

Memory Complexity of Estimating Entropy and Mutual Information

Tomer Berg[✉], Or Ordentlich[✉], and Ofer Shayevitz[✉], *Senior Member, IEEE*

Abstract—We observe an infinite sequence of independent identically distributed random variables X_1, X_2, \dots drawn from an unknown distribution p over $[n]$, and our goal is to estimate the entropy $H(p) = -\mathbb{E}[\log p(X)]$ within an ε -additive error. To that end, at each time point we are allowed to update a finite-state machine with S states, using a possibly randomized but time-invariant rule, where each state of the machine is assigned an entropy estimate. Our goal is to characterize the minimax memory complexity S^* of this problem, which is the minimal number of states for which the estimation task is feasible with probability at least $1 - \delta$ asymptotically, uniformly in p . Specifically, we show that there exist universal constants C_1 and C_2 such that $S^* \leq C_1 \cdot \frac{n(\log n)^4}{\varepsilon^2 \delta}$ for ε not too small, and $S^* \geq C_2 \cdot \max\{n, \frac{\log n}{\varepsilon}\}$ for ε not too large. The upper bound is proved using approximate counting to estimate the logarithm of p , and a finite memory bias estimation machine to estimate the expectation operation. The lower bound is proved via a reduction of entropy estimation to uniformity testing. We also apply these results to derive bounds on the memory complexity of mutual information estimation.

Index Terms—Memory complexity, sample complexity, entropy estimation, mutual information estimation, finite memory algorithms.

I. INTRODUCTION

THE problem of inferring properties of an underlying distribution given sample access is called *statistical property estimation*. A typical setup is as follows: given independent samples X_1, \dots, X_n from an unknown distribution p , the objective is to estimate a property $g(p)$ (e.g., entropy, support size, L_p norm, etc.) under some resource limitation. A prominent example of such a limitation is the amount of available samples, and this limitation gives rise to the notion of sample complexity, namely the minimal number of samples one needs to see in order to estimate $g(p)$ with some given accuracy. Many real-world machine learning and data analysis tasks are limited by insufficient samples, and the challenge of inferring properties of a distribution given a small sample

size is encountered in a variety of settings, including text data, customer data, and the study of genetic mutations across a population. The sample complexity of property estimation and, specifically, of entropy estimation, have therefore received much attention in the literature (see Section II for details).

However, in many contemporary settings, collecting enough samples for accurate estimation is less of a problem, and the bottleneck shifts to the computational resources available for the task and, in particular, the available memory size. In this work, we therefore focus on the problem of estimation under memory constraints, and, in particular, entropy estimation. In order to isolate the effect that finite memory has on the fundamental limits of the problem, we let the number of samples we process be arbitrarily large.

Formally, the problem is defined as follows. Let Δ_n be the collection of all distributions over $[n]$. The Shannon entropy of $p \in \Delta_n$ is $H(p) = -\sum_{x \in [n]} p(x) \log p(x)$. Given independent samples X_1, X_2, \dots from an unknown $p \in \Delta_n$, we would like to accurately estimate $H(p)$ using limited memory. To that end, an *S-state entropy estimator* is a finite-state machine with S states, defined by two functions: The (possibly randomized) memory update function $f : [S] \times [n] \rightarrow [S]$, describing the transition between states as a function of an input sample, and the entropy estimate function $\hat{H} : [S] \rightarrow [0, \log n]$, assigning an entropy estimate to each state. Letting M_t denote the state of the memory at time t , this finite-state machine evolves according to the rule:

$$M_0 = s_{\text{init}}, \quad (1)$$

$$M_t = f(M_{t-1}, X_t) \in [S], \quad (2)$$

for some predetermined initial state $s_{\text{init}} \in [S]$. If the machine is stopped at time t , it outputs the estimation $\hat{H}(M_t)$. We restrict the discussion to time-invariant memory update function f , since storing the time index necessarily incurs a memory cost, and, furthermore, since the number of samples is unbounded, simply storing the code generating a time-varying algorithm may require unbounded memory. We say that an ε -error occurred at time t if our estimate $\hat{H}(M_t)$ is ε -far from the correct entropy. Our figure of merit for the estimator is taken to be its worst-case asymptotic ε -error probability:

$$P_e(f, \hat{H}, \varepsilon) = \sup_{p \in \Delta_n} \limsup_{t \rightarrow \infty} \Pr(|\hat{H}(M_t) - H(p)| > \varepsilon). \quad (3)$$

We are interested in the *minimax memory complexity* $S^*(n, \varepsilon, \delta)$, defined as the smallest integer S for which there exist (f, \hat{H}) such that $P_e(f, \hat{H}, \varepsilon) \leq \delta$.

Received 10 June 2024; revised 21 February 2025; accepted 21 February 2025. Date of publication 4 March 2025; date of current version 23 April 2025. This work was supported by the Israel Science Foundation (ISF) under Grant 1641/21 and Grant 1766/22. (Corresponding author: Tomer Berg.)

Tomer Berg and Ofer Shayevitz are with the Department of Electrical Engineering—Systems, Tel Aviv University, Tel Aviv-Yafo 69978, Israel (e-mail: tomerberg@mail.tau.ac.il; ofersha@eng.tau.ac.il).

Or Ordentlich is with the School of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem 91904, Israel (e-mail: or.ordentlich@mail.huji.ac.il).

Communicated by J. Chen, Associate Editor for Shannon Theory and Information Measures.

Digital Object Identifier 10.1109/TIT.2025.3547871

0018-9448 © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

Authorized licensed use limited to: Hebrew University. Downloaded on June 24, 2025 at 19:27:42 UTC from IEEE Xplore. Restrictions apply.

Our main result is an upper bound on $S^*(n, \varepsilon, \delta)$, which shows that $\log \frac{n}{\varepsilon^2} + o(\log n)$ bits suffice for entropy estimation when $\varepsilon > 10^{-5}$, thus improving upon the best known upper bounds thus far ([1], [2]). While our focus here is on minimizing the memory complexity of the problem in the limit of infinite number of available samples, we further show that the estimation algorithm attaining this memory complexity upper bound only requires $\tilde{O}(n^c)$ samples, for any $c > 1$.¹ Thus, in entropy estimation one can achieve almost optimal sample complexity and memory complexity, simultaneously. Our proposed algorithm approximates the logarithm of $p(x)$, for a given $x \in [n]$, using a *Morris counter* [3]. The inherent structure of the Morris counter is particularly suited for constructing a nearly-unbiased estimator for $\log p(x)$, making it a natural choice for memory efficient entropy estimation. In order to compute the mean of these estimators, $\mathbb{E}[\log \widehat{p}(X)]$, in a memory efficient manner, a finite-memory bias estimation machine (e.g., [4], [5]) is leveraged for simulating the expectation operator. The performance of a scheme based on this high-level idea is analyzed, and yields the following upper bound on the memory complexity:

Theorem 1: For any $c > 1$, $\beta > 0$, $0 < \delta < 1$ and $\varepsilon = 10^{-5} + \beta + \psi_c(n)$, we have

$$S^*(n, \varepsilon, \delta) \leq n \left(\frac{8(c \log n + 2)^4}{\beta^2 \delta} + 4(c \log n + 2)^2 \right), \quad (4)$$

where

$$\begin{aligned} \psi_c(n) &= (e + 1)n^{-(c-1)+v_n(1)} \\ &+ \min \left\{ 1, Cn^{-\frac{c-1}{2}+v_n(1/2)} \right\} + \frac{100(c \log n + 2)}{n^c(1 - 0.5n^{-c})^2} \\ &= O \left(2\sqrt{\log n} \cdot n^{-\frac{c-1}{2}} \right), \end{aligned} \quad (5)$$

and we set $C = 2(e + 1)10^8$ and $v_n(\alpha) \triangleq \sqrt{\frac{2c\alpha^3}{\log n}} + \frac{\alpha}{\log n}$.

Moreover, there is an algorithm that attains (4) when the number of samples is $\Omega \left(\frac{n^c \cdot \text{poly}(\log n)}{\delta} \cdot \text{poly}(\log(1/\delta)) \right)$, and returns an estimation of $H(p)$ within an ε -additive error with probability at least $1 - 3\delta$.

Note that the additive term $\psi_c(n)$ only becomes negligible when 10^8 is much smaller than $n^{-\frac{c-1}{2}}$, thus the regime in which our results are significant is the asymptotic regime. Furthermore, while $\psi_c(n)$ vanishes for large n , our bound is always limited to $\varepsilon > 10^{-5}$. This small bias is due to inherent properties of the Morris counter, on which we elaborate in Section III. As we are more interested in the case where the entropy grows with the alphabet size n , the limitation of the attainable additive error to values above 10^{-5} is typically a very moderate one for the sizes of n we consider. While attaining good sample complexity is not the main focus of our work, we also note that if n is large and ε not too small, one can choose c arbitrarily close to 1, resulting in an algorithm whose sample complexity has similar dependence on n as those of the limited-memory entropy estimation algorithms proposed in [1] and [2], while requiring less memory states. This result might be of practical interest for applications in which memory is a scarcer resource than samples, e.g., a limited

memory high-speed router that leverages entropy estimation to monitor IP network traffic [6]. Finally, we note that initial simulation results are more optimistic than the prediction of the theorem. Specifically, for a uniformly distributed input over $n = 1000$ with parameters $c = 1.5, \beta = 0.1, \delta = 0.1$, the sample complexity of the algorithm for these parameters, as prescribed by lemma 12 is $L \approx 4 \cdot 10^{14}$. However, by running the entropy machine for $t = 10^{11}$ samples, we were able to obtain an additive error of about 0.15, much smaller than the infinite sample additive error predicted by Theorem 1, which is approximately $\beta + \min\{1, C \cdot n^{-1/2}\} + n^{-c} \cdot 100(c \log n + 2) \approx 1.18$. It seems that $\psi_c(n)$, while negligible for large n , is still marginally loose.

Furthermore, we derive two lower bounds on the memory complexity. The first lower bound shows that when $H(p)$ is close to $\log n$, the memory complexity cannot be too small. This bound is obtained via a reduction of entropy estimation to uniformity testing, by noting that thresholding the output of a good entropy estimation machine around $\log n$ can be used to decide whether p is close to the uniform distribution or not. The bound then follows from the $\Omega(n)$ lower bound of [7] on uniformity testing. The second lower bound follows from the observation that, if the number of states is too small, there must be some value of the entropy at distance greater than ε from all estimate value hence, for this value of the entropy, our entropy estimator will be ε -far from the real value with probability 1. Combining these lower bounds yields the following.

Theorem 2: For any $\varepsilon > 0$, we have

$$S^*(n, \varepsilon, \delta) \geq \frac{\log n}{2\varepsilon}. \quad (6)$$

Furthermore, if $\varepsilon < \frac{1}{4 \ln 2}$, then

$$S^*(n, \varepsilon, \delta) \geq n(1 - 2\sqrt{\varepsilon \ln 2}). \quad (7)$$

One of several open problems posed by the authors of [1] is to prove a lower bound on the space requirement of a sample optimal algorithm for entropy estimation. Theorem 2 answers this question by giving a lower bound on the memory size needed when the number of samples is infinite, which clearly also holds for any finite number of samples. In the concluding section of the paper, we extend our results to the mutual information estimation problem. Let $(X, Y) \sim p_{XY}$, where p_{XY} is an unknown distribution over $[n] \times [m]$ such that the marginal distribution of X is p_X and the marginal distribution of Y is p_Y . The mutual information between X and Y is given as $I(X; Y) = H(X) + H(Y) - H(X, Y)$. We derive the following bounds on the memory complexity of mutual information estimation, namely the minimal number of states needed to estimate $I(X; Y)$ with additive error at most ε with probability of at least $1 - \delta$, which we denote as $S_{\text{MI}}^*(n, m, \varepsilon, \delta)$.

Theorem 3: For any $c > 1$, $\beta > 0$ and $\varepsilon = 3 \cdot 10^{-5} + \beta + O \left(\min \left\{ 2\sqrt{\log n} \cdot n^{-\frac{1}{2}(c-1)}, 2\sqrt{\log m} \cdot m^{-\frac{1}{2}(c-1)} \right\} \right)$,

$$\begin{aligned} S_{\text{MI}}^*(n, m, \varepsilon, \delta) &\leq nm \left(\frac{288 \cdot (c \log nm + 2)^6}{\beta^2 \delta} \right. \\ &\quad \left. + 16(c \log nm + 2)^4 \right). \end{aligned} \quad (8)$$

¹The \tilde{O} suppresses poly-logarithmic terms.

For $\varepsilon < \frac{1}{12 \ln 2}$, and if $\frac{n}{\log^3 n} = \Omega(\log^7 m)$ and $\frac{m}{\log^3 m} = \Omega(\log^7 n)$ both hold, then

$$S_{\text{MI}}^*(n, m, \varepsilon, \delta) = \Omega\left(\frac{n \cdot m}{\log^3 n \cdot \log^3 m}\right). \quad (9)$$

II. RELATED WORK

The study of estimation under memory constraints has received far less attention than the sample complexity of statistical estimation. References [8] and [9] studied this setting for hypothesis testing with finite memory, and [10], [4] have studied estimating the bias of a coin using a finite state machine. It has then been largely abandoned, but recently there has been a revived interest in space-sample trade-offs in statistical estimation, and many works have addressed different aspects of the learning under memory constraints problem over the last few years. See, e.g., [11], [12], [13], [14], [15], [16], [17], [18], and [19] for a non exhaustive list of recent works.

The problem of estimating the entropy with limited independent samples from the distribution has a long history. It was originally addressed by [20], who suggested the simple and natural empirical plug-in estimator. This estimator outputs the entropy of the empirical distribution of the samples, and its sample complexity [21] is $\Theta\left(\frac{n}{\varepsilon} + \frac{\log^2 n}{\varepsilon^2}\right)$. Reference [21] showed that the plug-in estimator is always consistent, and the resulting sample complexity was shown to be linear in n . In the last two decades, many efforts were made to improve the bounds on the sample complexity. Paninski [22] and [23] was the first to prove that it is possible to consistently estimate the entropy using sublinear sample size. While the scaling of the minimal sample size of consistent estimation was shown to be $\frac{n}{\log n}$ in the seminal results of [24], [25], the optimal dependence of the sample size on both n and ε was not completely resolved until recently. In particular, $\Omega\left(\frac{n}{\varepsilon \log n}\right)$ samples were shown to be necessary, and the best upper bound on the sample complexity was relied on an estimator based on linear programming that can achieve an additive error ε using $O\left(\frac{n}{\varepsilon^2 \log n}\right)$ samples [26]. This gap was partially amended in [27] by a different estimator, which requires $O\left(\frac{n}{\varepsilon \log n}\right)$ samples but is only valid when ε is not too small. The sharp sample complexity was shown by [28] and [29] to indeed be

$$\Theta\left(\frac{n}{\varepsilon \log n} + \frac{\log^2 n}{\varepsilon^2}\right). \quad (10)$$

The space-complexity (which is the minimal memory in bits needed for the algorithm) of estimating the entropy of the empirical distribution of the data stream is well-studied for worst-case data streams of a given length, see [30], [6], and [31]. Reference [32] addressed the problem of deciding if the entropy of a distribution is above or beyond than some predefined threshold, using algorithms with limited memory. The trade-off between sample complexity and space/communication complexity for the entropy estimation of a distribution is the subject of a more recent line of work. The earliest work on the subject is [1], where the authors constructed an algorithm which is guaranteed to work with $O(n/\varepsilon^3 \cdot \text{polylog}(1/\varepsilon))$ samples and any memory size

$b \geq 20 \log\left(\frac{n}{\varepsilon}\right)$ bits (which corresponds to $O(n^{20}/\varepsilon^{20})$ memory states in our setup). Their upper bound on the sample complexity was later improved by [2] to $O(n/\varepsilon^2 \cdot \text{polylog}(1/\varepsilon))$ with space complexity of $O(\log\left(\frac{n}{\varepsilon}\right))$ bits. We note that in both works above the constant in the space complexity upper bound can be reduced from 20 to 5 by a careful analysis. The work of [1] is based on an empirical estimator. Consider the following simple approach: we draw N samples from the distribution to approximate \hat{p}_x and then take the average of $\log(1/\hat{p}_x)$ over R iterations. This approach gives the desired memory bound, but uses too many samples. To improve the sample complexity, the authors suggest to partition $[0, 1]$ into T disjoint intervals, and perform the simple approach above separately for probabilities within each interval. The essence of the algorithm is that when p_x is large inside the interval, fewer samples are needed to estimate p_x (small N), and if p_x is small inside the interval, fewer iterations are needed (small R). The algorithm of [2] is based on the observation that Y , the number of additional draws needed to see x exactly t more times (where t is an algorithm parameter) is a negative binomial random variable, $Y \sim \text{NB}(t, p_x)$. As $\mathbb{E}(Y) = t/p_x$, taking $\log(Y/t)$ as the estimate of $\log(1/p_x)$, and adding a few more samples to correct the bias, allows the authors to improve the sample complexity by a factor of $1/\varepsilon$. Both of the approaches above can be referred to as “natural” approaches, i.e., realizing algorithms that work well for the unconstrained setup in a memory limited framework. It has been pointed out recently in [33] that this approach for obtaining upper bounds in a memory limited scheme can be strictly suboptimal in the large sample regime. As an example, consider the natural statistic for estimating the parameter of a coin, which is counting the number of 1’s in a stream. This results in a quadratic risk of $O(1/\sqrt{S})$, as in order to count the number of 1’s in a stream of length k we must keep a clock that counts to k , thus overall the number of states used is $S = O(k^2)$. However, it is known from the works of [4] and [5] that the best achievable quadratic risk is $O(1/S)$, and it can be achieved by randomized or deterministic constructions. Unlike the works of [1] and [2], where the authors try to estimate p_x by drawing some related r.v. (Binomial or Negative-Binomial) and then taking its normalized logarithm, the algorithm we propose directly estimates $\log p_x$ by leveraging properties of Morris counters. This allows our algorithm to save memory, yet it comes at a price of an added periodic term, inherent to Morris counters, that can be only bounded by 10^{-5} . Whether this term is a real bottleneck or an artifact of our algorithm is a topic for further research.

III. PRELIMINARIES

In this section, we introduce mathematical notations and some relevant background for the paper.

A. Notation

We write $[n]$ to denote the set $\{1, \dots, n\}$, and consider discrete distributions over $[n]$. We use the notation p_i to denote the probability of element i in distribution p . When X is a random variable on $[n]$, p_X denotes the random variable




Fig. 1. The original Morris counter.

obtained by evaluating p in location X . The entropy of p is defined as $H(p) = -\sum_{x \in [n]} p_x \log p_x = \mathbb{E}_{X \sim p}(-\log p_X)$, where $H(p) = 0$ for a single mass distribution and $H(p) = \log n$ a uniform distribution over $[n]$. The total variation distance between distributions p and q is defined as half their ℓ^1 distance, i.e., $d_{TV}(p, q) = \frac{1}{2} \|p - q\|_1 = \frac{1}{2} \sum_{i=1}^n |p_i - q_i|$, and their KL (Kullback-Leibler) divergence is defined as $D_{KL}(p \| q) = \sum_{i=1}^n p_i \log \frac{p_i}{q_i}$. Logarithms are taken to base 2.

B. Morris Counter

Suppose one wishes to implement a counter that counts up to m . Maintaining this counter exactly can be accomplished using $\log m$ bits. In the first example of a non-trivial streaming algorithm, Morris gave a randomized “approximate counter”, which allows one to retrieve a constant multiplicative approximation to m with high probability using only $O(\log \log m)$ bits (see [3]). The Morris Counter was later analyzed in more detail by Flajolet [34], who showed that $O(\log \log m + \log(1/\epsilon) + \log(1/\delta))$ bits of memory are sufficient to return a $(1 \pm \epsilon)$ approximation with success probability $1 - \delta$. A recent result of [35] shows that $O(\log \log m + \log(1/\epsilon) + \log \log(1/\delta))$ bits suffice for the same task.

The original Morris counter is a random state machine with the following simple structure (described in the algorithm below): At each state $s = 1, 2, 3, \dots$, an increment causes the counter to transition to state $s + 1$ with probability 2^{-s} , and to remain in state s with probability $1 - 2^{-s}$.

Algorithm 1 IncrementMorrisCounter

Input: Previous memory state s

Output: Next memory state s

- 1: $B \leftarrow \text{Draw from a Bern}(2^{-s})$ distribution
 - 2: $s \leftarrow s + B$
-

This is formally the discrete time pure birth process of Figure 1:

The performance of the above counter was characterized by Flajolet, who proved the following theorem.

Theorem 4 [34]: Let C_m be the value of the Morris counter after m increments. It holds that

$$\mathbb{E}(C_m) = \log m + \mu + g(\log m) + \phi(m), \quad (11)$$

where

$$\mu = \frac{\gamma}{\ln 2} + \frac{1}{2} - \sum_{i=1}^{\infty} \frac{1}{2^i - 1}, \quad (12)$$

$$\gamma = \lim_{n \rightarrow \infty} \left(-\log n + \sum_{k=1}^n \frac{1}{k} \right) \text{ (Euler's constant)}, \quad (13)$$

$g(\cdot)$ is a periodic function of amplitude less than 10^{-5} , $|\phi(m)| \leq \min \left\{ 1, \frac{2^{\sqrt{16 \log m \cdot (\log m)^{4.5}}}}{2m} \right\}$ and the expectation is over the randomness of the counter.²

In his paper, Flajolet approximated $\mathbb{E}(C_m)$ with the Mellin integral transform of some function ϕ related to the marginal distribution of the counter, and then used Cauchy's residue theorem in order to compute the integral. The constant μ arises from the residual of the function at 0, where it has a double pole. Thus, if we are interested in approximating $\log m$ using the counter, then using $C_m - \mu$ as our approximation guarantees that on average our additive error will not be more than $10^{-5} + \phi(m)$, a property that we leverage in our entropy estimation algorithm.

C. Finite-State Bias Estimation Machine

In the bias estimation problem, we are given access to i.i.d. samples drawn from the $\text{Bern}(p)$ distribution, and we wish to estimate the value of p under the expected quadratic loss (also known as mean squared error distortion measure). The S -state randomized bias estimation algorithm presented below was purposed by [10], and the Markov chain induced by algorithm is described in Figure 2.

Algorithm 2 IncrementBiasEstimation

Input: Number of states S , previous memory state s , a sample $X \sim \text{Bern}(p)$

Output: Next memory state s , parameter estimate \hat{p}

- 1: $B \leftarrow \text{Draw from a Bern}(\frac{s-1}{S-1})$ distribution
 - 2: **if** $X = 1$ **then**
 - 3: $s \leftarrow s + \bar{B}$
 - 4: **else**
 - 5: $s \leftarrow s - B$
 - 6: **end if**
 - 7: $\hat{p} \leftarrow \frac{s-1}{S-1}$
-

The performance of this algorithm was carefully analyzed by [4] where it was shown, using the Markov chain tree theorem, to asymptotically induce a $\text{Binomial}(S - 1, \theta)$ stationary distribution on the memory state space, which implies that $\mathbb{E}(\hat{p} - p)^2 \leq O(1/S)$. In the same paper, it was further showed that the machine is order-optimal, by proving a lower bound of $\mathbb{E}(\hat{p} - p)^2 \geq \Omega(1/S)$ for any finite-state estimator. For completeness, we provide a simple proof for the MSE achieved by this construction in the appendix.

Lemma 1: Let $\hat{p}(k) = \frac{k-1}{S-1}$ be the estimate of p given state k in the bias estimation machine of Figure 2. Then we have

$$\text{MSE} = \lim_{t \rightarrow \infty} \mathbb{E}(\hat{p}(M_t) - p)^2 \leq \frac{1}{S-1}. \quad (14)$$

IV. UPPER BOUND - ENTROPY ESTIMATION ALGORITHM

In this section we prove Theorem 1, that is, we show the existence of an S -state randomized entropy estimation machine

²In [34], Flajolet bounded $\phi(m)$ with $O(m^{-0.98})$. Here, we carefully follow the constants in his derivation and provide an explicit upper bound on the error terms, since we are interested in bounds that can be applied for finite m .




Fig. 2. Bias estimation machine of [10] ($q = 1 - p$).

with $S \leq \frac{(c+1)^4 n (\log n)^4}{\beta^2 \delta}$ states that achieves additive error ε of at most $10^{-5} + \beta$. The basic idea is to let nature draw some X from p and use a Morris counter to approximate $-\log p_X$, then, since we are looking for $H(p) = \mathbb{E}(-\log p_X)$, use a bias estimation machine to simulate the averaging operation, by randomly generating coin tosses with bias that is proportionate to our estimate of $-\log p_X$. The bias estimation machine is incremented whenever a count is concluded in a randomized clock, which is simulated by another Morris counter. For a sufficiently large number of samples, this averaging converges (approximately) to the mean of $-\log p_X$, and thus outputs an approximation to the true underlying entropy. We divide our presentation to four parts: in the first part we describe the algorithm; in the second part we count the total number of states used by the algorithm; in the third part we assume the bias estimation machine is fed with an infinite number of i.i.d. samples and analyze the performance of the algorithm; and in the fourth part we relax this assumption by studying the mixing time of the Markovian process induced by our bias estimation machine. This allows us to prove an upper bound on the number of samples the developed algorithm requires.

A. Description of the Algorithm

- 1) The algorithm receives an *accuracy parameter* $\beta > 0$ and an *overhead parameter* $c > 1$.
- 2) In each iteration of the algorithm we collect a fresh sample $X \sim p$, and store its value. Assuming the received sample is x , we proceed to estimate $\log p_x$ based on more fresh samples using Morris counters.
- 3) We use two Morris counters - one that approximates a clock, and one that approximates a count for x values.
 - The first counter is randomly incremented whenever a new sample is observed, and it stops when it reaches state $M = B + 1$ states, where B is the smallest integer k such that $\lceil n^c \rceil \leq 2^k$. We denote the state of this counter as C_N . Let N denote the *random* time it takes the counter to arrive at state M . We will show in the sequel that N is expected to be around $\lceil n^c \rceil$ (up to small factors), thus this counter essentially approximates a clock that counts until $\lceil n^c \rceil$ samples are observed.
 - The second counter is randomly incremented whenever x is observed, and it stops when the first counter reaches state M . We denote the state of this counter as C_{N_x} . We allow this counter to have $2M$ states to make sure the

Algorithm 3 Entropy Estimation With Morris Counters

Input: A data stream $X_1, X_2 \dots \sim p$, alphabet size n , run time t , error probability $\delta, \beta > 0, c > 1$, constant μ of (12)

Output: Entropy estimate \hat{H}

1: Set

$B \leftarrow \min\{k \in \mathbb{N} : \lceil n^c \rceil \leq 2^k\}$, $M \leftarrow B + 1$
 $\eta \leftarrow$ Monte Carlo estimate of $\mathbb{E}(\log N)$ for

$$N = \sum_{k=1}^{M-1} \tau_k, \text{ where } \tau_k \sim \text{Geo}(2^{-k})$$

$$a \leftarrow 1 - \frac{\mu + \eta}{M}, \quad S_{\text{bias}} \leftarrow \left\lceil \frac{4M^2}{\beta^2 \delta} \right\rceil + 1$$

$$C_N \leftarrow 1, \quad C_{N_x} \leftarrow 1, \quad s \leftarrow 1$$

```

2: for  $i = 1, \dots, t$  do
3:   if  $C_N = 1$  then
4:      $x_{\text{test}} \leftarrow X_i$ 
5:      $C_{N_x} \leftarrow 1$ 
6:   else
7:      $C_N = \text{IncrementMorrisCounter}(C_N)$ 
8:     if  $X_i = x_{\text{test}}$  then
9:        $C_{N_x} \leftarrow \text{IncrementMorrisCounter}(C_{N_x})$ 
10:    end if
11:    if  $C_{N_x} < 2M$  then
12:      if  $C_N = M$  then
13:         $\theta_{N_x} \leftarrow a - \frac{C_{N_x} - (\mu + \eta)}{2M}$ 
14:         $(s, \hat{\theta}) \leftarrow \text{IncrementBiasEstimation}(S_{\text{bias}}, s, \theta_{N_x})$ 
15:         $C_N \leftarrow 1$ 
16:      end if
17:    else
18:       $C_N \leftarrow 1$ 
19:    end if
20:  end if
21: end for
22:  $\hat{H} \leftarrow 2M(\hat{\theta} - a)$ 

```

probability it ends before the first counter is sufficiently small. In the event that the counter indeed reaches state $2M$ before the first counter, we draw a fresh sample and initialize both counters. This counter approximates the logarithm of the number of observed x values in the length N window.

- 4) Denoting the number of observed x values in the previous stage as N_x , we define $\bar{C}_{N_x} \triangleq C_{N_x} - \mu - \mathbb{E} \log N$ to be the centralized output of the second counter. As $\mathbb{E}(\log N)$ is only a function of n , it can be calculated offline using Monte-Carlo simulation with the desired resolution (see Appendix). As we argue below, this is an almost unbiased estimator for $-\log p_x$.
- 5) We now increment a bias estimation machine with $S_{\text{bias}} = \left\lceil \frac{4M^2}{\beta^2 \delta} \right\rceil + 1$ states whose purpose is to simulate the expectation operation. Specifically, each time the first Morris counter concludes a count, we generate a $\text{Ber}(\theta_{N_x})$ random variable, with $\theta_{N_x} = a - \frac{\bar{C}_{N_x}}{2M}$, and use it as the input to our bias estimation machine. The offset

$a \triangleq 1 - \frac{\mathbb{E}(\log N) + \mu}{2M}$ guarantees that $\theta_{N_x} \in [0, 1)$ with probability 1, as $\theta_{N_x} = 1 - \frac{C_{N_x}}{2M}$ and $1 \leq C_{N_x} \leq 2M$ as it is the output of a Morris counter restricted to $2M$ states. Our estimator for the entropy \hat{H} is the bias estimate of the machine, after subtraction of the known offset a and multiplication by $2M$, that is, $\hat{H} = 2M(\hat{\theta} - a)$.

B. Number of States in Our Machine

As $n, t, \beta, \delta, M, \eta, a$ and S_{bias} are program constants, we do not count them in the memory consumption of the algorithm. At each time point, our algorithm keeps the value of x , the state of the Morris counter approximating the clock, the state of the Morris counter approximating the logarithm of the x counter, and the state of the bias estimation machine. Thus, the total number of states is the product of the individual number of states needed at each step, and recalling that $M = B + 1 \leq c \log n + 2$, the total number of states is

$$\begin{aligned} S &= n \cdot M \cdot 2M \cdot S_{\text{bias}} \\ &\leq n \cdot 2M^2 \cdot \left(\frac{4M^2}{\beta^2 \delta} + 2 \right) \\ &= n \left(\frac{8(c \log n + 2)^4}{\beta^2 \delta} + 4(c \log n + 2)^2 \right). \end{aligned} \quad (15)$$

C. Analysis of the Algorithm for $t = \infty$

Let X be the fresh sample collected at the start of an algorithm iteration. Our analysis is based on characterizing the joint distribution of (N, N_x, C_{N_x}) , where N is the number of observed samples until C_N arrives at state M , N_x is the number of times X appeared in the N sample window, and C_{N_x} is the final state of the second counter. We first bound the probability that N diverges significantly from n^c in Lemma 2, and then apply the result to upper bound $\mathbb{E}(N^{-\alpha})$ in Lemma 3 for any $\alpha \in (0, 1]$. We proceed to bound the difference between the expectation of $\bar{C}_{N_x}^\infty = C_{N_x}^\infty - \mu - \mathbb{E} \log N$ and $-H(p)$ in Lemma 4, where $C_{N_x}^\infty$ denotes the state of an *infinite* memory Morris counter. We show in Lemma 7 that the expected difference between C_{N_x} and $C_{N_x}^\infty$ is $O((\log n)/n^c)$, thus proving that the absolute difference between $\mathbb{E}(\bar{C}_{N_x})$ and the entropy is at most $\psi_c(n)$. Lemma 8 proves that the input to the bias estimation machine is an i.i.d. sequence of Bernoulli random variables, with parameter θ that equals to $\mathbb{E}(\bar{C}_{N_x})$, after scaling and shifting it to be in the interval $(0, 1]$. This implies, by lemma 9, that a judicious linear transformation of the output of the bias estimation machine is at most β far from $\mathbb{E}(\bar{C}_{N_x})$ with probability at least $1 - \delta$, thus giving the (ε, δ) guarantee of Theorem 1.

Lemma 2: For any $m \leq 2^\ell$ for some $1 \leq \ell \leq M - 1$, it holds that

$$\Pr(N < m) \leq e \cdot 2^{-\frac{1}{2} \cdot (M - \ell - 1)^2}. \quad (16)$$

Furthermore, for any $m \geq \alpha \cdot 4n^c$, it holds that

$$\Pr(N > m) \leq 5e^{-\alpha}. \quad (17)$$

Proof: Let τ_k be the time it takes to move from state k to state $k + 1$ in the first Morris counter and note that

$\tau_k \sim \text{Geo}(2^{-k})$. Thus, $N = \sum_{k=1}^{M-1} \tau_k$. The moment generating function of τ_k is

$$\mathbb{M}_{\tau_k}(s) \triangleq \mathbb{E}(e^{s\tau_k}) = \frac{1}{1 + 2^k(e^{-s} - 1)}, \quad (18)$$

and it is defined for all $s < -\ln(1 - 2^{-k})$. The moment generating function of N is therefore

$$\mathbb{M}_N(s) = \prod_{k=1}^{M-1} \mathbb{M}_{\tau_k}(s) = \prod_{k=1}^{M-1} \frac{1}{1 + 2^k(e^{-s} - 1)}, \quad (19)$$

and is defined for all $s < -\ln(1 - 2^{-(M-1)}) = -\ln(1 - 2^{-B})$. We apply Chernoff bound to prove both results. For the first bound, we have $\Pr(N < m) \leq e^{-sm} \cdot \mathbb{M}_N(s)$ for any $s < 0$. Setting $s = -\ln(1 + 1/m)$, we get

$$\Pr(N < m) \leq \left(1 + \frac{1}{m}\right)^m \cdot \prod_{k=1}^{M-1} \frac{1}{1 + \frac{2^k}{m}} \quad (20)$$

$$\leq e \cdot 2^{-\sum_{k=1}^{M-1} \log\left(1 + \frac{2^k}{m}\right)}. \quad (21)$$

As $m \leq 2^\ell$ for some $1 \leq \ell \leq M - 1$, we have $2^k/m \geq 2^{k-\ell}$, so we can lower bound the exponent above with

$$\sum_{k=1}^{M-1} \log(1 + 2^{k-\ell}) = \sum_{k=0}^{\ell-1} \log(1 + 2^{-k}) + \sum_{k=\ell}^{M-1} \log(1 + 2^k) \quad (22)$$

$$\geq \sum_{k=1}^{M-\ell-1} k \quad (23)$$

$$\geq \frac{1}{2} \cdot (M - \ell - 1)^2. \quad (24)$$

For the second bound, we have $\Pr(N > m) \leq e^{-sm} \cdot \mathbb{M}_N(s)$ for any $s > 0$. Setting $s = -\ln(1 - 1/4n^c)$, for which $\mathbb{M}_N(t)$ is well defined as $2^B \leq 2n^c$, we get

$$\Pr(N > m) \leq \left(1 - \frac{1}{4n^c}\right)^m \cdot \prod_{j=1}^{M-1} \frac{1}{1 - \frac{2^j}{4n^c}} \quad (25)$$

$$\leq \left(1 - \frac{1}{4n^c}\right)^m \cdot \prod_{j=1}^{M-1} \frac{1}{1 - 2^{-j}} \quad (26)$$

$$\leq 5 \exp\left\{-\frac{m}{4n^c}\right\} = 5e^{-\alpha} \quad (27)$$

where in (26) we used the fact that $\frac{2^j}{4n^c} \leq \frac{2^j}{2^T} = 2^{j-M}$, and in (27) we used the bound $\prod_{j=1}^M (1 - 2^{-j}) \geq \frac{1}{4} + \frac{1}{2^{M+1}}$, which can be proved via induction. \square

Lemma 3: Let $v_n(\alpha) \triangleq \sqrt{\frac{2c\alpha^3}{\log n}} + \frac{\alpha}{\log n}$. Then for any $0 < \alpha \leq 1$, we have that

$$\mathbb{E}(N^{-\alpha}) \leq (e + 1)n^{-c\alpha + v_n(\alpha)}. \quad (28)$$

Proof: Appealing to Lemma 2, for any $1 \leq \ell \leq M - 1, m \leq 2^\ell$, we have

$$\mathbb{E}(N^{-\alpha}) \leq \Pr(N < m) + m^{-\alpha} \cdot \Pr(N \geq m) \quad (29)$$

$$\leq e \cdot 2^{-\frac{1}{2} \cdot (M - \ell - 1)^2} + 2^{-\ell\alpha}. \quad (30)$$

Setting $\ell = M - 1 - \lceil \sqrt{2\alpha \cdot (M - 1)} \rceil$ and recalling that $n^c \leq 2^{M-1}$, we get

$$\mathbb{E}(N^{-\alpha}) \leq e \cdot 2^{-\frac{1}{2} \lceil \sqrt{2\alpha(M-1)} \rceil^2}$$

$$+ 2^{-\alpha(M-1) \cdot \left(1 - \frac{\lceil \sqrt{2\alpha(M-1)} \rceil}{M-1}\right)} \quad (31)$$

$$\leq e \cdot 2^{-\alpha(M-1)} + 2^{-\alpha(M-1) \cdot \left(1 - \frac{\sqrt{2\alpha(M-1)}+1}{M-1}\right)} \quad (32)$$

$$\leq e \cdot n^{-c \cdot \alpha} + n^{-c \cdot \alpha + v_n(\alpha)}. \quad (33)$$

□

According to Theorem 4, the value of the infinite memory Morris counter after m updates is close to $\log m$ in expectation, up to some small bias. We would like to show that, given N and N_X , the expectation of our counter is close to the expectation of $\log(N_X)$, which is the expectation of the logarithm of a Binomial variable, and that centering and taking the expectation over (N, X) gives us approximately $-H(p)$. To that end, we first analyze the algorithm under the assumption that the second counter is an infinite memory Morris counter, which we denote as $C_{N_X}^\infty$, and then we prove that the expected gap between $C_{N_X}^\infty$ and C_{N_X} is small.

Lemma 4: Let $\phi_c(n) = (e + 1)n^{-(c-1)+v_n(1)} + \min\left\{1, C \cdot n^{-\frac{1}{2} \cdot (c-1) + v_n(1/2)}\right\}$, where $C = 2(e + 1) \cdot 10^8$. Then

$$|\mathbb{E}(\bar{C}_{N_X}^\infty) + H(p)| \leq 10^{-5} + \phi_c(n). \quad (34)$$

Proof: As $C_{N_X}^\infty$ is the state of the infinite memory Morris counter, we have that $\bar{C}_{N_X}^\infty = C_{N_X}^\infty - \mu - \mathbb{E} \log N$. Given N and $X = x$, the number of x observations in the N sample window is distributed Binomial(N, p_x). We show that $\mathbb{E}(\bar{C}_{N_X}^\infty)$ is close to the expected logarithm of the normalized Binomial random variable, which then gives us $-H(p)$ plus some bias. From Theorem 4, we have

$$\begin{aligned} \mathbb{E}(\bar{C}_{N_X}^\infty | X, N_X) &= \mathbb{E}(C_{N_X}^\infty - \mu - \mathbb{E}(\log N) | X, N_X) \\ &= \log N_X - \mathbb{E}(\log N) + g(\log N_X) + \phi(N_X). \end{aligned} \quad (35)$$

$$= \log N_X - \mathbb{E}(\log N) + g(\log N_X) + \phi(N_X). \quad (36)$$

Note that

$$\begin{aligned} \mathbb{E}(\log N_X | X = x) - \mathbb{E}(\log N) &= \sum_{k=1}^{\infty} \Pr(N_X = k) \log k - \sum_{m=1}^{\infty} \Pr(N = m) \log m \\ &= \sum_{m=1}^{\infty} \Pr(N = m) \sum_{k=1}^m \Pr(N_X = k | N = m) \log \frac{k}{m} \end{aligned} \quad (37)$$

$$= \sum_{m=1}^{\infty} \Pr(N = m) \sum_{k=1}^m \Pr(N_X = k | N = m) \log \frac{k}{m} \quad (38)$$

$$= \mathbb{E}\left(\frac{N_X}{N} | X = x\right). \quad (39)$$

Thus, letting $\gamma_{N_X} = g(\log N_X) + \phi(N_X)$, we have

$$\begin{aligned} \mathbb{E}(\bar{C}_{N_X}^\infty | X) &= \mathbb{E}\left(\log \frac{N_X}{N} | X\right) + \mathbb{E}(\gamma_{N_X} | X) \\ &= \log p_X + \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X} | X\right) + \mathbb{E}(\gamma_{N_X} | X), \end{aligned} \quad (40)$$

$$= \log p_X + \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X} | X\right) + \mathbb{E}(\gamma_{N_X} | X), \quad (41)$$

implying $\mathbb{E}(\bar{C}_{N_X}^\infty) = -H(p) + \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right) + \mathbb{E}(\gamma_{N_X})$. We conclude the proof by bounding $\mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right)$ in Lemma 5, and then bounding $\mathbb{E}(\gamma_{N_X})$ in Lemma 6. □

Lemma 5: It holds that

$$0 \leq \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right) \leq (e + 1)n^{-(c-1)+v_n(1)}. \quad (42)$$

Proof: We first show that $\mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right) \geq 0$. By Jensen's inequality and convexity of $t \mapsto -\log(t)$

$$\begin{aligned} \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right) &= \mathbb{E}_{X,N}\left[\mathbb{E}_{N_X|N,X}\left(-\log \frac{N \cdot p_X}{N_X}\right)\right] \\ &\geq -\mathbb{E}_{X,N}\left[\log\left(N \cdot p_X \cdot \mathbb{E}_{N_X|N,X}\left[\frac{1}{N_X}\right]\right)\right]. \end{aligned} \quad (43)$$

$$\geq -\mathbb{E}_{X,N}\left[\log\left(N \cdot p_X \cdot \mathbb{E}_{N_X|N,X}\left[\frac{1}{N_X}\right]\right)\right]. \quad (44)$$

To establish non-negativity of $\mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right)$, it therefore suffices to show that $\mathbb{E}_{N_X|N,X=x}\left[\frac{1}{N_X}\right] \leq \frac{1}{p_X \cdot N}$. To that end, recall that given $X = x$ and N , we have $N_X \sim \text{Bin}(N - 1, p_x) + 1$. Thus, we indeed have

$$\mathbb{E}_{N_X|N,X=x}\left[\frac{1}{N_X}\right] \quad (45)$$

$$\begin{aligned} &= \sum_{m=0}^{N-1} \frac{1}{m+1} \binom{N-1}{m} p_x^m (1-p_x)^{N-m-1} \\ &= \sum_{m=0}^{N-1} \frac{1}{p_x \cdot N} \binom{N}{m+1} p_x^{m+1} (1-p_x)^{N-m-1} \\ &= \frac{1 - (1-p_x)^N}{p_x \cdot N} \\ &\leq \frac{1}{p_x \cdot N}. \end{aligned} \quad (46)$$

To upper bound $\mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right)$, we use Jensen's inequality and the concavity of $t \mapsto \log t$, to obtain

$$\begin{aligned} \mathbb{E}_{N_X|N,X=x}\left(\log \frac{N_X}{N \cdot p_X}\right) &\leq \log\left(\frac{\mathbb{E}_{N_X|N,X=x}[N_X]}{N \cdot p_X}\right) \\ &= \log\left(1 + \frac{1-p_x}{N \cdot p_X}\right) \leq \frac{1}{N \cdot p_X}. \end{aligned} \quad (47)$$

$$= \log\left(1 + \frac{1-p_x}{N \cdot p_X}\right) \leq \frac{1}{N \cdot p_X}. \quad (48)$$

Thus, overall,

$$\mathbb{E}\left[\log \frac{N_X}{N \cdot p_X}\right] \leq \mathbb{E}_{N,X}\left[\frac{1}{N \cdot p_X}\right] \quad (49)$$

$$= \mathbb{E}_X\left[\frac{1}{p_X}\right] \mathbb{E}_N\left[\frac{1}{N}\right] \quad (50)$$

$$= n \cdot \mathbb{E}_N\left[\frac{1}{N}\right] \quad (51)$$

and appealing to Lemma 3 with $\alpha = 1$, we have $\mathbb{E}\left[\log \frac{N_X}{N \cdot p_X}\right] \leq (e + 1)n^{-(c-1)+v_n(1)}$. □

Lemma 6: It holds that

$$\mathbb{E}(\gamma_{N_X}) \leq 10^{-5} + \min\left\{1, C \cdot n^{-\frac{1}{2} \cdot (c-1) + v_n(1/2)}\right\}. \quad (52)$$

Proof: Note that $\mathbb{E}(g(\log N_X)) \leq 10^{-5}$ is explicit in Theorem 4 for any $x \in [n]$, and in particular, $\mathbb{E}(g(\log N_X)) \leq 10^{-5}$.

Thus, it remains to upper bound $\mathbb{E}(\phi(N_X))$. It is straightforward to verify that $\phi(x) \leq \min\left\{1, \frac{2 \cdot 10^8}{\sqrt{x}}\right\}$ for all $x \geq 1$, and consequently,

$$\mathbb{E}(\phi(N_X)) \leq \mathbb{E}\left[\min\left\{1, \frac{2 \cdot 10^8}{\sqrt{N_X}}\right\}\right] \quad (53)$$

$$\leq \min\left\{1, 2 \cdot 10^8 \mathbb{E}\left[\sqrt{\frac{1}{N_X}}\right]\right\}. \quad (54)$$

From Jensen's inequality, concavity of $t \mapsto \sqrt{t}$, and equation (46),

$$\mathbb{E}\left[\sqrt{\frac{1}{N_X}}\right] = \mathbb{E}_{N,X}\left[\mathbb{E}_{N_X|N,X}\left[\sqrt{\frac{1}{N_X}}\right]\right] \quad (55)$$

$$\leq \mathbb{E}_{N,X}\left[\sqrt{\mathbb{E}_{N_X|N,X}\left[\frac{1}{N_X}\right]}\right] \quad (56)$$

$$\leq \mathbb{E}_{N,X}\left[\sqrt{\frac{1}{p_X \cdot N}}\right] = \mathbb{E}_N\left[\sqrt{\frac{1}{N}}\right] \mathbb{E}_X\left[\sqrt{\frac{1}{p_X}}\right]. \quad (57)$$

Note that, again using Jensen's inequality and concavity of $t \mapsto \sqrt{t}$, we have

$$\mathbb{E}_X\left[\sqrt{\frac{1}{p_X}}\right] = \sum_{x=1}^n \sqrt{p_x} \leq n \sqrt{\frac{1}{n} \sum_{x=1}^n p_x} = \sqrt{n}. \quad (58)$$

Appealing to Lemma 3 with $\alpha = 0.5$, we have

$$\mathbb{E}(N^{-0.5}) \leq (e+1)n^{-\frac{c}{2} + v_n(1/2)}. \quad (59)$$

Thus, substituting (58) and (59) into (57) and then into (54), and recalling that $C = 2(e+1)10^8$, we obtain the claimed result. \square

Lemma 7 below bounds the absolute difference between the expectation of the Morris counter and the expectation of the truncated Morris counter of the algorithm by $O((\log n)/n^c)$. The proof is relegated to the appendix.

Lemma 7: We have

$$|\mathbb{E}(C_{N_X}^\infty) - \mathbb{E}(C_{N_X})| \leq n^{-c} \cdot \frac{100(c \log n + 2)}{(1 - 0.5n^{-c})^2} \quad (60)$$

We now turn to analyzing the bias estimation phase of the algorithm.

Lemma 8: Let Y_{N_1}, Y_{N_2}, \dots denote the sequence of Bernoulli random variables fed to the bias estimation machine. Then

$$Y_{N_1}, Y_{N_2}, \dots \stackrel{i.i.d.}{\sim} \text{Bern}(\theta), \quad (61)$$

where $\theta = \frac{H(p)+b}{2M} + a$ and $|b| \leq 10^{-5} + \psi_c(n)$.

Proof: The sequence of samples is i.i.d. since each sample is a function of the i.i.d. series $\{X_i\}_{i=1}^\infty$ and the statistics of the Morris counters, which are initialized at every incrementation. Given $(X, N, N_X, C_{N_X}) = (x, m, n_x, C_{N_X})$, we set the Bernoulli parameter $\theta_{N_X} = a - \frac{C_{N_X}}{2M}$. Thus the unconditioned parameter θ is a mixture of θ_{N_X} over the joint distribution of (X, N, N_X, C_{N_X}) , that is,

$$\theta = \mathbb{E}(\theta_{N_X}) = a - \frac{\mathbb{E}(C_{N_X})}{2M} = a + \frac{H(p)+b}{2M}, \quad (62)$$

where we used Lemma 4 and Lemma 7. \square

Lemma 9: We have

$$\Pr(|\hat{H} - (H(p) + b)| > \beta) \leq \delta. \quad (63)$$

Proof: Recall that $\hat{H} = 2M(\hat{\theta} - a)$, and that $\mathbb{E}(\hat{\theta} - \theta)^2 \leq \frac{1}{S_{\text{bias}} - 1}$. As $S_{\text{bias}} = \left\lceil \frac{4M^2}{\beta^2 \delta} \right\rceil + 1$, we have

$$\mathbb{E}(\hat{H} - (H(p) + b))^2 = 4M^2 \cdot \mathbb{E}(\hat{\theta} - \theta)^2 \leq \beta^2 \delta, \quad (64)$$

thus, from Chebyshev's inequality,

$$\Pr(|\hat{H} - (H(p) + b)| > \beta) \leq \frac{\mathbb{E}(\hat{H} - (H(p) + b))^2}{\beta^2} \leq \delta. \quad (65)$$

\square

Note that our upper bound on the additive error in estimation of $H(p)$ is $\beta + |b| \leq \beta + 10^{-5} + \psi_c(n)$, which limits our results to estimation error $\varepsilon > 10^{-5} + \psi_c(n)$.

D. Analysis of the Algorithm for $t < \infty$

In the previous analysis, the number of observed samples was assumed to be unbounded. In practice we only need to observe $O(t_{\text{mix}}(\theta))$ samples, where $t_{\text{mix}}(\theta)$ is the mixing time of our machine whenever the input is $\text{Bern}(\theta)$ samples, i.e., the minimal time it takes for the total variation distance between the marginal distribution and the limiting distribution to be small. Lemma 10 and Lemma 11 bound the number of samples needed at the Morris counting phase and characterize the mixing time of the bias estimation machine, respectively. Combining the previous results, Lemma 12 shows that the total run time of the algorithm needed to obtain an ε additive approximation of the entropy with probability at least $1 - 3\delta$ is as prescribed by Theorem 1.

Specifically, recall that the bias estimation machine is only incremented after an iteration of the first Morris counter is completed, and the run time of each iteration is a random variable that is only bounded in expectation. We note that this in fact implies the existence of a good algorithm that has a bounded sample complexity; namely, running our entropy estimation algorithm on L samples is equivalent to running the bias estimation machine from [10] on a random number of samples $k = k(L)$ times with $\theta = \mathbb{E}(\theta_{N_X})$. The randomness in $k(L)$ follows since the runtime N_i of each iteration of the Morris counter procedure is a random variable. We use Chernoff's bound to upper bound the probability that $k(L)$ is small. This event is considered as an error in our analysis. We now upper bound the mixing time of the bias estimation machine from [10]. Whenever $k(L)$ is greater than this mixing time, the error of our algorithms with L samples is close to its asymptotic value.

To upper bound the mixing time, we use the *coupling method*. Recall that the transition matrix P of a Markov process $\{X_t\}_{t=0}^\infty$ supported on \mathcal{X} is a matrix whose elements are $\Pr(X_{t+1} = x' | X_t = x) = P(x, x')$, for any $x, x' \in \mathcal{X} \times \mathcal{X}$. We define a coupling of Markov chains with transition matrix P to be a process $\{X_t, Y_t\}_{t=0}^\infty$ with the property that both $\{X_t\}_{t=0}^\infty$ and $\{Y_t\}_{t=0}^\infty$ are Markov chains with transition matrix P , although the two chains may be correlated and have different initial distributions. Given a Markov chain on \mathcal{X} with transition matrix P , a *Markovian coupling* of two P -chains is a Markov

chain $\{X_t, Y_t\}_{t=0}^{\infty}$ with state space $\mathcal{X} \times \mathcal{X}$, which satisfies, for all x, y, x', y' ,

$$\Pr(X_{t+1} = x' | X_t = x, Y_t = y) = P(x, x'), \quad (66)$$

$$\Pr(Y_{t+1} = y' | X_t = x, Y_t = y) = P(y, y'). \quad (67)$$

Let $P^t(x_0)$ be the marginal distribution of the chain at time t when initiated at x_0 , and let π be the unique stationary distribution. Define the δ -mixing time as

$$t_{\delta}^* \triangleq \min\{t : d_{TV}(P^t(x_0), \pi) \leq \delta\}, \quad (68)$$

and $t_{\text{mix}} \triangleq t_{1/4}^*$. We now show that the bias estimation machine with S states mixes in $\Theta(S \log S)$ time, uniformly for all $\theta \in (0, 1]$.

Lemma 10: Let $t_{\text{mix}}(p)$ denote the mixing time of the bias estimation machine with S states when the input is i.i.d. $\text{Bern}(p)$, and define the *worst-case* mixing time to be $t^* = \max_{p \in (0,1]} t_{\text{mix}}(p)$. Then

$$\ln(2) \cdot (S-1) \log(S-1) \leq t^* \leq 4S \log S. \quad (69)$$

Proof: The transition probabilities of the bias estimation machine of Figure 2 are given, for $1 < k < S$, as

$$X_{t+1}|X_t=k = \begin{cases} k+1, & \text{w.p. } \frac{S-k}{S-1} \cdot p, \\ k, & \text{w.p. } \frac{k-1}{S-1} \cdot p + \frac{S-k}{S-1} \cdot q, \\ k-1, & \text{w.p. } \frac{k-1}{S-1} \cdot q, \end{cases} \quad (70)$$

and for the extreme states $\{1, S\}$ as

$$X_{t+1}|X_t=1 = \begin{cases} 2, & \text{w.p. } p, \\ 1, & \text{w.p. } q, \end{cases} \quad (71)$$

and

$$X_{t+1}|X_t=S = \begin{cases} S, & \text{w.p. } p, \\ S-1, & \text{w.p. } q. \end{cases} \quad (72)$$

We construct a Markovian coupling in which the two chains stay together at all times after their first simultaneous visit to a single state, that is

$$\text{if } X_s = Y_s \text{ then } X_t = Y_t \text{ for all } t \geq s. \quad (73)$$

The following theorem is due to [36] (Theorem 5.4), will give us an upper bound on the mixing time using this coupling.

Theorem 5: Let $\{(X_t, Y_t)\}$ be a Markovian coupling satisfying (73), for which $X_0 = x_0$ and $Y_0 = y_0$. Let τ_{couple} be the coalescence time of the chains, that is,

$$\tau_{\text{couple}} \triangleq \min\{t : X_t = Y_t\}. \quad (74)$$

Then

$$t_{\text{mix}} \leq 4 \max_{x_0, y_0 \in \mathcal{X}} \mathbb{E}(\tau_{\text{couple}}). \quad (75)$$

Assume w.l.o.g. that $x_0 < y_0$ and let U_t be an i.i.d. sequence drawn according to the $\text{Unif}(0, 1)$ distribution. We construct a coupling on (X_t, Y_t) such that, at each time point $t < \tau_{\text{couple}}$, X_t and Y_t are incremented in the following manner:

$$X_{t+1}|X_t=i = \begin{cases} i+1, & \text{if } U_t \leq \frac{S-i}{S-1} \cdot p, \\ i, & \text{if } \frac{S-i}{S-1} \cdot p \leq U_t \leq 1 - \frac{i-1}{S-1} \cdot q, \\ i-1, & \text{if } 1 - \frac{i-1}{S-1} \cdot q \leq U_t \leq 1, \end{cases} \quad (76)$$

and

$$Y_{t+1}|Y_t=j = \begin{cases} j+1, & \text{if } U_t \leq \frac{S-j}{S-1} \cdot p, \\ j, & \text{if } \frac{S-j}{S-1} \cdot p \leq U_t \leq 1 - \frac{j-1}{S-1} \cdot q, \\ j-1, & \text{if } 1 - \frac{j-1}{S-1} \cdot q \leq U_t \leq 1. \end{cases} \quad (77)$$

One can validate that the transition probabilities are the correct ones, for example

$$\begin{aligned} \Pr(X_{t+1} = i | X_t = i) &= \Pr\left(\frac{S-i}{S-1} \cdot p \leq U_t \leq 1 - \frac{i-1}{S-1} \cdot q\right) \\ &= 1 - \frac{i-1}{S-1} \cdot q - \frac{S-i}{S-1} \cdot p \end{aligned} \quad (78)$$

$$= 1 - \frac{i-1}{S-1} \cdot q - \frac{S-i}{S-1} \cdot p \quad (79)$$

$$= \frac{i-1}{S-1} \cdot p + \frac{S-i}{S-1} \cdot q, \quad (80)$$

and, similarly, $\Pr(Y_{t+1} = j | Y_t = j) = \frac{j-1}{S-1} \cdot p + \frac{S-j}{S-1} \cdot q$. The other transition probabilities are easily calculated. Note that $i < j$ implies $\frac{S-j}{S-1} < \frac{S-i}{S-1}$, thus Y_t cannot move right unless X_t moves right and X_t cannot move left unless Y_t moves left. Moreover, since $x_0 < y_0$, we have $i < j$ for all $t < \tau_{\text{couple}}$. This follows from construction, since $\frac{S-i}{S-1} \cdot p$ is always smaller than $1 - \frac{j-1}{S-1} \cdot q$, implying that X_t cannot jump over Y_t when they are one-state apart. Thus, the *distance* process $D_t \triangleq Y_t - X_t$, is a non-increasing function of t , with initial state $D_0 = y_0 - x_0$, that can only decrease by one unit at a time or stay unchanged. We have

$$\begin{aligned} \Pr(D_{t+1} = D_t - 1) &= \Pr(X_{t+1} = X_t + 1, Y_{t+1} = Y_t) \\ &\quad + \Pr(Y_{t+1} = Y_t - 1, X_{t+1} = X_t) \end{aligned} \quad (81)$$

$$\begin{aligned} &= \Pr\left(\frac{S-Y_t}{S-1} \cdot p \leq U_t \leq \frac{S-X_t}{S-1} \cdot p\right) \\ &\quad + \Pr\left(1 - \frac{Y_t-1}{S-1} \cdot q \leq U_t \leq 1 - \frac{X_t-1}{S-1} \cdot q\right) \end{aligned} \quad (82)$$

$$= \frac{Y_t - X_t}{S-1} \cdot p + \frac{Y_t - X_t}{S-1} \cdot q \quad (83)$$

$$= \frac{D_t}{S-1}. \quad (84)$$

The expected coupling time is now the expected time it takes for D_t to decrease from D_0 to D_t , thus in order to maximize it under the given coupling, we need to maximize D_0 , which corresponds to setting $X_0 = 1, y_0 = S$. For $D_0 = S-1$, consider the process $M_t \triangleq D_0 - D_t$, which is a non-decreasing function of t that goes from 0 to $S-1$ and has $\Pr(M_{t+1} = M_t + 1) = \Pr(D_{t+1} = D_t - 1) = \frac{D_t}{S-1} = 1 - \frac{M_t}{S-1}$. Then this process is no other than the *Coupon Collector* process with $S-1$ coupons, and the expected coupling time in our chain is identical to the expected number of coupons collected until the set contains all $S-1$ types, which according to [36], Proposition 2.3., is

$$\mathbb{E}(\tau_{\text{couple}}) = (S-1) \cdot \sum_{k=1}^{S-1} \frac{1}{k} \quad (85)$$

$$\leq (S-1)(\ln(S-1) + 1) \quad (86)$$

$$\leq S \log(S). \quad (87)$$

To show that this upper bound is indeed tight, consider the case of $p = 1$. In this case, the chain of Figure 2 is simply the

Coupon Collector process with $S - 1$ coupons, thus, letting τ be the (random) time it takes to collect all coupons, we have

$$\mathbb{E}(\tau) = (S - 1) \cdot \sum_{k=1}^{S-1} \frac{1}{k} \quad (88)$$

$$\geq \ln(2) \cdot (S - 1) \log(S - 1). \quad (89)$$

□

From [36], Eq. (4.34), we have that the δ -mixing time t_δ^* can be upper bounded in terms on the mixing time by

$$t_\delta^* \leq \left\lceil \log \left(\frac{1}{\delta} \right) \right\rceil \cdot t_{\text{mix}}. \quad (90)$$

Set

$$k = 4 \left\lceil \log \left(\frac{1}{\delta} \right) \right\rceil \left(\frac{4(c \log n + 2)^2}{\beta^2 \delta} + 1 \right) \times \log \left(\frac{4(c \log n + 2)^2}{\beta^2 \delta} + 1 \right), \quad (91)$$

and note that from equation (90), Lemma 10, and substituting $S_{\text{bias}} = \frac{4M^2}{\beta^2 \delta} + 1$, we have that the δ -mixing time of the bias estimation machine is at most k . Let N_1, N_2, \dots, N_k be the first k i.i.d. Morris counter running times, which are all distributed as N in the analysis from Section IV. Lemma 11 uses the concentration of N to show that, with probability $1 - \delta$, the number of samples we need to observe until the bias machine mixes is not large.

Lemma 11: Let $m = 4n^c \cdot \ln \left(\frac{5k}{\delta} \right)$. Then

$$\Pr \left(\sum_{i=1}^k N_i > k \cdot m \right) \leq \delta. \quad (92)$$

Proof: Appealing to Lemma 2 we have $\Pr(N > m) \leq \delta/k$. Consequently, the probability that at least one of the random variables N_1, \dots, N_k is greater than m is at most $1 - \left(1 - \frac{\delta}{k}\right)^k \leq \delta$. □

We conclude with the following lemma, which connects Lemma 10 and Lemma 11 to show that our entropy estimator performs well even if the number of input samples is limited to $\tilde{O}(n^c/\delta)$.

Lemma 12: Let the algorithm of Theorem 1 run on $L = k \cdot m$ samples, and output the estimate \hat{H}_{M_L} . Then with probability at least $1 - 3\delta$, \hat{H}_{M_L} is within ε -additive error from $H(p)$.

Proof: Lemma 11 implies that, with probability at least $1 - \delta$, after observing $k \cdot m$ samples, the bias estimation machine has been incremented at least k times. Recall that, by definition, after $t \geq t_\delta^*$ increments of the bias estimation machine, we have that $d_{\text{TV}}(P^t(x_0), \pi) \leq \delta$, and that our S -states entropy estimator has $\sum_{i \in \hat{H}_\varepsilon} \pi_i < \delta$, where $\hat{H}_\varepsilon = \{i \in [S] : |\hat{H}_i - H(p)| > \varepsilon\}$. Thus, from a union bound, a fraction of 2δ of the distribution $P^t(x_0)$ (at most) is supported on \hat{H}_ε . Putting it all together, we have that a finite-time algorithm that outputs an estimate $\hat{H}(M_L)$ after

$$L = k \cdot m = \Omega \left(\frac{n^c \cdot \text{poly}(\log n)}{\delta} \cdot \text{poly}(\log(1/\delta)) \right) \quad (93)$$

will be ε -far from the correct entropy with probability at most 3δ .³ □

V. LOWER BOUNDS

In this section we prove Theorem 2. The $\Omega(n)$ bound is proved via reduction to uniformity testing. For the $\frac{\log n}{2\varepsilon}$ bound, we use a simple quantization argument. Assume that $S < \frac{\log n}{2\varepsilon}$. Then there must be two consecutive estimate values $\hat{H}_1, \hat{H}_2 \in [0, \log n]$ such that $\hat{H}_2 - \hat{H}_1 > 2\varepsilon$. This implies that $H = (\hat{H}_1 + \hat{H}_2)/2$ has $|H - \hat{H}_1| = |H - \hat{H}_2| > \varepsilon$. Thus, for this value of the entropy, we have $\Pr(|\hat{H}(M_t) - H| > \varepsilon) = 1$ for all $t \in \mathbb{N}$.

A. Proof of the $(1 - 2\sqrt{\varepsilon \ln 2})n$ Bound

An (ε, δ) uniformity tester can distinguish (with probability $0 < \delta < 1/2$) between the case where p is uniform and the case where p is ε -far from uniform in total variation. Assume we have an (ε, δ) entropy estimator. Then we can obtain an $(\tilde{\varepsilon} = \sqrt{\varepsilon \ln 2}, \delta)$ uniformity tester using the following protocol: the tester declares that p is uniform if $\hat{H} > \log n - \varepsilon$, and that p is $\tilde{\varepsilon}$ -far from uniform if $\hat{H} < \log n - \varepsilon$. We now argue that this is indeed an $(\tilde{\varepsilon}, \delta)$ uniformity tester, in which case the $(1 - 2\tilde{\varepsilon})n$ lower bound will follow immediately from the lower bound on uniformity testing of [7]. If $p = u$, where u is the uniform distribution over $[n]$, then $H(p) = \log n$ and $\hat{H} > \log n - \varepsilon$ with probability at least $1 - \delta$, so our tester will correctly declare “uniform” with probability at least $1 - \delta$. If $d_{\text{TV}}(p, u) > \sqrt{\varepsilon \ln 2}$, then from Pinsker’s inequality ([37], Lemma 11.6.1),

$$2\varepsilon < \frac{2}{\ln 2} d_{\text{TV}}(p, u)^2 \leq D(p||u) = \log n - H(p), \quad (94)$$

which implies $H(p) < \log n - 2\varepsilon$ and $\hat{H} < \log n - \varepsilon$ with probability at least $1 - \delta$. Thus, our tester will correctly declare “far from uniform” with probability at least $1 - \delta$.

VI. MEMORY COMPLEXITY OF MUTUAL INFORMATION ESTIMATION

We extend our results to the problem of mutual information estimation. The upper bound follows by a slight tweaking of our entropy estimation machine, and the lower bound follows by noting the close relation between mutual information and joint entropy, and lower bounding the memory complexity of the latter.

A. Upper Bound Achieving Algorithm

- 1) The algorithm receives an *accuracy parameter* $\beta > 0$ and an *overhead parameter* $c > 1$.
- 2) In each iteration of the algorithm we collect a fresh pair of samples $(X, Y) \in [n] \times [m]$ according to p_{XY} , and store their values. Assuming the received sample is x , we proceed to estimate $\log(p_x p_y / p_{xy})$ based on more fresh samples.

³Note that the probability that the second Morris counter saturates even once in k iterations is less than $O \left(n^{-c} \cdot \text{poly}(\log n) \frac{\text{poly}(\log(1/\delta))}{\delta} \right)$, thus is negligible for any $\delta \gg n^{-c}$.

- 3) We use *four* Morris counters - one that approximates a clock, one that approximates a count for x values, one that approximates a count for y values, and one that approximates a count for the pair (x, y) . The first of these counters have $M = B + 1$ states, where B is the smallest integer k such that $\lceil (n \cdot m)^c \rceil \leq 2^k$. This counter (denoted as C_N) approximates a clock that counts until $\lceil (n \cdot m)^c \rceil$ samples from the distribution are observed. The second, third and fourth counters run in parallel to the first one and approximate a counter for x , a counter for y , and a counter for the pair (x, y) , and we denote their outputs as C_{N_x} , C_{N_y} , and $C_{N_{xy}}$, respectively. These counters each have $2M$ states, to guarantee they do not exceed the first counter with high probability. In the event that any of them reaches state $2M$ before the first counter, we draw a fresh sample and initialize all counters.
- 4) We define $C_{MI} = C_{N_x} + C_{N_y} - C_{N_{xy}}$, and let $\bar{C}_{MI} = C_{MI} - \mu - \mathbb{E} \log N$ be the centralized version of C_{MI} . This is an almost unbiased estimator for $-\log(p_x p_y / p_{xy})$.
- 5) We now increment a bias estimation machine with $S_{bias} = \left\lceil \frac{36M^2}{\beta^2 \delta} \right\rceil + 1$ states whose purpose is to simulate the expectation operation. Specifically, each time the first Morris counter concludes a count, we generate a $\text{Ber}(\theta_{N_{xy}})$ random variable, with $\theta_{N_{xy}} = a - \frac{\bar{C}_{MI}}{6M}$, and use it as the input to our bias estimation machine. The offset $a \triangleq \frac{4M - \mathbb{E}(\log N) - \mu}{6M}$ guarantees that $\theta_{N_{xy}} \in [0, 1)$ with probability 1, as $-2M \leq C_{MI} \leq 4M$ since $C_{N_x}, C_{N_y}, C_{N_{xy}}$ are the outputs of Morris counters with $2M$ states. Our estimator for the mutual information \hat{I} is the bias estimate of the machine, after subtraction of the known offset a and multiplication by $6M$, that is, $\hat{I} = 6M(\hat{\theta} - a)$.

B. Number of States of Mutual Information Estimator

$n, m, t, \beta, \delta, c, M, \eta, a$, and S_{bias} are program constants, so we do not count them in the memory consumption of the algorithm. At each time point, our algorithm keeps the value of a pair (x, y) , which requires $n \cdot m$ states, the state of the Morris counter approximating the clock, the state of the Morris counter approximating the logarithm of the x counter, and the state of the bias estimation machine. Thus, the total number of states is the product of the individual number of states needed at each step, and recalling that $M = B + 1 \leq c \log nm + 2$, the total number of states is

$$S \leq nm \cdot M \cdot (2M)^3 \cdot \left(\frac{36M^2}{\beta^2 \delta} + 2 \right) \\ = nm \left(\frac{288 \cdot (c \log nm + 2)^6}{\beta^2 \delta} + 16(c \log nm + 2)^4 \right). \quad (95)$$

C. Analysis of the Algorithm for $t = \infty$

Let (X, Y) be the fresh sample pair collected at the start of an algorithm iteration. We begin our analysis by showing that \bar{C}_{MI} is close in expectation to $I(X; Y)$. Let $C_{N_x}^\infty, C_{N_y}^\infty$ and $C_{N_{xy}}^\infty$ denote the corresponding infinite Morris counters. We first analyze the algorithm for these counters, and then appeal

Algorithm 4 Mutual Information Estimation With Morris Counters

Input: A data stream $(X_1, Y_1), (X_2, Y_2) \dots \sim p_{XY}$, alphabet size n , alphabet size m , run time t , error probability $\delta, \beta > 0, c > 1$, constant μ of (12)

Output: Mutual Information estimate \hat{I}

1: Set

$$B \leftarrow \min\{k \in \mathbb{N} : \lceil (nm)^c \rceil \leq 2^k\}, M \leftarrow B + 1$$

$$\eta \leftarrow \text{Monte Carlo estimate of } \mathbb{E}(\log N) \text{ for}$$

$$N = \sum_{k=1}^{M-1} \tau_k, \text{ where } \tau_k \sim \text{Geo}(2^{-k})$$

$$a \leftarrow \frac{2}{3} - \frac{\mu + \eta}{6M}, S_{bias} \leftarrow \left\lceil \frac{36M^2}{\beta^2 \delta} \right\rceil + 1$$

$$C_N \leftarrow 1, C_{N_x}, C_{N_y}, C_{N_{xy}} \leftarrow 1, s \leftarrow 1$$

2: **for** $i = 1, \dots, t$ **do**

3: **if** $C_N = 1$ **then**

$$4: \quad (x_{\text{test}}, y_{\text{test}}) \leftarrow (X_i, Y_i)$$

$$5: \quad C_{N_x}, C_{N_y}, C_{N_{xy}} \leftarrow 1$$

6: **else**

$$7: \quad C_N = \text{IncrementMorrisCounter}(C_N)$$

8: **if** $X_i = x_{\text{test}}$ **then**

$$9: \quad C_{N_x} \leftarrow \text{IncrementMorrisCounter}(C_{N_x})$$

10: **if** $Y_i = y_{\text{test}}$ **then**

$$11: \quad C_{N_y} \leftarrow \text{IncrementMorrisCounter}(C_{N_y})$$

$$12: \quad C_{N_{xy}} \leftarrow \text{IncrementMorrisCounter}(C_{N_{xy}})$$

13: **end if**

14: **else if** $Y_i = y_{\text{test}}$ **then**

$$15: \quad C_{N_y} \leftarrow \text{IncrementMorrisCounter}(C_{N_y})$$

16: **end if**

17: **if** $\max\{C_{N_x}, C_{N_y}, C_{N_{xy}} < 2M\}$ **then**

18: **if** $C_N = M$ **then**

$$19: \quad C_{MI} \leftarrow C_{N_x} + C_{N_y} - C_{N_{xy}}$$

$$20: \quad \theta_{N_{xy}} \leftarrow a - \frac{C_{MI} - (\mu + \eta)}{6M}$$

$$21: \quad s \leftarrow \text{IncrementBiasEstimation}(S_{bias}, s, \theta_{N_{xy}})$$

$$22: \quad C_N \leftarrow 1$$

23: **end if**

24: **else**

$$25: \quad C_N \leftarrow 1$$

26: **end if**

27: **end if**

28: **end for**

$$29: \quad \hat{\theta}_{MI} \leftarrow \frac{s-1}{S_{bias}-1}$$

$$30: \quad \hat{I} \leftarrow 6M(\hat{\theta}_{MI} - a)$$

to Lemma 7 to bound the deviation of the limited memory counters used in our algorithm.

Lemma 13: Denote $d_c(n, \alpha) = n^{-\alpha(c-1) + v_n(\alpha)}$ and let

$$\phi_c(n, m)$$

$$= 2(e+1) \max\{d_c(n, 1), d_c(m, 1)\}$$

$$+ 3 \min \left\{ 1, C \max \left\{ d_c \left(n, \frac{1}{2} \right), d_c \left(m, \frac{1}{2} \right), d_c \left(nm, \frac{1}{2} \right) \right\} \right\}.$$

We have

$$|\mathbb{E}(\bar{C}_{MI}^\infty) + I(X; Y)| \leq 3 \cdot 10^{-5} + \phi_c(n, m). \quad (96)$$

Proof: Following the proof of Lemma 4, we write

$$\begin{aligned} \mathbb{E}(\bar{C}_{\text{MI}}^\infty | X, N_X) \\ = \mathbb{E}(C_{N_X}^\infty + C_{N_Y}^\infty - C_{N_{XY}}^\infty - \mu - \mathbb{E}(\log N) | X, N_X) \end{aligned} \quad (97)$$

$$= \log N_X + \log N_Y - \log N_{XY} - \mathbb{E}(\log N) + \gamma_{N_{XY}}, \quad (98)$$

where $\gamma_{N_{XY}} = \gamma_{N_X} + \gamma_{N_Y} + \gamma_{N_{XY}}$. We then have

$$\begin{aligned} \mathbb{E}(\bar{C}_{\text{MI}}^\infty | X) \\ = \mathbb{E}\left(\log \frac{N_X N_Y / N_{XY}}{N} | X\right) + \mathbb{E}(\gamma_{N_{XY}} | X) \end{aligned} \quad (99)$$

$$\begin{aligned} = \log \frac{p_X p_Y}{p_{XY}} + \mathbb{E}\left(\log \frac{N_X N_Y / N_{XY}}{N \cdot p_X p_Y / p_{XY}} | X\right) \\ + \mathbb{E}(\gamma_{N_X} | X), \end{aligned} \quad (100)$$

implying $\mathbb{E}(\bar{C}_{\text{MI}}^\infty) = -I(X; Y) + \mathbb{E}\left(\log \frac{N_X N_Y / N_{XY}}{N \cdot p_X p_Y / p_{XY}}\right) + \mathbb{E}(\gamma_{N_{XY}})$. Decomposing

$$\begin{aligned} \mathbb{E}\left(\log \frac{N_X N_Y / N_{XY}}{N \cdot p_X p_Y / p_{XY}}\right) \\ = \mathbb{E}\left(\log \frac{N_X}{N \cdot p_X}\right) + \mathbb{E}\left(\log \frac{N_Y}{N \cdot p_Y}\right) \\ - \mathbb{E}\left(\log \frac{N_{XY}}{N \cdot p_{XY}}\right) \end{aligned} \quad (101)$$

and applying Lemma 5 to each term separately, we have

$$-d_c(nm, 1) \leq \mathbb{E}\left(\log \frac{N_X N_Y / N_{XY}}{N \cdot p_X p_Y / p_{XY}}\right) \quad (102)$$

$$\leq 2(e+1) \max\{d_c(n, 1), d_c(m, 1)\}, \quad (103)$$

and, similarly, recalling Lemma 6, we have that

$$\begin{aligned} \mathbb{E}(\gamma_{N_{XY}}) \leq 3 \cdot 10^{-5} \\ + \min\{1, C \max\{d_c(n, 1/2), d_c(m, 1/2), d_c(nm, 1/2)\}\}. \end{aligned} \quad (104)$$

This implies that in the counting phase of the algorithm we obtain an estimate for (minus) the mutual information that has an average bias bounded from above by

$$\begin{aligned} 3 \cdot 10^{-5} + \phi_c(n, m) = 3 \cdot 10^{-5} \\ + O\left(\min\left\{2\sqrt{\log n} \cdot n^{-\frac{1}{2}(c-1)}, 2\sqrt{\log m} \cdot m^{-\frac{1}{2}(c-1)}\right\}\right). \end{aligned}$$

□

The additive expected error resulting from the truncation of the Morris counters $C_{N_X}^\infty$, $C_{N_Y}^\infty$ and $C_{N_{XY}}^\infty$ at state $2M$ is upper bounded according to Lemma 7 and the triangle inequality by $\frac{300(c \log nm + 2)}{(nm)^{c(1-0.5(nm)^{-c})^2}}$, which is asymptotically negligible w.r.t $\phi_c(n, m)$. In a similar fashion to Lemma 8, the input sequence to the bias estimation machine is an i.i.d. sequence with distribution $\text{Bern}(\theta_{\text{MI}})$, where

$$\theta_{\text{MI}} = \mathbb{E}(\theta_{N_{XY}}) = \mathbb{E}\left(a - \frac{\bar{C}_{\text{MI}}}{6M}\right) = a + \frac{I(X; Y) + b_{\text{MI}}}{6M}. \quad (105)$$

As we set $S_{\text{bias}} = \left\lceil \frac{36M^2}{\beta^2 \delta} \right\rceil + 1$ and $\hat{I} = 6M(\hat{\theta}_{\text{MI}} - a)$, we have

$$\begin{aligned} \mathbb{E}(\hat{I} - (I(X; Y) + b_{\text{MI}}))^2 \\ = 36M^2 \cdot \mathbb{E}(\hat{\theta}_{\text{MI}} - \theta_{\text{MI}})^2 \leq \beta^2 \delta, \end{aligned} \quad (106)$$

and we obtain the (ε, δ) guarantee from Chebyshev's inequality, i.e., $\Pr(|\hat{I} - (I(X; Y) + b_{\text{MI}})| > \beta) \leq \delta$.

D. Lower Bound

For simplicity of proof, let $\varepsilon, \delta \geq \frac{1}{300}$, and recall that $\varepsilon < \frac{1}{12 \ln 2}$. Our lower bound from Theorem 2 implies that for joint entropy estimation of $H(X, Y)$ where $(X, Y) \in [n] \times [m]$, the memory complexity is $\Omega(n \cdot m)$. Assume that we have a mutual information estimation machine that returns an estimate of $I(X; Y)$ with additive error at most ε with probability at least $1 - \delta$ using $S_{\text{MI}}^*(n, m, \varepsilon, \delta)$ states. We show below an algorithm that uses this machine as a closed box and estimates $H(X, Y) = H(X) + H(Y) - I(X; Y)$ with additive error of at most 3ε with probability at least $1 - 3\delta$ using $S_{\text{MI}}^* \cdot O(\log^3 n \cdot \log^3 m)$ states. Since estimation of $H(X, Y)$ requires $S^*(n \cdot m, 3\varepsilon, 3\delta) = \Omega(n \cdot m)$, this must imply that

$$S_{\text{MI}}^*(n, m, \varepsilon, \delta) > \Omega\left(\frac{n \cdot m}{\log^3 n \cdot \log^3 m}\right). \quad (107)$$

We now describe such an algorithm. The algorithm has 3 modes. It starts in mode 1, in which $H(X)$ is estimated. It then moves to mode 2, in which $H(Y)$ is estimated, and finally it moves to mode 3 in which $I(X; Y)$ is estimated. The current mode is stored using $S_1 = 3$ states. The estimation of each of the 3 quantities above is done using

$$\tilde{S} = \max\{S^*(n, \varepsilon, \delta), S^*(m, \varepsilon, \delta), S_{\text{MI}}^*(n, m, \varepsilon, \delta)\} \quad (108)$$

states. Those states are “reused” once the algorithm switches its mode of operation. The algorithm is as follows:

- 1) Start in Mode 1.
- 2) Increment a Morris counter with $S_2 = O(\log \log n)$ states at each observation of X . This counter determines the run time of mode 1 and we denote it *RunModeX*.
- 3) Estimate $H(X)$ using the Morris-counter entropy estimator we introduced in Section IV with $S^*(n, \varepsilon, \delta)$ states.
- 4) As *RunModeX* arrives at state S_2 , save the estimate $\hat{H}(X)$ of $H(X)$ using $S_3 = O(\log^2 n)$ states.
- 5) Switch to Mode 2.
- 6) Increment a Morris counter with $S_4 = O(\log \log m)$ states at each observation of Y . This counter determines the run time of mode 2 and we denote it *RunModeY*.
- 7) Estimate $H(Y)$ using the entropy estimator with $S^*(n, \varepsilon, \delta)$ states.
- 8) As *RunModeY* arrives at state S_4 , save the estimate $\hat{H}(Y)$ of $H(Y)$ using $S_5 = O(\log^2 m)$ states.
- 9) Switch to Mode 3.
- 10) Estimate $I(X; Y)$ using the closed-box machine with $S_{\text{MI}}^*(n, m, \varepsilon, \delta)$ states.
- 11) From this time onward, estimate $H(X, Y)$ as $\hat{H}(X) + \hat{H}(Y) - \hat{I}(X; Y)$, where $\hat{I}(X; Y)$ is the current estimate of the closed-box machine.

The idea here is that, after a long enough time, the entropy estimator output will be accurate enough, at which point we can store that value and switch modes. In order to decide if enough time has passed, we must ensure that the bias estimation machine, which outputs our entropy estimates, is sufficiently mixed. From Lemma 10, we have that the mixing time of the bias estimation machine is at most $4S_{\text{Bias}} \log S_{\text{Bias}} \leq O(\log^3 n)$ samples, as $S_{\text{Bias}} = O(\log^2 n)$ states. Thus, it suffices to run the machine for $\log^k n$ samples

of independent $\text{Ber}(\theta)$ random variables for $k \gg 1$ and then stop it, which would guarantee it is sufficiently mixed. In order to save memory we use another Morris counter with $S_2 = O(\log \log^k n) = O(\log \log n)$ states that determines when the mode run ends. We then store the state of the bias estimation machine, which corresponds to our estimate $\hat{H}(X)$ of $H(X)$, using $S_3 = S_{\text{Bias}} = O(\log^2 n)$ states. At this point, the algorithm switches to mode 2, and estimates $H(Y)$ with $S^*(m, \varepsilon, \delta)$ states. As in mode 1, we use a Morris counter of $S_4 = O(\log \log m)$ states to determine when the machine is sufficiently mixed and can be stopped, and store the state of the bias estimation machine, which corresponds to the estimate $\hat{H}(Y)$ of $H(Y)$ this time, using $S_5 = O(\log^2 m)$ states. The process then moves to state 3 where $I(X; Y)$ is estimated using the closed-box machine and, subsequently, the machine estimates $H(X, Y)$ as $\hat{H}(X) + \hat{H}(Y) - \hat{I}(X; Y)$, where $\hat{I}(X; Y)$ is the current estimate of the closed box machine. All in all, this algorithm produces a $(3\varepsilon, 3\delta)$ (recall that we assumed $\delta, \varepsilon \geq 1/100$) estimate of $H(X, Y)$ using

$$S \leq \tilde{S} \prod_{i=1}^5 S_i = \tilde{S} \cdot O(\log^3 n \cdot \log^3 m), \quad (109)$$

which implies that

$$\tilde{S} = \Omega\left(\frac{S}{\log^3 n \log^3 m}\right) \quad (110)$$

$$= \Omega\left(\frac{S^*(n, m, 3\varepsilon, 3\delta)}{\log^3 n \log^3 m}\right) \quad (111)$$

$$= \Omega\left(\frac{n \cdot m}{\log^3 n \log^3 m}\right). \quad (112)$$

Finally, since Theorem 1 states that $S^*(n, \varepsilon, \delta) = O(n \cdot \log^4 n)$ and $S^*(m, \varepsilon, \delta) = O(m \cdot \log^4 m)$, and we assumed that $\frac{n}{\log^3 n} = \Omega(\log^7 m)$ and $\frac{m}{\log^3 m} = \Omega(\log^7 n)$, we must therefore have that

$$S_{\text{MI}}^*(n, m, \varepsilon, \delta) = \Omega\left(\frac{n \cdot m}{\log^3 n \cdot \log^3 m}\right). \quad (113)$$

VII. CONCLUSION AND OPEN PROBLEMS

Due to the limitation $\varepsilon > 10^{-5}$, our upper bound is not informative when very small additive error is required. Indeed, the Morris counter seems to be inadequate in these regimes and, despite many follow up works, we are not aware of an improved analysis that cancels out the 10^{-5} term of [34]. A natural question to ask then is whether this is a true limitation arising as a result of the bounded memory or an artifact of the Morris counter. This gives rise to two potential directions for future research:

- Is there a counting algorithm with the same memory consumption as the Morris counter that does not suffer from this bias?
- Can we find an entropy estimator with similar number of states without this lower bound on the attainable additive error?

Another interesting research direction is to close the $\text{poly}(\log n)$ gap between our upper and lower bounds w.r.t the dependence on n . It seems plausible to us that the upper bound

is tight, i.e., that the real dependence on n is $n \text{poly}(\log n)$, as n is the minimal number of states needed to save one sample, and we must save our running entropy estimate as well. One possible reason for this mismatch between the bounds might be that our lower bound relies on reduction to the uniformity testing problem, which does not fully utilize the properties of a finite-state entropy estimator. In particular, the reduction is from estimation to binary hypothesis testing (testing uniform vs. ε -far from uniform), whereas in ε -additive entropy estimation we effectively have $\frac{\log n}{2\varepsilon}$ hypotheses. Particularly, it would seem that the binary test of $H(p) = \log n$ vs. $H(p) \leq \log n - \varepsilon$ is easier as there is only one distribution with $H(p) = \log n$ (uniform). Hence, another preliminary approach for lower bounds might be

- Solve the binary hypothesis testing problem $H(p) = \alpha \log n$ vs. $H(p) \leq \alpha \log n - \varepsilon$ for some $0 < \alpha < 1$.

As there are many distributions with entropy $\alpha \log n$, and since solving this problem immediately implies a lower bound on entropy estimation, this approach might help in improve upon our lower bound.

APPENDIX

A. Proof of Lemma 1

Let $p_{i,j}$ be the transition probability from state i to state j , and let π_k be the unique stationary distribution of state $k \in [S]$. We first show that $\pi_k = \mu_k$, where μ_k is the Binomial($S-1, \theta$) distribution, that is,

$$\mu_k = \binom{S-1}{k-1} p^{k-1} q^{S-k}. \quad (114)$$

For brevity, denote $\text{Bin}_p^S(k) = \binom{S-1}{k} p^k q^{S-k-1}$. As $\binom{S-1}{k-1} = \frac{k}{S-k} \binom{S-1}{k}$, we have $\text{Bin}_p^S(k-1) = \text{Bin}_p^S(k) \cdot \frac{k}{S-k} \cdot \frac{q}{p}$. Recall that if μ is the stationary distribution if and only if $\sum_{i=1}^S \mu_i p_{i,k+1} = \mu_{k+1}$ for any $k \in [S-1]$. Write

$$\begin{aligned} & \sum_{i=1}^S \mu_i p_{i,k+1} \\ &= \mu_k p_{k,k+1} + \mu_{k+1} p_{k+1,k+1} + \mu_{k+2} p_{k+2,k} \end{aligned} \quad (115)$$

$$\begin{aligned} &= \text{Bin}_p^S(k-1) \cdot \frac{S-k}{S-1} p \\ &+ \text{Bin}_p^S(k) \left(\frac{k}{S-1} p + \frac{S-(k+1)}{S-1} q \right) \\ &+ \text{Bin}_p^S(k+1) \cdot \frac{k+1}{S-1} q \end{aligned} \quad (116)$$

$$\begin{aligned} &= \text{Bin}_p^S(k) \left(\frac{k}{S-k} \cdot \frac{S-k}{S-1} q + \frac{k}{S-1} p \right. \\ &+ \left. \frac{S-(k+1)}{S-1} q + \frac{S-(k+1)}{k+1} \cdot \frac{k+1}{S-1} p \right) \end{aligned} \quad (117)$$

$$= \text{Bin}_p^S(k) (p + q) = \mu_{k+1}. \quad (118)$$

Now, due to the Ergodicity of the chain, when the machine is initiated with Bern(p) samples and run for a long enough time, eq. (114) implies that $M_t - 1$ is distributed Binomial($S-1, p$), thus the estimate $\hat{p}(M_t)$ has $\mathbb{E}(\hat{p}(M_t)) = p$ and $\mathbb{E}(\hat{p}(M_t) - p)^2 = \text{Var}(\hat{p}(M_t)) = \frac{pq}{S-1} \leq \frac{1}{S-1}$.

B. Monte Carlo Guarantee

Lemma 14: Monte Carlo simulation provides an α -additive estimation for $\mathbb{E}(\log N)$ with probability $1 - \delta$ with $L = \frac{(c \log n + 4)^2 + 1}{\alpha^2 \delta}$ samples.

Proof: Firstly, we have

$$\mathbb{E}(N) = \sum_{k=1}^{M-1} \mathbb{E}(\tau_k) = \sum_{k=1}^{M-1} 2^k \leq 2^M, \quad (119)$$

and from Jensen's inequality $\mathbb{E}(\log N) \leq \log \mathbb{E}(N) \leq M$. The proof follows from the Taylor series expansion of $\log N$.

$$\begin{aligned} \log N &= -\log \left(1 - \frac{N-1}{N} \right) \\ &= \sum_{k=1}^{\infty} \frac{1}{k} \left(\frac{N-1}{N} \right)^k \end{aligned} \quad (120)$$

$$= \sum_{k=1}^{\mathbb{E}(N)-1} \frac{1}{k} \left(1 - \frac{1}{N} \right)^k + \sum_{k=\mathbb{E}(N)}^{\infty} \frac{1}{k} \left(1 - \frac{1}{N} \right)^k \quad (121)$$

$$\leq \sum_{k=1}^{\mathbb{E}(N)-1} \frac{1}{k} + \frac{1}{\mathbb{E}(N)} \sum_{k=0}^{\infty} \left(1 - \frac{1}{N} \right)^k \quad (122)$$

$$\leq \log(\mathbb{E}(N)) + 1 + \frac{N}{\mathbb{E}(N)}. \quad (123)$$

Thus we can write

$$\begin{aligned} \log^2(N) &\leq \log^2(2\mathbb{E}(N)) \\ &\quad + \frac{2N}{\mathbb{E}(N)} \cdot \log(2\mathbb{E}(N)) + \frac{N^2}{\mathbb{E}^2(N)}, \end{aligned} \quad (124)$$

and, taking expectation, we have

$$\begin{aligned} \mathbb{E}(\log^2(N)) &\leq \log^2(2\mathbb{E}(N)) + 2\log(2\mathbb{E}(N)) + \frac{\text{Var}(N) + \mathbb{E}^2(N)}{\mathbb{E}^2(N)} \end{aligned} \quad (125)$$

$$\leq (M+1)^2 + 2(M+1) + 2 = (M+2)^2 + 1. \quad (126)$$

Finally, from Chebyshev's inequality, we have

$$\begin{aligned} \Pr \left(\left| \frac{1}{L} \sum_{j=1}^L \log(N_j) - \mathbb{E}(\log N) \right| > \alpha \right) &\leq \frac{\text{Var}(\log N)}{L\alpha^2} \end{aligned} \quad (127)$$

$$\leq \frac{\mathbb{E}(\log^2(N))}{L\alpha^2} \quad (128)$$

$$\leq \frac{(c \log n + 4)^2 + 1}{L\alpha^2}, \quad (129)$$

thus taking $L = \frac{(c \log n + 4)^2 + 1}{\alpha^2 \delta}$ achieves the result.

C. Proof of Lemma 7

From total probability

$$\begin{aligned} \mathbb{E}(C_{N_X}^{\infty}) &= \Pr(C_{N_X}^{\infty} < 2M) \mathbb{E}(C_{N_X}^{\infty} | C_{N_X}^{\infty} < 2M) \\ &\quad + \Pr(C_{N_X}^{\infty} \geq 2M) \mathbb{E}(C_{N_X}^{\infty} | C_{N_X}^{\infty} \geq 2M) \end{aligned} \quad (130)$$

$$\begin{aligned} &= \Pr(C_{N_X}^{\infty} < 2M) \mathbb{E}(C_{N_X}) \\ &\quad + \Pr(C_{N_X}^{\infty} \geq 2M) \mathbb{E}(C_{N_X}^{\infty} | C_{N_X}^{\infty} \geq 2M), \end{aligned} \quad (131)$$

which implies

$$\begin{aligned} \mathbb{E}(C_{N_X}^{\infty}) - \Pr(C_{N_X}^{\infty} \geq 2M) \mathbb{E}(C_{N_X}^{\infty} | C_{N_X}^{\infty} \geq 2M) \\ \leq \mathbb{E}(C_{N_X}) \leq \frac{\mathbb{E}(C_{N_X}^{\infty})}{\Pr(C_{N_X}^{\infty} < 2M)}. \end{aligned} \quad (132)$$

Note that

$$\Pr(C_{N_X}^{\infty} = k | C_{N_X}^{\infty} \geq 2M) = \frac{\Pr(C_{N_X}^{\infty} = k)}{\Pr(C_{N_X}^{\infty} \geq 2M)} \quad (133)$$

for $k \geq 2M$ and zero otherwise. Hence,

$$\begin{aligned} \Pr(C_{N_X}^{\infty} \geq 2M) \mathbb{E}(C_{N_X}^{\infty} | C_{N_X}^{\infty} \geq 2M) \\ = \sum_{k=2M}^{\infty} k \Pr(C_{N_X}^{\infty} = k). \end{aligned} \quad (134)$$

We thus have the upper bound

$$\begin{aligned} |\mathbb{E}(C_{N_X}) - \mathbb{E}(C_{N_X}^{\infty})| &\leq \\ \max \left\{ \sum_{k=2M}^{\infty} k \Pr(C_{N_X}^{\infty} = k), \frac{\mathbb{E}(C_{N_X}^{\infty}) \Pr(C_{N_X}^{\infty} \geq 2M)}{\Pr(C_{N_X}^{\infty} < 2M)} \right\}. \end{aligned} \quad (135)$$

We first bound the second term. Write

$$\mathbb{E}(C_{N_X}^{\infty}) = \mathbb{E}(\log N_X) + \mu + \mathbb{E}(\gamma_{N_X}) \quad (136)$$

$$\leq \log(\mathbb{E}(N_X)) + \mu + 1 + 10^{-5} \quad (137)$$

$$= \log(\mathbb{E}(N \cdot p_X)) + \mu + 1 + 10^{-5} \quad (138)$$

$$< c \log n + 3, \quad (139)$$

where (136) follows from Theorem 4 and the smoothing theorem, (137) follows from Lemma 6 and Jensen's inequality, and (139) follows as $\mu \approx -0.3$ and from $\mathbb{E}(N) = \sum_{k=1}^{M-1} 2^k = 2^M - 2 \leq 4n^c$. Now, note that if $C_{N_X}^{\infty} \geq 2M$ after $N = m$ samples, then the second counter must have moved from state $2M - 1$ to state $2M$ in less than m steps, implying that

$$\Pr(C_{N_X}^{\infty} \geq 2M | N = m) \leq m \cdot 2^{-(2M-1)}. \quad (140)$$

Since $2^{-M} \leq 1/(2n^c)$, we have

$$\Pr(C_{N_X}^{\infty} \geq 2M) = \mathbb{E}(\Pr(C_{N_X}^{\infty} \geq 2M | N)) \quad (141)$$

$$\leq \mathbb{E}(N \cdot 2^{-(2M-1)}) \quad (142)$$

$$\leq 2n^{-c}. \quad (143)$$

Combining the above, we get

$$\frac{\mathbb{E}(C_{N_X}^{\infty}) \Pr(C_{N_X}^{\infty} \geq 2M)}{\Pr(C_{N_X}^{\infty} < 2M)} \leq \frac{2(c \log n + 3)}{n^c(1 - 2n^{-c})}. \quad (144)$$

□ We now proceed to carefully bound $\sum_{k=2M}^{\infty} k \Pr(C_{N_X}^{\infty} = k)$. For any $X = x$ we have

$$\begin{aligned} &\sum_{k=2M}^{\infty} k \Pr(C_{N_X}^{\infty} = k) \\ &= \sum_{m=1}^{\infty} \Pr(N = m) \sum_{n_x=1}^m \Pr(N_x = n_x | N = m) \\ &\quad \times \sum_{k=2M}^{\infty} \Pr(C_{N_X}^{\infty} = k | N = m, N_x = n_x) \end{aligned} \quad (145)$$

$$\begin{aligned}
&= \sum_{m=1}^{\infty} \Pr(N = m) \sum_{n_x=1}^m \Pr(N_x = n_x \mid N = m) \\
&\times \sum_{k=2M}^{\infty} \Pr(C_{N_x}^{\infty} = k \mid N = m). \tag{146}
\end{aligned}$$

We divide the computation into sample-state blocks, where each sample block is of length $\lceil 4n^c \rceil$ and each state block is of length M . Clearly $\Pr(N \in [\ell \cdot \lceil 4n^c \rceil, (\ell + 1) \cdot \lceil 4n^c \rceil]) \leq \Pr(N \geq \ell \cdot \lceil 4n^c \rceil)$, and $k \leq \alpha M$ in the interval $[(\alpha - 1)M, \alpha M)$. Thus,

$$\begin{aligned}
&\sum_{k=2M}^{\infty} k \Pr(C_{N_x}^{\infty} = k) \\
&\leq \sum_{m=1}^{\infty} \Pr(N = m) \sum_{n_x=1}^m \Pr(N_x = n_x \mid N = m) \\
&\times \sum_{\alpha=2}^{\infty} (\alpha + 1)M \max_{\alpha M \leq k \leq (\alpha+1)M} \Pr(C_{N_x}^{\infty} = k \mid N = m) \tag{147} \\
&\leq \sum_{\ell=0}^{\infty} \Pr(N \geq \ell \cdot \lceil 4n^c \rceil) \sum_{\alpha=2}^{\infty} (\alpha + 1)M \\
&\times \max_{\substack{\ell \cdot \lceil 4n^c \rceil \leq m \leq (\ell+1) \cdot \lceil 4n^c \rceil \\ 1 \leq n_x \leq m \\ \alpha M \leq k \leq (\alpha+1)M}} \Pr(C_{N_x}^{\infty} = k \mid N = m). \tag{148}
\end{aligned}$$

First, we have from Lemma 2 that $\Pr(N \geq \ell \cdot \lceil 4n^c \rceil) \leq 5e^{-\ell}$. Now, note that if $C_{N_x}^{\infty} = k$, then the second counter must have moved from state $k - 1$ to state k in less than n_x steps. Thus we have

$$\begin{aligned}
&\max_{\substack{\ell \cdot \lceil 4n^c \rceil \leq m \leq (\ell+1) \cdot \lceil 4n^c \rceil \\ 1 \leq n_x \leq m \\ \alpha M \leq k \leq (\alpha+1)M}} \Pr(C_{N_x}^{\infty} = k \mid N = m) \\
&\leq \max_{\substack{\ell \cdot \lceil 4n^c \rceil \leq m \leq (\ell+1) \cdot \lceil 4n^c \rceil \\ 1 \leq n_x \leq m \\ \alpha M \leq k \leq (\alpha+1)M}} n_x \cdot 2^{-(k-1)} \tag{149}
\end{aligned}$$

$$\leq \max_{\substack{\ell \cdot \lceil 4n^c \rceil \leq m \leq (\ell+1) \cdot \lceil 4n^c \rceil \\ \alpha M \leq k \leq (\alpha+1)M}} m \cdot 2^{-(k-1)} \tag{150}$$

$$\leq (\ell + 1) \cdot \lceil 4n^c \rceil \cdot 2^{-(\alpha M - 1)}. \tag{151}$$

Plugging back the above, we have

$$\begin{aligned}
&\sum_{k=2M}^{\infty} k \Pr(C_{N_x}^{\infty} = k) \\
&\leq 10M \cdot \lceil 4n^c \rceil \left(\sum_{\ell=0}^{\infty} (\ell + 1)e^{-\ell} \right) \left(\sum_{\alpha=2}^{\infty} (\alpha + 1)2^{-\alpha M} \right) \tag{152}
\end{aligned}$$

$$\leq 10M \cdot \lceil 4n^c \rceil \cdot \frac{1}{(1 - e^{-1})^2} \cdot \frac{3 \cdot 2^{-2M}}{(1 - 2^{-M})^2} \tag{153}$$

$$< \frac{100(c \log n + 2)}{n^c(1 - 0.5n^{-c})^2}, \tag{154}$$

where we used the identity

$$\sum_{n=N_1}^{\infty} nq^{n-1} = \frac{N_1 q^{N_1-1} - (N_1 - 1)q^{N_1}}{(1 - q)^2} < \frac{N_1 q^{N_1-1}}{(1 - q)^2}. \tag{155}$$

thus, overall,

$$|\mathbb{E}(C_{N_x}) - \mathbb{E}(C_{N_x}^{\infty})| \leq \frac{100(c \log n + 2)}{n^c(1 - 0.5n^{-c})^2}. \tag{156}$$

REFERENCES

- [1] J. Acharya, S. Bhadane, P. Indyk, and Z. Sun, "Estimating entropy of distributions in constant space," 2019, *arXiv:1911.07976*.
- [2] M. Aliakbarpour, A. McGregor, J. Nelson, and E. Waingarten, "Estimation of entropy in constant space with improved sample complexity," 2022, *arXiv:2205.09804*.
- [3] R. Morris, "Counting large numbers of events in small registers," *Commun. ACM*, vol. 21, no. 10, pp. 840–842, Oct. 1978.
- [4] F. Leighton and R. Rivest, "Estimating a probability using finite memory," *IEEE Trans. Inf. Theory*, vol. IT-32, no. 6, pp. 733–742, Nov. 1986.
- [5] T. Berg, O. Shayevitz, and O. Ordentlich, "Deterministic finite-memory bias estimation," in *Proc. Conf. Learn. Theory*, Jul. 2021, pp. 566–585.
- [6] A. Chakrabarti, K. Do Ba, and S. Muthukrishnan, "Estimating entropy and entropy norm on data streams," *Internet Math.*, vol. 3, no. 1, pp. 63–78, Jan. 2006.
- [7] T. Berg, O. Ordentlich, and O. Shayevitz, "On the memory complexity of uniformity testing," in *Proc. Conf. Learn. Theory*, Jan. 2022, pp. 3506–3523.
- [8] T. M. Cover, "Hypothesis testing with finite statistics," *Ann. Math. Statist.*, vol. 40, no. 3, pp. 828–835, Jun. 1969.
- [9] M. E. Hellman and T. M. Cover, "Learning with finite memory," *Ann. Math. Statist.*, vol. 41, no. 3, pp. 765–782, 1970.
- [10] F. Samanigo, "Estimating a binomial parameter with finite memory," *IEEE Trans. Inf. Theory*, vol. IT-19, no. 5, pp. 636–643, Sep. 1973.
- [11] J. Steinhardt and J. C. Duchi, "Minimax rates for memory-bounded sparse linear regression," in *Proc. Conf. Learn. Theory*, Jun. 2015, pp. 1564–1587.
- [12] J. Steinhardt, G. Valiant, and S. Wager, "Memory, communication, and statistical queries," in *Proc. Conf. Learn. Theory*, 2016, pp. 1490–1516.
- [13] R. Raz, "Fast learning requires good memory: A time-space lower bound for parity learning," *J. ACM*, vol. 66, no. 1, p. 3, 2018.
- [14] Y. Dagan and O. Shamir, "Detecting correlations with little memory and communication," in *Proc. Conf. Learn. Theory*, 2018, pp. 1145–1198.
- [15] A. Jain and H. Tyagi, "Effective memory shrinkage in estimation," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2018, pp. 1071–1075.
- [16] Y. Dagan, G. Kur, and O. Shamir, "Space lower bounds for linear prediction in the streaming model," in *Proc. Conf. Learn. Theory*, Feb. 2019, pp. 929–954.
- [17] V. Sharan, A. Sidford, and G. Valiant, "Memory-sample tradeoffs for linear regression with small error," in *Proc. 51st Annu. ACM SIGACT Symp. Theory Comput.*, Jun. 2019, pp. 890–901.
- [18] S. Garg, P. K. Kothari, P. Liu, and R. Raz, "Memory-sample lower bounds for learning parity with noise," in *Proc. 24th Int. Conf. Approximation Algorithms Combinat. Optim. Problems, APPROX 25th Int. Conf. Randomization Comput.*, Jan. 2021, p. 60.
- [19] A. Pensia, V. Jog, and P.-L. Loh, "Communication-constrained hypothesis testing: Optimality, robustness, and reverse data processing inequalities," 2022, *arXiv:2206.02765*.
- [20] G. P. Basharin, "On a statistical estimate for the entropy of a sequence of independent random variables," *Theory Probab. Appl.*, vol. 4, no. 3, pp. 333–336, Jan. 1959.
- [21] A. Antos and I. Kontoyiannis, "Convergence properties of functional estimates for discrete distributions," *Random Struct. Algorithms*, vol. 19, nos. 3–4, pp. 163–193, 2001.
- [22] L. Paninski, "Estimation of entropy and mutual information," *Neural Comput.*, vol. 15, no. 6, pp. 1191–1253, Jun. 2003.
- [23] L. Paninski, "Estimating entropy on m bins given fewer than m samples," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 2200–2203, Sep. 2004.
- [24] G. Valiant and P. Valiant, "A clt and tight lower bounds for estimating entropy," *Electron. Colloq. Comput. Complex.*, vol. 17, p. 179, Nov. 2010.
- [25] G. Valiant and P. Valiant, "Estimating the unseen: An $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs," in *Proc. 43rd Annu. ACM Symp. Theory Comput.*, Jun. 2011, pp. 685–694.
- [26] P. Valiant and G. Valiant, "Estimating the unseen: Improved estimators for entropy and other properties," in *Proc. NIPS*, vol. 26, Dec. 2013, pp. 2157–2165.
- [27] G. Valiant and P. Valiant, "The power of linear estimators," in *Proc. IEEE 52nd Annu. Symp. Found. Comput. Sci. (FOCS)*, Oct. 2011, pp. 403–412.
- [28] J. Jiao, K. Venkat, Y. Han, and T. Weissman, "Minimax estimation of functionals of discrete distributions," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2835–2885, May 2015.

- [29] Y. Wu and P. Yang, "Minimax rates of entropy estimation on large alphabets via best polynomial approximation," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3702–3720, Jun. 2016.
- [30] A. Lall, V. Sekar, M. Ogihara, J. Xu, and H. Zhang, "Data streaming algorithms for estimating entropy of network traffic," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 34, no. 1, pp. 145–156, Jun. 2006.
- [31] S. Guha, A. McGregor, and S. Venkatasubramanian, "Sublinear estimation of entropy and information distances," *ACM Trans. Algorithms*, vol. 5, no. 4, pp. 1–16, Oct. 2009.
- [32] S. Chien, K. Ligett, and A. McGregor, "Space-efficient estimation of robust statistics and distribution testing," in *Proc. ICS*, Jan. 2010, pp. 251–265.
- [33] T. Berg, O. Ordentlich, and O. Shayevitz, "Statistical inference with limited memory: A survey," *IEEE J. Sel. Areas Inf. Theory*, vol. 5, pp. 623–644, 2024.
- [34] P. Flajolet, "Approximate counting: A detailed analysis," *BIT Numer. Math.*, vol. 25, no. 1, pp. 113–134, Mar. 1985.
- [35] J. Nelson and H. Yu, "Optimal bounds for approximate counting," 2020, *arXiv:2010.02116*.
- [36] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times*, vol. 107. Providence, RI, USA: American Mathematical Soc., 2017.
- [37] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2012.

Tomer Berg received the B.Sc. degree (summa cum laude) in biomedical engineering from the Ben Gurion University of the Negev in 2010 and the M.Sc. degree (magna cum laude) in electrical engineering from Tel Aviv University in 2019, where he is currently pursuing the Ph.D. degree in electrical engineering. His research interests include problems in information theory, statistics, and communication. He was a recipient of the 2020 Weinstein Award for Excellence in graduate Studies, the 2020 ISIT Jack Keil Wolf Student Paper Award, the 2021 David Burshtein Paper Award, the 2022 Weinstein Paper Award, and the 2023 Feder Award for best student work in communications technology for works on Inference under Memory Constraints.

Or Ordentlich received the B.Sc. (cum laude), M.Sc. (summa cum laude), and Ph.D. degrees in electrical engineering from Tel Aviv University, Israel, in 2010, 2011, and 2016, respectively. From 2015 to 2017, he was a Post-Doctoral Fellow with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology (MIT), and the Department of Electrical and Computer Engineering, Boston University. He is currently an Associate Professor with the School of Computer Science and Engineering, Hebrew University of Jerusalem. His research interests include information theory and its application to modern problems in communication, compression, and data science. He has been serving as an Associate Editor for *Signal Processing and Source Coding* in IEEE TRANSACTIONS ON INFORMATION THEORY since 2021.

Ofer Shayevitz (Senior Member, IEEE) received the B.Sc. degree from the Technion Institute of Technology, Haifa, Israel, and the M.Sc. and Ph.D. degrees from Tel-Aviv University, Tel Aviv-Yafo, Israel, all in electrical engineering. He is currently an Associate Professor with the Department of EE—Systems, Tel Aviv University. Before joining Tel Aviv University, he was a Post-Doctoral Fellow with the Information Theory and Applications (ITA) Center, University of California, San Diego, and a Quantitative Analyst with the D.E. Shaw Group, New York. Prior to his graduate studies, he was a Digital Communication Engineer and the Team Leader of Israeli Defense Forces, and worked in statistical signal processing in Israeli hi-tech industry. He is also actively involved in Israeli hi-tech industry and regularly consults to various companies. His research interests include wide cross-section of problems in information theory, data science, statistical inference, and discrete mathematics. He was a recipient of various research grants, including European Research Council (ERC) Starting Grant, several Israel Science Foundation (ISF) grants, and the Marie Curie Grant.