sequence to be translated to large deviation estimates for the (\tilde{X}_t) sequence.

We will use only the fact that $p_* > 0$ in the proof, which is weaker than the flag-state condition. However, the stronger assumption $q_* > 0$, which is equivalent to the flag-state condition, will be needed for Lemma 2 in the next section.

Construction of the (\tilde{X}_t) sequence

- For each $w \in \mathcal{L}(M)$ fix a partition P_w of $[0, 1]$ into subintervals $I_{w,x}$, $x \in \mathcal{X}$, such that:
 1. The length of each $I_{w,x}$ is $\mathbb{P}(X_{-|w|-1} = x | X_{-|w|}^{-1} = w)$.
 2. All length 0 intervals $I_{w,x}$ are taken to be the empty set, rather than points.
 3. The leftmost interval $I_{w,x}$ is $I_{w,y_{k(w)}} = [0, \mathbb{P}(X_{-|w|-1} = y_{k(w)} | X_{-|w|}^{-1} = w)]$ where $k(w) = \min\{k : \mathbb{P}_k(w) \geq \mathbb{P}_j(w), \forall j \in \mathcal{S}\}$.
- Let $(U_t)_{t \in -\mathbb{N}}$ be an i.i.d. sequence of uniform $([0,1])$ random variables.
- The sequence $(\tilde{X}_t)_{t \in -\mathbb{N}}$ is defined inductively as follows:
 1. $\tilde{X}_{-1} = x$ if and only if $U_{-1} \in I_{\lambda,x}$, (where λ is the empty word).
 2. Conditioned on $\tilde{X}_t^{-1} = w$ ($t \leq -1$), $\tilde{X}_{t-1} = x$ if and only if $U_{t-1} \in I_{w,x}$.

By induction on the length of w , it is easily seen that $\mathbb{P}(X_{-|w|}^{-1} = w) = \mathbb{P}(\tilde{X}_{-|w|}^{-1} = w)$ for any word $w \in \mathcal{X}^*$, so $(\tilde{X}_t)_{t \in -\mathbb{N}} \stackrel{d.}{=} (X_t)_{t \in -\mathbb{N}}$.

Proof of Lemma 1. If $w \in \mathcal{L}(M)$ and $k \in \mathcal{S}$ satisfy $\mathbb{P}_k(w) \geq \mathbb{P}_j(w)$, for all j , then by Bayes Theorem

$$\mathbb{P}(S_{-|w|} = k | X_{-|w|}^{-1} = w) \geq r_* \cdot \mathbb{P}(S_{-|w|} = j | X_{-|w|}^{-1} = w), \quad \text{for all } j.$$

Thus, since $X_{-\infty}^{t-1}$ and X_t^∞ are independent given the current state S_t , it follows that for any such w and k

$$\begin{aligned} \mathbb{P}(X_{-|w|-1} = y_k | X_{-|w|}^{-1} = w) &\geq \mathbb{P}(S_{-|w|} = k | X_{-|w|}^{-1} = w) \cdot \mathbb{P}(X_{-|w|-1} = y_k | S_{-|w|} = k) \\ &\geq \left(\frac{1}{|\mathcal{S}|} \cdot \min_{j \in \mathcal{S}} \left\{ \frac{\mathbb{P}(S_{-|w|} = k | X_{-|w|}^{-1} = w)}{\mathbb{P}(S_{-|w|} = j | X_{-|w|}^{-1} = w)} \right\} \right) \cdot \mathbb{P}(X_{-|w|-1} = y_k | S_{-|w|} = k) \\ &\geq \frac{r_*}{|\mathcal{S}|} \cdot p_*. \end{aligned}$$

So, by definition of the intervals $I_{w,x}$, we have

$$I_{w,y_{k(w)}} = [0, \mathbb{P}(X_{-|w|-1} = y_{k(w)} | X_{-|w|}^{-1} = w)] \supseteq [0, p_* r_* / |\mathcal{S}|], \quad \text{for all } w \in \mathcal{X}^*. \quad (7)$$

Now, let $K_0 \equiv 1$ and, for $t \in \mathbb{N}$, let $K_t \equiv k(\tilde{X}_{-t}^{-1})$. That is, $K_t = k(w)$ on the event $\tilde{X}_{-t}^{-1} = w$. Also, define N_t and N'_t by

$$\begin{aligned} N_t &\equiv |\{0 \leq \tau \leq t - 1 : \tilde{X}_{-\tau-1} = y_{K_\tau}\}|, \\ N'_t &\equiv |\{0 \leq \tau \leq t - 1 : U_{-\tau-1} \leq p_* r_* / |\mathcal{S}|\}|. \end{aligned}$$

By (7) we know that the random interval $I_{\tilde{X}_{-\tau}^{-1}, y_{K_\tau}}$ always contains $[0, p_* r_* / |\mathcal{S}|]$. So,

$$U_{-\tau-1} \leq p_* r_* / |\mathcal{S}| \implies U_{-\tau-1} \in I_{\tilde{X}_{-\tau}^{-1}, y_{K_\tau}} \implies \tilde{X}_{-\tau-1} = y_{K_\tau}$$

and, therefore, N_t is always lower bounded by N'_t :

$$N_t \geq N'_t, \quad \text{for each } t \in \mathbb{N}.$$

But, N'_t is simply $\sum_{\tau=0}^{t-1} Z_\tau$, where $Z_\tau \equiv \mathbb{1}_{\{U_{-\tau-1} \leq p_* r_* / |\mathcal{S}|\}}$. Since the Z_τ s are i.i.d. with $\mathbb{P}(Z_\tau = 1) = p_* r_* / |\mathcal{S}|$, $\mathbb{P}(Z_\tau = 0) = 1 - p_* r_* / |\mathcal{S}|$ we may, therefore, apply Hoeffding's inequality to conclude that

$$\mathbb{P}(N_t < \eta t) \leq \mathbb{P}(N'_t < \eta t) \leq \alpha'_1.$$

The lemma follows since $\tilde{X}_{-t}^{-1} \stackrel{d.}{=} X_{-t}^{-1} \stackrel{d.}{=} X_0^{t-1}$, for each $t \in \mathbb{N}$, which implies

$$\mathbb{P}(G_t^c) \equiv \mathbb{P}\left(\{x_0^{t-1} \in \mathcal{L}_t(M) : N(x_0^{t-1}) < \eta t\}\right) \leq \mathbb{P}(N_t < \eta t). \quad \square$$

3.1.2. Pair chain coupling bound

In this subsection we prove the following lemma.

Lemma 2. For any $t \in \mathbb{N}$, $x_0^{t-1} \in G_t$, and states $k, \hat{k} \in \mathcal{S}(x_0^{t-1})$

$$\|\phi_k(x_0^{t-1}) - \phi_{\hat{k}}(x_0^{t-1})\|_{TV} \leq \alpha_2^t$$

where $\|\mu - \nu\|_{TV} \equiv \frac{1}{2} \|\mu - \nu\|_1$ is the total variational norm between two probability distributions (probability vectors) μ and ν .

The proof is based on a coupling argument for an auxiliary Markov chain $(R_\tau, \widehat{R}_\tau)_{\tau=0}^t$ on state space $|\mathcal{S}| \times |\mathcal{S}|$, which we call the *pair chain*.

The pair chain

For $x_0^{t-1} \in G_t$ and states $k, \hat{k} \in \mathcal{S}(x_0^{t-1})$, the *pair chain* $(R_\tau, \widehat{R}_\tau)_{\tau=0}^t$ is the time-inhomogeneous Markov chain on $|\mathcal{S}| \times |\mathcal{S}|$ with initial state $(R_0, \widehat{R}_0) = (k, \hat{k})$ and transition probabilities

$$\begin{aligned} \mathbb{P}(R_{\tau+1} = j, \widehat{R}_{\tau+1} = \widehat{j} | R_\tau = i, \widehat{R}_\tau = \widehat{i}) \\ = \begin{cases} \mathbb{P}(S_{\tau+1} = j | S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1}) \cdot \mathbb{P}(S_{\tau+1} = \widehat{j} | S_\tau = \widehat{i}, X_\tau^{t-1} = x_\tau^{t-1}), & \text{if } i \neq \widehat{i} \\ \mathbb{P}(S_{\tau+1} = j | S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1}), & \text{if } i = \widehat{i} \text{ and } j = \widehat{j} \\ 0, & \text{if } i = \widehat{i} \text{ and } j \neq \widehat{j} \end{cases} \end{aligned}$$

for $0 \leq \tau \leq t - 1$. By marginalizing it follows that the state sequences (R_τ) and (\widehat{R}_τ) are each individually (time-inhomogeneous) Markov chains with transition probabilities

$$\begin{aligned} \mathbb{P}(R_{\tau+1} = j | R_\tau = i) &= \mathbb{P}(S_{\tau+1} = j | S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1}), \quad \text{and} \\ \mathbb{P}(\widehat{R}_{\tau+1} = \widehat{j} | \widehat{R}_\tau = \widehat{i}) &= \mathbb{P}(S_{\tau+1} = \widehat{j} | S_\tau = \widehat{i}, X_\tau^{t-1} = x_\tau^{t-1}). \end{aligned}$$

Thus, for any $r_0^t, \widehat{r}_0^t \in \mathcal{S}^{t+1}$,

$$\begin{aligned} \mathbb{P}(R_0^t = r_0^t) &= \mathbb{P}(S_0^t = r_0^t | S_0 = k, X_0^{t-1} = x_0^{t-1}), \quad \text{and} \\ \mathbb{P}(\widehat{R}_0^t = \widehat{r}_0^t) &= \mathbb{P}(S_0^t = \widehat{r}_0^t | S_0 = \widehat{k}, X_0^{t-1} = x_0^{t-1}). \end{aligned}$$

So, by [23, Proposition 4.7], we have the following coupling bound:

$$\begin{aligned} \|\phi_k(x_0^{t-1}) - \phi_{\hat{k}}(x_0^{t-1})\|_{TV} &\equiv \|\mathbb{P}_k(S_t | X_0^{t-1} = x_0^{t-1}) - \mathbb{P}_{\hat{k}}(S_t | X_0^{t-1} = x_0^{t-1})\|_{TV} \\ &\leq \mathbb{P}(R_t \neq \widehat{R}_t). \end{aligned} \tag{8}$$

This bound will be used below to prove the lemma.

Proof of Lemma 2. Fix $x_0^{t-1} \in G_t$ and states $k, \hat{k} \in \mathcal{S}(x_0^{t-1})$, and define the time set Γ by

$$\Gamma \equiv \{0 \leq \tau \leq t - 1 : x_\tau = y_\ell, \text{ for some } \ell \in \mathcal{S}, \text{ and } \mathbb{P}_\ell(x_{\tau+1}^{t-1}) \geq \mathbb{P}_j(x_{\tau+1}^{t-1}), \forall j \in \mathcal{S}\}.$$

Also, for $\tau \in \Gamma$, let ℓ_τ be the state ℓ such that $x_\tau = y_\ell$. Then, for each $\tau \in \Gamma$ and $i \in \mathcal{S}(x_\tau^{t-1})$,

$$\begin{aligned} \mathbb{P}(X_{\tau+1}^{t-1} = x_{\tau+1}^{t-1} | S_\tau = i, X_\tau = y_{\ell_\tau}) &\leq \max_{j \in \mathcal{S}} \mathbb{P}(X_{\tau+1}^{t-1} = x_{\tau+1}^{t-1} | S_{\tau+1} = j) \\ &= \mathbb{P}(X_{\tau+1}^{t-1} = x_{\tau+1}^{t-1} | S_{\tau+1} = \ell_\tau). \end{aligned}$$

So, by Bayes Theorem, for any such $\tau \in \Gamma$ and $i \in \mathcal{S}(x_\tau^{t-1})$ we have

$$\begin{aligned} \mathbb{P}(S_{\tau+1} = \ell_\tau | S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1}) &= \frac{\mathbb{P}(S_{\tau+1} = \ell_\tau, S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1})}{\mathbb{P}(S_\tau = i, X_\tau^{t-1} = x_\tau^{t-1})} \\ &= \frac{\mathbb{P}(S_\tau = i, X_\tau = y_{\ell_\tau}) \cdot \mathbb{P}(S_{\tau+1} = \ell_\tau | S_\tau = i, X_\tau = y_{\ell_\tau}) \cdot \mathbb{P}(X_{\tau+1}^{t-1} = x_{\tau+1}^{t-1} | S_{\tau+1} = \ell_\tau)}{\mathbb{P}(S_\tau = i, X_\tau = y_{\ell_\tau}) \cdot \mathbb{P}(X_{\tau+1}^{t-1} = x_{\tau+1}^{t-1} | S_\tau = i, X_\tau = y_{\ell_\tau})} \\ &\geq \mathbb{P}(S_{\tau+1} = \ell_\tau | S_\tau = i, X_\tau = y_{\ell_\tau}) \geq q_*. \end{aligned}$$

Combining this inequality with the pair chain coupling bound (8) gives

$$\begin{aligned} \|\phi_k(x_0^{t-1}) - \phi_{\widehat{k}}(x_0^{t-1})\|_{\text{TV}} &\leq \mathbb{P}(R_t \neq \widehat{R}_t) \\ &= \prod_{\tau=0}^{t-1} \mathbb{P}(R_{\tau+1} \neq \widehat{R}_{\tau+1} | R_\tau \neq \widehat{R}_\tau) \\ &\leq \prod_{\tau \in \Gamma} \mathbb{P}(R_{\tau+1} \neq \widehat{R}_{\tau+1} | R_\tau \neq \widehat{R}_\tau) \\ &\leq \prod_{\tau \in \Gamma} \max_{i, \widehat{i} \in \mathcal{S}(x_\tau^{t-1}), i \neq \widehat{i}} \mathbb{P}(R_{\tau+1} \neq \widehat{R}_{\tau+1} | R_\tau = i, \widehat{R}_\tau = \widehat{i}) \\ &\leq \prod_{\tau \in \Gamma} \max_{i, \widehat{i} \in \mathcal{S}(x_\tau^{t-1}), i \neq \widehat{i}} (1 - \mathbb{P}(R_{\tau+1} = \widehat{R}_{\tau+1} = \ell_\tau | R_\tau = i, \widehat{R}_\tau = \widehat{i})) \\ &\leq \prod_{\tau \in \Gamma} (1 - q_*^2) \\ &\leq (1 - q_*^2)^\eta \\ &= \alpha_2^t. \end{aligned}$$

The inequality in the second to last line follows from the fact that $x_0^{t-1} \in G_t$, which implies $|\Gamma| \geq \eta t$. (Note: We have assumed in the proof that $\mathbb{P}(R_\tau \neq \widehat{R}_\tau) > 0$, for all $0 \leq \tau \leq t - 1$. If this is not the case then $\mathbb{P}(R_t \neq \widehat{R}_t) = 0$, so the conclusion follows trivially from (8).) \square

3.1.3. Convergence of entropy rate approximations

In this subsection we will use the bounds given in the previous two subsections for $\mathbb{P}(G_t^c)$ and the difference in induced distributions $\|\phi_k(x_0^{t-1}) - \phi_{\widehat{k}}(x_0^{t-1})\|_{\text{TV}}$, $x_0^{t-1} \in G_t$ and $k, \widehat{k} \in \mathcal{S}(x_0^{t-1})$, to establish Theorem 1. First, however, we will need two more simple lemmas.

Lemma 3. For $k \in \mathcal{S}$, let $\mathbb{P}_k(X_0)$ be the distribution of X_0 given the initial state is $S_0 = k$, and for a probability distribution μ on \mathcal{S} , let $\mathbb{P}_\mu(X_0) = \sum_k \mu_k \mathbb{P}_k(X_0)$ be the distribution of X_0 when S_0 is chosen according to μ . Then, for any two probability distributions μ and ν on \mathcal{S}

$$\|\mathbb{P}_\mu(X_0) - \mathbb{P}_\nu(X_0)\|_{\text{TV}} \leq \|\mu - \nu\|_{\text{TV}}.$$

Proof. It is equivalent to prove the statement for 1-norms. In this case, we have

$$\begin{aligned} \|\mathbb{P}_\mu(X_0) - \mathbb{P}_\nu(X_0)\|_1 &= \left\| \sum_k \mu_k \cdot \mathbb{P}_k(X_0) - \sum_k \nu_k \cdot \mathbb{P}_k(X_0) \right\|_1 \\ &\leq \sum_k |\mu_k - \nu_k| \cdot \|\mathbb{P}_k(X_0)\|_1 \\ &= \|\mu - \nu\|_1. \quad \square \end{aligned}$$

Lemma 4. Let $\mu = (\mu_1, \dots, \mu_N)$ and $\nu = (\nu_1, \dots, \nu_N)$ be two probability measures on the finite set $\{1, \dots, N\}$. If $\|\mu - \nu\|_{TV} \leq \epsilon$, with $\epsilon \in [0, 1/e]$, then

$$|H(\mu) - H(\nu)| \leq N\epsilon \log_2(1/\epsilon).$$

Proof. If $\|\mu - \nu\|_{TV} \leq \epsilon$ then $|\mu_k - \nu_k| \leq \epsilon$, for all k . Thus,

$$\begin{aligned} |H(\mu) - H(\nu)| &= \left| \sum_k \nu_k \log_2(\nu_k) - \mu_k \log_2(\mu_k) \right| \\ &\leq \sum_k |\nu_k \log_2(\nu_k) - \mu_k \log_2(\mu_k)| \\ &\leq N \cdot \max_{\epsilon' \in [0, \epsilon]} \max_{\xi \in [0, 1-\epsilon']} |(\xi + \epsilon') \log_2(\xi + \epsilon') - \xi \log_2(\xi)| \\ &= N \cdot \max_{\epsilon' \in [0, \epsilon]} \epsilon' \log_2(1/\epsilon') \\ &= N \cdot \epsilon \log_2(1/\epsilon). \end{aligned}$$

The last two equalities, for $0 \leq \epsilon' \leq \epsilon \leq 1/e$, may be verified using single variable calculus techniques for maximization (recalling our conventions for continuous extensions of the functions $\xi \log(\xi)$ and $\xi \log(1/\xi)$ given in Section 2.1). \square

Proof of Theorem 1. If $\alpha_2 = 0$, let $t_0 = 1$. Otherwise, let $t_0 = \lceil \log_{\alpha_2}(1/e) \rceil$. We claim, first, that for each $t \geq t_0$ and $x_0^{t-1} \in G_t$,

$$H(X_t|X_0^{t-1} = x_0^{t-1}) - H(X_t|X_0^{t-1} = x_0^{t-1}, S_0) \leq |\mathcal{X}| \cdot \alpha_2^t \cdot \log_2(1/\alpha_2^t). \tag{9}$$

To see this, note that for any $k \in \mathcal{S}(x_0^{t-1})$ Lemmas 2 and 3 imply

$$\begin{aligned} &\|\mathbb{P}_k(X_t|X_0^{t-1} = x_0^{t-1}) - \mathbb{P}(X_t|X_0^{t-1} = x_0^{t-1})\|_{TV} \\ &\leq \|\mathbb{P}_k(S_t|X_0^{t-1} = x_0^{t-1}) - \mathbb{P}(S_t|X_0^{t-1} = x_0^{t-1})\|_{TV} \\ &\leq \max_{\hat{k} \in \mathcal{S}(x_0^{t-1})} \|\mathbb{P}_k(S_t|X_0^{t-1} = x_0^{t-1}) - \mathbb{P}_{\hat{k}}(S_t|X_0^{t-1} = x_0^{t-1})\|_{TV} \\ &= \max_{\hat{k} \in \mathcal{S}(x_0^{t-1})} \|\phi_k(x_0^{t-1}) - \phi_{\hat{k}}(x_0^{t-1})\|_{TV} \\ &\leq \alpha_2^t. \end{aligned}$$

So, by Lemma 4,

$$H(X_t|X_0^{t-1} = x_0^{t-1}) - H(X_t|X_0^{t-1} = x_0^{t-1}, S_0 = k) \leq |\mathcal{X}| \cdot \alpha_2^t \cdot \log_2(1/\alpha_2^t)$$

for each $k \in \mathcal{S}(x_0^{t-1})$, as $t \geq t_0$. The claim (9) follows since the entropy $H(X_t|X_0^{t-1} = x_0^{t-1}, S_0)$ is, by definition, a weighted average of the entropies $H(X_t|X_0^{t-1} = x_0^{t-1}, S_0 = k)$, $k \in \mathcal{S}(x_0^{t-1})$:

$$\begin{aligned} & H(X_t|X_0^{t-1} = x_0^{t-1}, S_0) \\ & \equiv \sum_{k \in \mathcal{S}(x_0^{t-1})} \mathbb{P}(S_0 = k | X_0^{t-1} = x_0^{t-1}) \cdot H(X_t|X_0^{t-1} = x_0^{t-1}, S_0 = k). \end{aligned}$$

Now, for any $t \in \mathbb{N}$,

$$h = \lim_{\tau \rightarrow \infty} H(X_0|X_{-\tau}^{-1}) \geq \lim_{\tau \rightarrow \infty} H(X_0|X_{-\tau}^{-1}, S_{-\tau}) = H(X_0|X_{-1}^{-1}, S_{-1}) = H(X_t|X_0^{t-1}, S_0).$$

Thus, using Lemma 1 and the estimate (9) the difference $h(t + 1) - h$ may be bounded as follows for all $t \geq t_0$:

$$\begin{aligned} h(t + 1) - h &= H(X_t|X_0^{t-1}) - h \\ &\leq H(X_t|X_0^{t-1}) - H(X_t|X_0^{t-1}, S_0) \\ &= \sum_{x_0^{t-1} \in \mathcal{L}_t(M)} \mathbb{P}(x_0^{t-1}) \cdot \left[H(X_t|X_0^{t-1} = x_0^{t-1}) - H(X_t|X_0^{t-1} = x_0^{t-1}, S_0) \right] \\ &\leq \mathbb{P}(G_t) \cdot (|\mathcal{X}| \cdot \alpha_2^t \cdot \log_2(1/\alpha_2^t)) + \mathbb{P}(G_t^c) \cdot \log_2 |\mathcal{X}| \\ &\leq 1 \cdot (|\mathcal{X}| \cdot \alpha_2^t \cdot \log_2(1/\alpha_2^t)) + \alpha_1^t \cdot \log_2 |\mathcal{X}|. \end{aligned}$$

The theorem follows directly from this inequality. \square

3.2. Under path-mergeable assumption

Building on the results of the previous section for flag-state HMMs, we now proceed to the proof of exponential convergence of the entropy rate block estimates for path-mergeable HMMs. The general approach is as follows:

1. We show that for any path-mergeable HMM M there is some $n \in \mathbb{N}$ such that the block model M^n (defined below) is flag-state.
2. We combine Point 1 with the exponential convergence bound for flag-state HMMs given by Theorem 1 to obtain the desired bound for path-mergeable HMMs.

The main theorem, Theorem 2, will be given in Section 3.2.2 after introducing the block models in Section 3.2.1.

3.2.1. Block models

Definition 8. Let $M = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ be an edge-emitting hidden Markov model. For $n \in \mathbb{N}$, the block model M^n is the triple $(\mathcal{S}, \mathcal{W}, \{\mathcal{Q}^{(w)}\})$ where:

- $\mathcal{W} = \mathcal{L}_n(M)$ is the set of length- n words of positive probability.
- $\mathcal{Q}_{ij}^{(w)} = \mathbb{P}_i(X_0^{n-1} = w, S_n = j)$ is the n -step transition probability from i to j on w .

One can show that if M is irreducible and n is relatively prime to the period of M 's graph, $\text{per}(M)$, then M^n is also irreducible. Further, in this case M and M^n have the same stationary distribution π and the stationary output process of M^n is the same (i.e. equal in distribution) to

the stationary output process for M , when the latter is considered over length- n blocks rather than individual symbols. That is, for any $w_0^{t-1} \in \mathcal{W}^t$,

$$\mathbb{P}^n(W_0^{t-1} = w_0^{t-1}) = \mathbb{P}(X_0^{tn-1} = w_0^{t-1})$$

where W_t denotes the t th output of the block model M^n , and \mathbb{P}^n is the probability distribution over the output sequence (W_t) when the initial state of the block model is chosen according to the stationary distribution π .

The following important lemma allows us to reduce questions for path-mergeable HMMs to analogous questions for flag-state HMMs by considering such block presentations.

Lemma 5. *If $M = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}^{(x)}\})$ is an edge-emitting HMM with path-mergeable states, then there exists some $n \in \mathbb{N}$, relatively prime to $\text{per}(M)$, such that the block model $M^n = (\mathcal{S}, \mathcal{W}, \{\mathcal{Q}^{(w)}\})$ is flag-state.*

Proof. Extending the flag symbol definition (4) to multi-symbol words, we will say a word $w \in \mathcal{L}(M)$ is a *flag word* for state k if all states that can generate w can transition to k on w :

$$k \in \delta_i(w), \quad \text{for all } i \text{ with } \mathbb{P}_i(w) > 0.$$

The following two facts are immediate from this definition:

1. If w is a flag word for state i and v is any word with $j \in \delta_i(v)$, then $wv \in \mathcal{L}(M)$ is a flag word for state j .
2. If w is a flag word for state j and v is any word with $wv \in \mathcal{L}(M)$, then wv is also a flag word for j .

Irreducibility of the HMM M along with Fact 1 ensure that if some state $i \in \mathcal{S}$ has a flag word then each state $j \in \mathcal{S}$ has a flag word. Moreover, by Fact 2, we know that in this case the flag words may all be chosen of some fixed length n , relatively prime to $\text{per}(M)$, which implies the block model M^n is flag-state. Thus, it suffices to show that there is a single state i with a flag word.

Below we will explicitly construct a word v^* , which is a flag word for some state i^* . For notational convenience, we will assume throughout that the state set is $\mathcal{S} = \{1, 2, \dots, |\mathcal{S}|\}$ and the alphabet is $\mathcal{X} = \{1, 2, \dots, |\mathcal{X}|\}$. Also, for $i, j \in \mathcal{S}$ we will denote by w_{ij} and k_{ij} the special word w and state k satisfying the path-mergeability condition (3), so that $k_{ij} \in \delta_i(w_{ij}) \cap \delta_j(w_{ij})$.

v^* is then constructed by the following algorithm:

- $t := 0, v_0 := \lambda$ (the empty word), $i_0 := 1, \mathcal{R} := \mathcal{S}/\{1\}$
- While $\mathcal{R} \neq \emptyset$ do:
 - $k_t := \min\{k : k \in \mathcal{R}\}$
 - $j_t := \min\{j : j \in \delta_{k_t}(v_t)\}$
 - $w_t := w_{i_t j_t}$
 - $v_{t+1} := v_t w_t$
 - $i_{t+1} := k_{i_t j_t}$
 - $\mathcal{R} := \mathcal{R} / (\{k : i_{t+1} \in \delta_k(v_{t+1})\} \cup \{k : \mathbb{P}_k(v_{t+1}) = 0\})$
 - $t := t + 1$
- $v^* := v_t, i^* := i_t$.

By construction we always have $i_{t+1} \in \delta_{k_t}(v_{t+1}) = \delta_{k_t}(v_t w_t)$. Thus, the state k_t is removed from the set \mathcal{R} in each iteration of the loop, so the algorithm must terminate after a finite number of steps. Now, let us assume that the loop terminates at time t with word $v^* = v_t =$

$w_0 w_1 \dots w_{t-1}$ and state $i^* = i_t$. Then, since $i_0 = 1$ and $i_{\tau+1} \in \delta_{i_\tau}(w_\tau)$ for each $\tau = 0, 1, \dots, t - 1$, we know $i^* \in \delta_1(v^*)$, and, hence, that $v^* \in \mathcal{L}(M)$. Further, since each state $k \neq 1$ must be removed from the list \mathcal{R} before the algorithm terminates, we know that for each $k \neq 1$ one of the following must hold:

1. $i_{\tau+1} \in \delta_k(v_{\tau+1})$, for some $0 \leq \tau \leq t - 1$, which implies $i^* \in \delta_k(v^*)$.
2. $\mathbb{P}_k(v_{\tau+1}) = 0$, for some $0 \leq \tau \leq t - 1$, which implies $\mathbb{P}_k(v^*) = 0$.

It follows that v^* is a flag word for state i^* . \square

3.2.2. Convergence of entropy rate approximations

Theorem 2. *Let M be an edge-emitting HMM, and let h and $h(t)$ denote the entropy rate and length- t entropy rate estimate for its output process. If the block model M^n is flag-state, for some n relatively prime to $\text{per}(M)$, then the estimates $h(t)$ converge exponentially with*

$$\limsup_{t \rightarrow \infty} \{h(t) - h\}^{1/t} \leq \alpha^{1/n} \tag{10}$$

where $0 < \alpha < 1$ is the convergence rate, given by [Theorem 1](#), for the entropy estimates of the block model. In particular, by [Lemma 5](#), the entropy estimates $h(t)$ always converge exponentially if M is path-mergeable.

- Remarks.**
1. If M is path-mergeable then [Lemma 5](#) ensures that there will always be some n , relatively prime to $\text{per}(M)$, such that the block model M^n is flag-state. However, to obtain an ideal bound on the speed of convergence of the entropy rate estimates for M , one should generally choose n as small as possible, and the construction given in the lemma is not always ideal for this.
 2. The theorem does not explicitly require the HMM M to be path-mergeable for exponential convergence to hold, only that the block model M^n is flag-state, for some n . However, while it is easy to check if a HMM M is path-mergeable or flag-state (see [Section 5.1](#) for the path-mergeability test), it is generally much more difficult to determine if the block model M^n is flag-state for some unknown n . Indeed, the computational time to do so may be super-exponential in the number of states. And, while asymptotically “almost all” HMMs are path-mergeable in the sense of [Proposition 2](#), “almost none” are themselves flag-state. Thus, the flag-state condition in practice is not as directly applicable as the path-mergeability condition.
 3. Another somewhat more useful condition is incompatibility. We say states i and j are *incompatible* if there exists some length m such that the sets of allowed words of length m ,³ that can be generated from states i and j have no overlap (i.e., for each $w \in \mathcal{L}_m(M)$, either $\mathbb{P}_i(w) = 0$, or $\mathbb{P}_j(w) = 0$, or both).

By a small modification of the construction used in [Lemma 5](#), it is easily seen that if each pair of distinct states i, j of a HMM M is either path-mergeable or incompatible then the block model M^n will be flag-state, for some n . This weaker sufficient condition may sometimes be useful in practice since incompatibility of state pairs, like path-mergeability, may be checked in reasonable computational time (using a similar algorithm).

Proof. The claim (10) is an immediate consequence of [Theorem 1](#) and the following simple lemma. \square

³ Hence, also, the sets of allowed words of all lengths $m' > m$.

Lemma 6. Let M be an edge-emitting HMM with block model M^n , for some n relatively prime to $\text{per}(M)$. Let h and $h(t)$ be the entropy rate and length- t block estimate for the output process of M , and let g and $g(t)$ be the entropy rate and length- t block estimate for the output process of M^n . Then:

- (i) $g = nh$.
- (ii) $\limsup_{t \rightarrow \infty} \{h(t) - h\}^{1/t} = [\limsup_{t \rightarrow \infty} \{g(t) - g\}^{1/t}]^{1/n}$.

Proof. (i) is a direct computation:

$$g = \lim_{t \rightarrow \infty} \frac{H(W_0^{t-1})}{t} = \lim_{t \rightarrow \infty} \frac{H(X_0^{nt-1})}{t} = \lim_{t \rightarrow \infty} n \cdot \frac{H(X_0^{nt-1})}{nt} = nh$$

where X_t and W_t denote, respectively, the t th output symbol of M and M^n . To show (ii), note that by the entropy chain rule (see, e.g., [10])

$$\begin{aligned} g(t+1) - g &= H(W_t | W_0^{t-1}) - nh = H(X_{nt}^{n(t+1)-1} | X_0^{nt-1}) - nh \\ &= \left\{ \sum_{\tau=nt}^{n(t+1)-1} H(X_\tau | X_0^{\tau-1}) \right\} - nh = \sum_{\tau=nt}^{n(t+1)-1} \{h(\tau+1) - h\}. \end{aligned}$$

The claim follows from this relation and the fact that the block estimates $h(t)$ are weakly decreasing (i.e. nonincreasing). \square

4. Relation to state-emitting HMMs

In Section 3 we established exponential convergence of the entropy rate estimates $h(t)$ for path-mergeable, edge-emitting HMMs. The following simple proposition shows that path-mergeability, as well as the flag-state property, both translate directly from state-emitting to edge-emitting HMMs. Thus, exponential convergence of the entropy rate estimates $h(t)$ also holds for path-mergeable, state-emitting HMMs.

Proposition 1. Let $M = (\mathcal{S}, \mathcal{X}, \mathcal{T}, \mathcal{O})$ be a state-emitting HMM, and let $M' = (\mathcal{S}, \mathcal{X}, \{\mathcal{T}'^{(x)}\})$ be the corresponding edge-emitting HMM defined by the conversion algorithm of Section 2.2.2 with $\mathcal{T}'_{ij}{}^{(x)} = \mathcal{T}_{ij}\mathcal{O}_{jx}$. Then:

- (i) M is path-mergeable if and only if M' is path-mergeable.
- (ii) M is a flag-state HMM if and only if M' is a flag-state HMM.

Proof. By induction on t , it is easily seen that for each $t \in \mathbb{N}$, $x_1^t \in \mathcal{X}^t$, and $s_0^t \in \mathcal{S}^{t+1}$

$$\mathbb{P}_i(S_0^t = s_0^t, X_1^t = x_1^t) = \mathbb{P}'_i(S_0^t = s_0^t, X_0^{t-1} = x_1^t)$$

where \mathbb{P}_i and \mathbb{P}'_i are, respectively, the measures on state-symbol sequences $(S_t, X_t)_{t \geq 0}$ for M and M' from initial state $S_0 = i$. This implies that the transition function $\delta_i(w)$ for M is the same as the transition function $\delta'_i(w)$ for M' , i.e. $\delta_i(w) = \delta'_i(w)$ for each $i \in \mathcal{S}$, $w \in \mathcal{X}^*$. Both claims follow immediately from the equivalence of transition functions. \square

5. Typicality and testability of the path-mergeability property

5.1. Testability

The following algorithm to test for path-mergeability of state pairs in an edge-emitting HMM is a small modification of the standard table filling algorithm for deterministic finite automata (DFAs) given in [18]. To test for path-mergeability in a state-emitting HMM one may first convert it to an edge-emitting HMM and then test the edge-emitting HMM for path-mergeability, as the result will be equivalent by Proposition 1.

Algorithm 1. Test for path-mergeable state pairs in an edge-emitting HMM.

1. Initialization step

- Create a table with boxes for each pair of distinct states (i, j) . Initially all boxes are unmarked.
- Then, for each state pair (i, j) , mark the box for pair (i, j) if there is some symbol x with $\delta_i(x) \cap \delta_j(x) \neq \emptyset$.

2. Inductive step

- If the box for pair (i', j') is already marked and $i' \in \delta_i(x)$, $j' \in \delta_j(x)$, for some symbol x , then mark the box for pair (i, j) .
- Repeat until no more new boxes can be marked this way.

By induction on the length of the minimum path-merging word w for a given state pair (i, j) it is easily seen that a state pair (i, j) ends up with a marked box under Algorithm 1 if and only if it is path-mergeable. Thus, the HMM is itself path-mergeable if and only if all state pairs (i, j) end up with marked boxes. This algorithm is also reasonably fast (polynomial time) if the inductions in the second step are carried out in an efficient manner. In particular, a decent encoding of the inductive step gives a maximum run time $O(m \cdot n^4)$, where n is the number of states and m is the number of symbols.⁴

5.2. Typicality

By an *edge-emitting HMM topology* we mean simply the directed, symbol-labeled graph, without assigned probabilities for the allowed transitions. Similarly, by a *state-emitting HMM topology* we mean the directed graph of allowed state transitions along with the sets of allowed output symbols for each state, without the probabilities of allowed transitions or emission probabilities of allowed symbols from each state. By a *labeled, n -state, m -symbol topology* (either edge-emitting or state-emitting) we mean a topology with n states $\{1, \dots, n\}$ and m symbols $\{1, \dots, m\}$, or some subset thereof. That is, we allow that not all the symbols are actually used.

Clearly, the path-mergeability property depends only on the topology of a HMM. The following proposition shows that this property is asymptotically typical in the space of HMM topologies.

⁴ For the standard DFA table filling algorithm an efficient encoding of the inductive step, as described in [18], gives a maximal run time $O(m \cdot n^2)$. However, DFAs, by definition, have deterministic transitions; for each state i and symbol x there is only 1 outgoing edge from state i labeled with symbol x . The increased run time for checking path-mergeability is due to the fact that each state i may have up to n outgoing transitions on any symbol x in a general HMM topology.

Proposition 2. *If a HMM topology is selected uniformly at random from all irreducible, labeled, n -state, m -symbol topologies then it will be path-mergeable with probability $1 - O(\alpha^n)$, uniformly in m , for some fixed constant $0 < \alpha < 1$.*

Remarks. 1. The claim holds for either a randomly selected edge-emitting topology or a randomly selected state-emitting topology, though we will present the proof only in the edge-emitting case as the other case is quite similar.

2. In fact, the proof will show that, with probability $1 - O(\alpha^n)$, a particular strong form of path-mergeability holds. Namely, that each pair of states i, j has a length-1 path-merging word, or path-merging symbol.
3. A similar statement also holds if one does not require the topology to be irreducible, and, in fact, the proof is somewhat simpler in this case. However, because we are interested only in irreducible HMMs where there is a unique stationary state distribution π and well defined entropy rate, we feel it is more appropriate to consider a randomly selected irreducible HMM topology.
4. For fixed $m \in \mathbb{N}$, there are exponentially more n -state, $(m + 1)$ -symbol topologies than n -state, m -symbol topologies, as $n \rightarrow \infty$. Thus, it follows from the proposition that a similar claim holds if one considers topologies with n states and exactly m symbols used, rather than the n -state, m -symbol topologies we consider in which some of the symbols may not be used.
5. We are considering labeled topologies, so that the state labels $1, \dots, n$ and symbol labels $1, \dots, m$ are not interchangeable. That is, two topologies that would be the same under a permutation of the symbol and/or state labels are considered distinct, not equivalent. We believe a similar statement should also hold if one considers a topology permutation-equivalence class chosen uniformly at random from all such equivalence classes with n states and m symbols. However, the proof in this case seems more difficult.

Proof (Edge-Emitting). Assume, at first, that a topology T is selected uniformly at random from all labeled, n -state, m -symbol topologies without requiring irreducibility, or even that each state has any outgoing edges. That is, for each pair of states i, j and symbol x we have, independently, a directed edge from state i to state j labeled with symbol x present with probability $1/2$ and absent with probability $1/2$.

Let $A_{i,x} \equiv \{j : \exists \text{ an edge from } i \text{ to } j \text{ labeled with } x\}$ be the set of states that state i can transition to on symbol x , and let $N_{i,x} \equiv |A_{i,x}|$ be the number of states that state i can transition to on symbol x . Also, define the following events:

$$\begin{aligned}
 E_{1,x} &\equiv \{N_{i,x} > n/4, \text{ for all } i\}, \quad x \in \mathcal{X} \\
 E_{2,x} &\equiv \{A_{i,x} \cap A_{j,x} \neq \emptyset, \text{ for all } i \neq j\}, \quad x \in \mathcal{X} \\
 E_3 &\equiv \{i \rightarrow j, \text{ for all } i, j\} = \{T \text{ is irreducible}\}
 \end{aligned}$$

where $i \rightarrow j$ means there is a path from i to j in the topology (directed graph) T . We note that:

1. By Hoeffding’s inequality,

$$\mathbb{P}(N_{i,x} \leq n/4) = \mathbb{P}(N_{i,x} - \mathbb{E}(N_{i,x}) \leq -n/4) \leq e^{-2n(1/4)^2} = e^{-n/8}.$$

So, for each x , $\mathbb{P}(E_{1,x}^c) \leq ne^{-n/8}$.

2. Since the sets $A_{i,x}, A_{j,x}$ are independent for all $i \neq j$, $\mathbb{P}(A_{i,x} \cap A_{j,x} = \emptyset | E_{1,x}) \leq (3/4)^{n/4}$, for all $i \neq j$. Hence, $\mathbb{P}(E_{2,x}^c | E_{1,x}) \leq \binom{n}{2} (3/4)^{n/4}$, for each x .

3. Since there are at least $\lceil n/4 \rceil - 1$ states in the set $A_{i,x}/\{i\}$ on the event $E_{1,x}$, we have

$$\mathbb{P}(i \not\rightarrow j | E_{1,x}) \leq \mathbb{P}(j \notin A_{k,x}, \text{ for all } k \in A_{i,x}/\{i\} | E_{1,x}) \leq (3/4)^{n/4-1}.$$

Hence, $\mathbb{P}(E_3^c | E_{1,x}) \leq n^2(3/4)^{n/4-1}$, for each x .

Now, the probability that an irreducible, labeled, n -state, m -symbol HMM topology selected uniformly at random is path-mergeable is $\mathbb{P}(T \text{ is path-mergeable} | E_3)$, where \mathbb{P} is, as above, the probability measure on the random topologies T given by selecting each directed symbol-labeled edge to be, independently, present or absent with probability $1/2$. Using points 1, 2, and 3, and fixing any symbol $x \in \mathcal{X}$, we may bound this probability as follows:

$$\begin{aligned} \mathbb{P}(T \text{ is path-mergeable} | E_3) &\geq \mathbb{P}(T \text{ is path-mergeable}, E_3) \\ &\geq \mathbb{P}(T \text{ is path-mergeable}, E_3, E_{1,x}) \\ &= \mathbb{P}(E_{1,x}) \cdot \mathbb{P}(T \text{ is path-mergeable}, E_3 | E_{1,x}) \\ &= \mathbb{P}(E_{1,x}) \cdot (1 - \mathbb{P}(T \text{ is not path-mergeable}, E_3 | E_{1,x}) - \mathbb{P}(E_3^c | E_{1,x})) \\ &\geq \mathbb{P}(E_{1,x}) \cdot (1 - \mathbb{P}(E_{2,x}^c | E_{1,x}) - \mathbb{P}(E_3^c | E_{1,x})) \\ &\geq (1 - ne^{-n/8}) \cdot \left(1 - \binom{n}{2} (3/4)^{n/4} - n^2(3/4)^{n/4-1}\right) \\ &= 1 - O(\alpha^n), \end{aligned}$$

for any $\alpha > (3/4)^{1/4}$. (Note that $(3/4)^{1/4} \approx 0.931 > e^{-1/8} \approx 0.882$.) \square

Acknowledgments

The author thanks Jim Crutchfield for helpful discussions. This work was partially supported by ARO grant W911NF-12-1-0234 and VIGRE grant DMS0636297.

References

- [1] R. Atar, O. Zeitouni, Lyapunov exponents for finite state nonlinear filtering, *SIAM J. Control Optim.* 35 (1995) 36–55.
- [2] R. Atar, O. Zeitouni, Exponential stability for nonlinear filtering, *Ann. Inst. H. Poincaré Probab. Statist.* 33 (6) (1997) 697–725.
- [3] L.R. Bahl, F. Jelinek, R.L. Mercer, A maximum likelihood approach to continuous speech recognition, *IEEE Trans. Pattern Anal. Mach. Intell. PAMI-5* (1983) 179–190.
- [4] P. Baldi, Y. Chauvin, T. Hunkapiller, M.A. McClure, Hidden Markov models of biological primary sequence information, *PNAS* 91 (1994) 1059–1063.
- [5] J. Birch, Approximations for the entropy for functions of Markov chains, *Ann. Math. Statist.* 33 (3) (1962) 930–938.
- [6] D. Blackwell, The entropy of functions of finite-state Markov chains, in: *Transactions of the First Prague Conference on Information Theory, Statistical Decision Functions, Random Processes*, Publishing House of the Czechoslovak Academy of Sciences, 1957, pp. 13–20.
- [7] D. Blackwell, L. Koopmans, On the identifiability problem for functions of finite Markov chains, *Ann. Math. Statist.* 28 (4) (1957) 1011–1015.
- [8] P. Chigansky, R. Lipster, Stability of nonlinear filters in nonmixing case, *Ann. Appl. Probab.* 14 (4) (2004) 2038–2056.
- [9] P. Collet, F. Leonardi, Loss of memory of hidden Markov models and Lyapunov exponents, 2009. arXiv/0908.0077.
- [10] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, second ed., Wiley-Interscience, Hoboken, NJ, 2006.
- [11] J.P. Crutchfield, D.P. Feldman, Regularities unseen, randomness observed: levels of entropy convergence, *Chaos* 13 (1) (2003) 25–54.
- [12] R. Douc, G. Fort, E. Moulines, P. Priouret, Forgetting the initial distribution for hidden Markov models, *Stoch. Process. Appl.* 119 (4) (2009) 1235–1256.
- [13] S.R. Eddy, Profile hidden Markov models, *Bioinformatics* 14 (9) (1998) 755–763.

- [14] S.R. Eddy, G. Mitchison, R. Durbin, Maximum discrimination hidden Markov models of sequence consensus, *J. Comput. Biol.* 2 (1) (1995) 9–23.
- [15] E.J. Gilbert, On the identifiability problem for functions of finite Markov chains, *Ann. Math. Statist.* 30 (3) (1959) 688–697.
- [16] G. Han, B. Marcus, Analyticity of entropy rate of hidden Markov chains, *IEEE Trans. Inform. Theory* 52 (12) (2006) 5251–5266.
- [17] B.M. Hochwald, P.R. Jelenkovic, State learning and mixing in entropy of hidden Markov processes and the Gilbert–Elliott channel, *IEEE Trans. Inform. Theory* 45 (1) (1999) 128–138.
- [18] J.E. Hopcroft, R. Motwani, J.D. Ullman, *Introduction to Automata Theory, Languages, and Computation*, second ed., Addison-Wesley, 2001.
- [19] F. Jelinek, Continuous speech recognition by statistical methods, *Proc. IEEE* 64 (1976) 532–536.
- [20] B.H. Juang, L.R. Rabiner, Hidden Markov models for speech recognition, *Technometrics* 33 (3) (1991) 251–272.
- [21] K. Karplus, C. Barrett, R. Hughey, Hidden Markov models for detecting remote protein homologies, *Bioinformatics* 14 (10) (1998) 846–856.
- [22] F. Le Gland, L. Mevel, Exponential forgetting and geometric ergodicity in hidden Markov models, *Math. Control Signals Systems* 13 (2000) 63–93.
- [23] D. Levin, Y. Peres, E.L. Wilmer, *Markov Chains and Mixing Times*, American Mathematical Society, Providence, RI, 2006.
- [24] H. Pfister, *On the capacity of finite state channels and analysis of convolutional accumulate- m codes* (Ph.D. thesis), University of California, San Diego, 2003.
- [25] L.R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, *IEEE Proc.* 77 (1989) 257–286.
- [26] A. Siepel, D. Haussler, Combining phylogenetic and hidden Markov models in biosequence analysis, *J. Comput. Biol.* 11 (2–3) (2004) 413–428.
- [27] R.B. Sowers, A.M. Makowski, Discrete-time filtering for linear systems in correlated noise with non-gaussian initial conditions: formulas and asymptotics, *IEEE Trans. Automat. Control* 37 (1992) 114–121.
- [28] N.F. Travers, J.P. Crutchfield, Asymptotic synchronization for finite-state sources, *J. Stat. Phys.* 145 (5) (2011) 1202–1223.
- [29] N.F. Travers, J.P. Crutchfield, Exact synchronization for finite-state sources, *J. Stat. Phys.* 145 (5) (2011) 1181–1201.