

Statistical Inference with Limited Memory: A Survey

Tomer Berg, Or Ordentlich, Ofer Shayevitz

Abstract

The problem of statistical inference in its various forms has been the subject of decades-long extensive research. Most of the effort has been focused on characterizing the behavior as a function of the number of available samples, with far less attention given to the effect of memory limitations on performance. Recently, this latter topic has drawn much interest in the engineering and computer science literature. In this survey paper, we attempt to review the state-of-the-art of statistical inference under memory constraints in several canonical problems, including hypothesis testing, parameter estimation, and distribution property testing/estimation. We discuss the main results in this developing field, and by identifying recurrent themes, we extract some fundamental building blocks for algorithmic construction, as well as useful techniques for lower bound derivations.

1 Introduction

Statistical inference is one of the most prominent pillars of modern science and engineering. Generally speaking, a statistical inference problem is the task of reliably estimating some property of interest from random samples, obtained from a partially or completely unknown distribution. Such properties can include for example the mean of the distribution, its higher moments, whether the distribution belongs to a certain class (e.g., is uniform or Gaussian) or is far from that class, the entropy of the distribution, or even the distribution itself. These types of problems arise constantly in many research areas, both theoretical and practical, and across multiple disciplines in engineering, computer science, physics, economics, biology, and many other fields and subfields within.

The extremely broad literature on statistical inference has been concerned almost exclusively with an idealized setting, in which no practical constraints are posed on the inference algorithm or the data collection process. Specifically, two implicit postulates underlying this classical setup are that inference is completely *centralized*, namely that all the samples are available to a single agent who carries

out the inference procedure, and that this agent has essentially *unlimited computational power*, hence can implement the optimal decision rule. These crucial assumptions, however, often fail to hold in real world settings, where data is either collected in a distributed fashion by multiple remotely-located agents who are subject to stringent communication constraints, or by a single agent who has limited computational capabilities (or both).¹ If the number of collected samples is very small relative to the computational power or the available communication bandwidth, then the computationally-limited / distributed setup can be easily reduced essentially to a computationally unlimited / centralized one with little overhead. This is however not possible when either communication or computation / memory become a bottleneck.

As a simple example, imagine a network of low-complexity remote sensors sampling their local environments and reporting back over a low-bandwidth channel to a base station, whose goal is to monitor irregular behavior in the network. The amount of local data observed by each sensor might be quite large (e.g., video, internet traffic, etc.) and cannot be fully processed / locally stored, nor completely conveyed to the base-station in time to facilitate fast anomaly detection. To meet their goal, the sensors would therefore need to send a much shorter and partial description of their samples. What is the shortest description they can send that still allows the base station to achieve its desired level of minimax risk? What is the effect of limited computational power on the attainable risk? Can these limits be efficiently achieved, and if so, how? This types of timely questions have gained much contemporary interest.

In this survey, we focus on the computational aspect of such problems, and in particular on the limited memory aspect, where an inference algorithm is restricted to store only a small number of bits in memory. The study of learning and estimation under memory constraints has been initiated in the late 1960s by Cover and Hellman [4, 5] (with a precursor by Robbins [6] on the two armed bandit problem), and remained an active research area for a few decades. It has then been largely abandoned, but has recently drawn much attention again, with many works addressing different aspects of inference under memory constraints in the last few years. Despite this ongoing effort, there are still gaps in our understanding of the effects that memory limitation has on inference accuracy in various problems, as well as its impact on the required sample size. In this survey, we attempt to provide a snapshot overview of this fast-evolving area, summarizing the main results in the field and the landscape around it. We also delineate a few high level ideas and techniques often used to obtain upper and lower bounds in this type of problems, point out some interesting phenomena and trade-offs, and discuss open problems.

¹Another common constraint is privacy, and in particular local differential privacy, where the central server is not allowed to learn much about the samples of any individual agent. We refer the interested reader to a survey and a couple of state of the art works on this topic [1, 2, 3].

Outline. In Section 2 we formally define the general inference setting, present the different types of inference problems we will encounter in the paper, and introduce the setting of inference under finite-state memory constraints. In Section 3, we discuss some assumptions on finite-state algorithms, including time-invariance and randomization. Section 4 is a short digression from the main theme of the paper, where we briefly survey the closely-related problems of inference in the data stream model, and inference under communication constraints. Section 5 introduces the main methodologies used in deriving lower bounds and upper bounds in the memory constrained setting. Sections 6 thru 9 give a thorough overview of the landscape and current status of knowledge in memory constrained hypothesis testing, parameter estimation, distribution testing and distribution property estimation problems. Finally, Section 10 addresses open problems and future research directions.

2 Problem Setting and Definitions

2.1 The Inference Problem

In classical statistical inference, one is often concerned with the following prototypical problem. Nature picks some unknown probability distribution P from a given family \mathcal{F} of possible distributions, and an agent wants to estimate some property $\pi(P) \in \mathcal{Y}$ of the distribution. The property space \mathcal{Y} can be taken in most cases to be some subset of \mathbb{R}^d for some d . To that end, the agent observes n i.i.d samples $X^n = (X_1, \dots, X_n)$ from P and computes an estimator $\hat{\pi} : \mathcal{X}^n \rightarrow \mathcal{Y}$ for $\pi(P)$. The quality of estimation is measured w.r.t some loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, which gives rise to the notion of the *expected risk*, defined as

$$R_n(P, \hat{\pi}) \triangleq \mathbb{E}[\ell(\pi(P), \hat{\pi}(X^n))]. \quad (1)$$

The agent's goal is to find an estimator $\hat{\pi}$ that is optimal in some sense. Various notions of optimality appear in the literature, including those based on the *Bayesian paradigm* and on the *unbiased estimation* approach. Here, we focus on the more robust *minimax paradigm*, since it makes no assumptions on the behavior of the data or the structure of the estimator. Under this paradigm, the agent chooses the estimator π_n^* that guarantees the minimal possible risk in the worst case, i.e., attaining the so-called *minimax risk* of the problem, defined as

$$R_n^* \triangleq \inf_{\hat{\pi}} \sup_{P \in \mathcal{F}} R_n(P, \hat{\pi}). \quad (2)$$

Most of the works in the machine learning literature have characterized the statistical hardness of the inference problem by the *minimax sample complexity*, which is the minimal number of samples n for which $R_n^* \leq \delta$ can be guaranteed uniformly

for all $P \in \mathcal{F}$, for some fixed target risk $\delta > 0$, i.e.,

$$\text{SC}(\delta) = \min\{n : R_n^* \leq \delta\}. \quad (3)$$

The theory of minimax inference is well understood in many statistical inference settings, and here is a short, non exhaustive list [7, 8, 9, 10, 11, 12, 13].

2.2 Problem Types

Let us briefly recall some inference problems that have been extensively studied in the traditional (computationally-unbounded) setting. We later revisit these problems under our computation restricted paradigms.

2.2.1 Simple Hypothesis Testing

In this problem, the samples are drawn from a member in a discrete family of distributions, $\mathcal{F} = \{P_0, \dots, P_{d-1}\}$ and we are interested in guessing the correct underlying distribution $\pi(P_i) = i$. Some common problems are ascertaining whether the samples come from a particular population (say, infected/healthy), or identifying whether they are generated by a signal suffused in noise in contrast to originating from pure noise (the RADAR problem). The frequently used loss function in this case is the 0/1 loss, $\ell(\pi, \hat{\pi}) = \mathbb{1}(\pi \neq \hat{\pi})$, which gives rise to the expected risk

$$\mathbb{E}[\ell(\pi, \hat{\pi})] = \mathbb{E}[\mathbb{1}(\pi \neq \hat{\pi})] = \Pr(\pi \neq \hat{\pi}), \quad (4)$$

which is precisely the error probability of the algorithm. The minimax risk is hence the smallest error probability that can be uniformly guaranteed for all hypotheses in \mathcal{F} . A secondary line of work, mostly appearing in the information theory literature, is interested in the decay rate of $\Pr(\pi \neq \hat{\pi})$ as a function of the number of samples n , and, particularly, in the optimal *error exponent* for Neyman–Pearson testing, which is the decay rate of the type II error probability subject to the constraint that the type I error probability is at most some specified constant. For a more symmetric approach, one may consider the Bayesian setting, in which prior probabilities are assigned to both hypotheses and we want to minimize the average error probability, or the minimax setting, in which we want to minimize the maximal error probability. The best achievable error exponent in both cases is the *Chernoff information* or *Chernoff Divergence* [14, 15, 16].

2.2.2 Parameter Estimation

In this setup, the family of distributions is parameterized $\mathcal{F} = \{P_\theta\}_{\theta \in \mathbb{R}^d}$, and we would like to approximate the value of the parameter itself, i.e., $\pi(P_\theta) = \theta$. Some common problems are estimating the bias of a coin and location estimation

in sensor networks. The loss function of choice here is the squared error loss, $\ell(\hat{\theta}, \theta) = (\hat{\theta} - \theta)^2$, which brings rise to the expected risk

$$\mathbb{E}[\ell(\hat{\theta}, \theta)] = \mathbb{E}[(\hat{\theta} - \theta)^2], \quad (5)$$

known as the *quadratic risk*. Thus we are looking for the *minimax quadratic risk*, also referred to as minimum mean squared error.

2.2.3 Composite Hypothesis Testing

In this more general problem the family of distributions is a finite disjoint union of sub-families $\mathcal{F} = \cup_{i=1}^{d-1} \mathcal{F}_i$, and we want to find out whether the underlying distribution belongs to one of the sub-families. The loss is again the 0/1 loss thus the expected risk is the minimax error probability. Of a particular interest to us is the *distribution property testing* variant, on which we elaborate below.

2.2.4 Distribution Property Testing

This is a special case of composite hypothesis testing, where we are interested in deciding whether the underlying distribution has a certain property or is ε -far in total variation from having this property. Common problem are testing whether the distribution is uniform, sparse, or has entropy below some threshold value. There are two composite hypotheses in this problem, where \mathcal{F}_0 contains all the distributions with the property, and \mathcal{F}_1 contains all the distributions P that satisfy

$$\inf_{Q \in \mathcal{F}_0} d_{\text{TV}}(P, Q) > \varepsilon,$$

where $d_{\text{TV}}(P, Q)$ is the total variation distance between P and Q , which is defined as half the ℓ_1 distance, i.e., $d_{\text{TV}}(P, Q) = \frac{\|P-Q\|_1}{2}$.

2.2.5 Distribution Property Estimation

The family of distributions is again parameterized by $\mathcal{F} = \{P_\theta\}_{\theta \in \mathbb{R}^d}$, but here we are trying to learn some “nice” property of the distribution induced by θ , e.g., some simple scalar function $h(P_\theta)$ instead of θ itself. Common problems are approximating the Shannon entropy of a distribution, or approximating the support size (also known as the distinct element problem). The loss function is often taken to be the ℓ_1 distance, which gives rise to the ℓ_1 risk,

$$\mathbb{E}[\ell(\hat{h}, h)] = \mathbb{E}|\hat{h} - h|. \quad (6)$$

In the machine learning literature works dealing with this problem, the loss function is sometimes taken to be the 0 – 1 loss w.r.t. an ε -error event in ℓ_1 , defined

as

$$\ell(\hat{h}, h) = \mathbb{1}(|\hat{h} - h| > \varepsilon), \quad (7)$$

where $\varepsilon > 0$ is a given design parameter. Namely, here the risk is the probability that our estimator is too far from the ground truth, i.e., $\mathbf{P}_e = \Pr(|\hat{h} - h| > \varepsilon)$. This is also known as the Probably Approximately Correct (PAC) setting [17].

2.3 Memory Constraints

In the limited-memory framework, an agent observes an i.i.d sequence $X_1, X_2, \dots \in \mathcal{X}$, and at each time point is allowed to update a memory register, or equivalently, a memory state. Throughout the survey, we will use the memory states terminology to refer to a memory limited algorithm, as it is consistent with many of the earlier works and some of the contemporary ones. Specifically, let $\mathbf{FSM}(S, \mathcal{X}, \mathcal{Y})$ denote the family of all *finite-state machines* over the state space $[S] = \{1, 2, \dots, S\}$, with input space \mathcal{X} and output space \mathcal{Y} . We will usually write $\mathbf{FSM}(S)$ for short, leaving the input and output spaces implicit. A finite-state machine, or finite-state algorithm, is a pair (f, g) of a *state-transition function* $f : [S] \times \mathcal{X} \rightarrow [S]$, and a *state-decision function* $g : [S] \rightarrow \mathcal{Y}$. Without loss of generality, we can assume the finite-state machine is initialized to state 1 at time $n = 0$. When fed with the input sequence X_1, X_2, \dots , the state of the machine evolves as follows:

$$T_0 = 1 \quad (8)$$

$$T_n = f(T_{n-1}, X_n). \quad (9)$$

Correspondingly, the state-decision at time n , which in our problem yields the estimator $\hat{\pi}$ for π at time n , is given by

$$\hat{\pi}(X^n) = g(T_n). \quad (10)$$

With a slight abuse of notation, we will write $\hat{\pi} \in \mathbf{FSM}(S)$ to mean that our estimator is of the above structure, i.e., is obtained as the state-decision of some finite-state machine with S states. We will refer to $\hat{\pi}$ as finite-state estimation algorithm, or sometimes for brevity, simply a finite-state estimator.

We emphasize that the state-transition and state-decision rules (f, g) we consider here are *time-invariant*, i.e., they are fixed and do not depend on the time parameter n . It is also possible to consider time-varying machines, namely to let (f_n, g_n) depend on n . We will discuss the time-varying case in Subsection 3.1, and explain why the restriction to time-invariant machines is reasonable.

We include *randomized* machines in our \mathbf{FSM} family, i.e., machines for which the transition function f or the decision function g depend on some exogenous randomness, independent of the samples. We will elaborate more on randomized

algorithms in Subsection 3.2, where we highlight some of the advantages randomized algorithms possess over deterministic algorithms.

In this survey, we are mainly interested in the minimax risk $R_n^*(S)$ of the memory-constrained inference problem. This risk is similar to the unconstrained minimax risk (2), with the additional structural constraint requiring the algorithm producing the estimator $\hat{\pi}$ to be finite-state. Formally:

$$R_n^*(S) = \inf_{\hat{\pi} \in \text{FSM}(S)} \sup_{P \in \mathcal{F}} R_n(P, \hat{\pi}). \quad (11)$$

In many of the works covered in this survey, one is interested in the asymptotic behaviour of a finite-state algorithm, that is, in the risk attained with fixed memory in the limit of many samples $n \rightarrow \infty$. To that end, we define the *asymptotic risk* of an estimator to be

$$R(P, \hat{\pi}) = \limsup_{n \rightarrow \infty} R_n(P, \hat{\pi}), \quad (12)$$

Note that using \limsup is required in order to cover FSM algorithms that have periodic components, for which the regular limit does not exist. In some of the works covered in the survey, this issue is circumvented using an alternative definition of the asymptotic risk as the asymptotic average of the instantaneous risks:

$$\bar{R}(P, \hat{\pi}) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n R_t(P, \hat{\pi}). \quad (13)$$

We note that for our purposes, there is no restriction in using $R(P, \hat{\pi})$ instead of $\bar{R}(P, \hat{\pi})$. This follows since on the one hand, $\bar{R}(P, \hat{\pi}) \leq R(P, \hat{\pi})$, so any lower bound on $\bar{R}(P, \hat{\pi})$ is also a lower bound on $R(P, \hat{\pi})$, and on the other hand, all algorithms appearing in this survey have no periodic components, and for such algorithms both definitions coincide.

We can now define the *asymptotic minimax risk* to be

$$R^*(S) = \inf_{\hat{\pi} \in \text{FSM}(S)} \sup_{P \in \mathcal{F}} R(P, \hat{\pi}). \quad (14)$$

While in general the algorithm $\hat{\pi}$ might be randomized, in some cases we will be interested in the performance of deterministic algorithms, i.e., algorithms that do not have access to external randomness (as we elaborate on in section 3.2). We thus define an additional notion of minimax risk, which we will refer to as the *asymptotic deterministic risk*, defined as

$$R_{\text{det}}^*(S) = \inf_{\hat{\pi} \in \text{DFSM}(S)} \sup_{P \in \mathcal{F}} R(P, \hat{\pi}), \quad (15)$$

where $\text{DFSM}(S)$ is the family of finite-state *deterministic* estimation algorithms. Let $\delta^* \geq 0$ be the smallest obtainable minimax risk, i.e.,

$$\delta^* = \min\{\delta : \exists S, n \in \mathbb{N}^+, R_n^*(S) \leq \delta\}.$$

For any $\delta^* \leq \delta$, we can define the *sample-memory* curve at level δ , which is a curve on the (n, S) axis, that represents the minimal number of samples n and the minimal number of states S needed to attain a minimax risk of δ . The curve can be characterized by the function

$$S^*(n, \delta) = \min\{S : R_n^*(S) \leq \delta\}, \quad (16)$$

which gives the minimal number of states S required for the minimax risk to be smaller than δ , or, equivalently, by

$$n^*(S, \delta) = \min\{n : R_n^*(S) \leq \delta\}, \quad (17)$$

which gives the minimal number of samples n required for the minimax risk to be smaller than δ . Note that these two functions are in an inverse relation. The sample-memory curve at level δ naturally yields two asymptotic quantities: First, the classical notion of sample complexity is obtained as the minimal number of samples required to attain a risk δ in the limit of infinite (or large enough) memory:

$$\text{SC}(\delta) = \lim_{S \rightarrow \infty} n^*(S, \delta), \quad (18)$$

and second, the notion of *memory complexity*, which we define here as the smallest number of states required to attain a risk δ , in the limit of an arbitrarily large number of samples:

$$\text{MC}(\delta) = \lim_{n \rightarrow \infty} S^*(n, \delta). \quad (19)$$

Note that we define the memory complexity as the number of *states* the algorithm uses, rather than the number of *bits*; this is just a matter of convenience, since often writing results using the number of states is more natural and readable, having the same “units” as number of samples. We will nevertheless occasionally use bits for memory size, in which case we will use the term *memory complexity in bits*. Note that the memory complexity in bits never exceeds $\text{SC}(\delta) \cdot \log |\mathcal{X}|$, as this is the number of bits needed to trivially achieve a risk of δ by naively saving the first $\text{SC}(\delta)$ samples.

Characterizing the sample-memory curve in general is a very difficult task, and we are currently unaware of any non-trivial problem for which the sample-memory curve at level δ is completely known. Indeed, most of the classical research was focused on only characterizing the two extreme points of $\text{SC}(\delta)$ and $\text{MC}(\delta)$, and

only in recent years have other points on the curve been addressed (as we expand upon in later sections).

3 Algorithm Types

In this section, we discuss different assumptions that can be made on finite memory algorithms. First, we consider the dependence of the algorithms on time, and argue that it makes sense to restrict the discussion to time-invariant algorithms. We then consider randomized algorithms, i.e. ones that have access to an external source of randomness, and contrast them with deterministic algorithms, where the only randomness comes from the samples.

3.1 Time-Varying vs. Time-Invariant

A finite-memory estimation algorithm $\hat{\pi}$ is said to be *time-varying*, if either the state-transition function f , or the state-decision function g (or both) are allowed to depend the time index n . Otherwise, namely if both f and g are fixed for all n , the algorithm is called time-invariant. We note that any time-varying algorithm can be trivially reduced to a time-invariant one with a $O(\log n)$ bits overhead, if the total number of available samples n is finite, by simply keeping a clock.² Thus, in our framework, a time-varying algorithm implicitly means that we assume an external clock is provided “for free”, i.e., without accounting for the memory required to store it. Since what we care about is the memory budget, it seems less preferable to make this kind of an assumption. Moreover, when the number of samples is unbounded, this assumption can lead to absurd results. For example, one can solve the two-armed Bandit problem with only 2 states [18], and the binary hypothesis testing problem using just 3 states and with arbitrarily small error (!) [4, 19]. Finally, as we shall see, some optimal time-invariant algorithms inherently do not keep a clock at all. For all these reasons, in this survey we will restrict our attention to time-invariant algorithms.

3.2 Randomized vs. Deterministic

A randomized algorithm is an algorithm whose state-update and/or state-decision functions depend not only on the samples, but also on some external randomness independent of the samples. A natural question is in what cases randomization even helps? Note that, trivially, there is no loss of generality in assuming that the state-decision function is deterministic; this follows simply since the risk incurred

²We do not account here for the complexity of describing the state transition function f . In time-varying algorithms, this description may grow as a function of the number of observed samples.

by a random decision rule is some average over risks obtained by deterministic ones, hence one can always pick the risk-minimizing deterministic rule.

Importantly however, this argument *does not* apply to the state-transition function, since picking the particular randomness realization that minimizes the risk over the choice of transition functions will generally result in a time-varying machine, which is not in $\text{FSM}(S)$. Indeed, for time-varying algorithms it is sufficient to consider deterministic algorithms as the optimal time-varying rule is deterministic (see [20]). For time-invariant algorithms, we will see in this survey that there are certain problems where randomization improves the limiting behaviour, characterized by the asymptotic notation $O(\cdot)$, of a finite-state machine, e.g., binary hypothesis testing under the 0 – 1 loss, and others where it does not, e.g., bias estimation under the squared loss measure (as we show in the sequel).

Now, while allowing the use of randomization in the estimation procedure may certainly help, this resource has a cost. Even if one has access to unlimited randomness (which is the case, for example, when observing an unlimited number of samples, since randomness can be extracted from the i.i.d. sequence X_1, X_2, \dots), storing this randomness places a toll on one's memory budget, which is taken into account in the deterministic setup. In general, we would seek upper bounds on deterministic algorithm as they are more restrictive than randomized ones, and lower bounds for randomized algorithms, which then carry on to deterministic algorithms due to the reasons mentioned above.

4 Related Problems

We briefly comment on two models related to inference from i.i.d. samples under memory constraints. The first is the data stream model, in which no statistical assumptions are made on the data, and the algorithms are designed for the worst case. The data stream model is hence fundamentally more difficult than the inference setting addressed in this paper. The second setup we consider in this section is the communication-constrained setup. Here, samples are collected by remotely located agents, and the inference bottleneck is imposed by a limited communication budget between the agents. There are several models that capture the limited communication links, some of which are easier than the memory-constrained setup. One such example is the *blackboard model*, where each agent has access to all the messages sent by previous agents, and each message can take S values. This model is easier than our memory constrained model, in which each agent can choose its message using a finite-state machine whose input is its available observation and the message sent by the previous agent. Please note that this section is optional and can be skipped.

4.1 The Data Stream Model

The setting that we cover in this survey is closely-related to the *data stream model* in theoretical computer science, yet there are principal differences between the two. For clarity, in this section we briefly highlight the similarities and pinpoint the differences between the models.

The data stream model characterizes algorithms with small memory size that make passes over the input in order; any item not explicitly stored is inaccessible to the algorithm in the same pass. In many cases the number of passes is limited to one. The figure of merit is taken to be the memory complexity in bits, usually referred to in these works as the *space complexity* of a problem. The memory complexity in bits of many different streaming algorithms have been extensively studied within theoretical computer science, and some well known examples include estimation of frequency moments of the data stream [21, 22], estimation of Shannon and Rényi entropy of the empirical distribution of the data stream [23, 24, 25], estimation of heavy hitters [26, 27, 28], and estimation of distinct elements [29, 30, 31, 32, 33]. Though some exceptions exist [34, 35], the vast majority of works are concerned with worst case inputs (i.e., adversarially generated data), and thus lower bounds for the data stream model do not generally apply in the statistical learning setup, where the data is often assumed to be drawn i.i.d. from some unknown underlying distribution.

In general, the data stream model is substantially more difficult than the stochastic setup. To demonstrate this point, we provide a couple of examples for problems where the memory complexity required for evaluating some function of the empirical distribution of a worst case input, tends to be significantly larger than the memory complexity for evaluating the same function of the underlying distribution in a stochastic setup. Consider, for example, the problem of entropy estimation of a data stream. To provide an ε multiplicative approximation of the empirical entropy of a stream of m values in a single pass, one must use at least $\Omega(\varepsilon^{-2}/\log^2(\varepsilon))$ memory bits, according to the work of Chakrabarti et al. [25] (the lower bound is proved for the assignment $\varepsilon^{-1} = 3\sqrt{m}(\log m + 1/2)$). In contrast, when approximating the entropy of a distribution over alphabet $\mathcal{X} = [k]$ with the same multiplicative factor, only $O(\log(k/\varepsilon))$ memory bits are needed (see [36, 37]). This entails an *exponential* increase in the number of necessary memory bits w.r.t the dependence on ε .

Another example follows from the problem of approximating the second frequency moment of a sequence, also known as the *repeat rate* or *Gini's Index of homogeneity*. Woodruff [38] has shown that any algorithm that ε approximates the second moment requires at least $\Omega(1/\varepsilon^2)$ memory bits, for any $\varepsilon = \Omega(1/\sqrt{m})$, where m is the size of the universe. The statistical variant of the problem is approximating the collision probability of a distribution over alphabet $\mathcal{X} = [k]$. A simple algorithm that performs this is the following: define a collision machine to be a machine that, at each time point, outputs 1 if the current sample is the same

as the preceding sample and 0 otherwise. This can be implemented using $\log k$ bits, since we are only storing the preceding sample, and it results in a Bernoulli process with parameter $p \geq 1/k$. As we show elsewhere in the survey, this parameter can be approximated using $O(\log(k/\varepsilon))$ bits, resulting in an overall $O(\log(k/\varepsilon))$ memory bits. Thus, again, we have an exponential increase with respect to the loss parameter ε .

There are, however, some interesting rare cases in which the adversarial setting does not cost more memory than the stochastic one. A specific example arises from the field of universal prediction with finite-state machines. In a surprising result, Meron and Feder [39] showed that a K -state randomized finite-memory predictor for worst case sequences achieves a $\asymp 1/K$ redundancy w.r.t to the least square loss, which is the same redundancy it incurs in the stochastic Bernoulli setting. Another example stems from the problem of distinct elements approximation: It is known, due to the works of Indyk and Woodruff [31] and Kane et al. [33], that the memory complexity in bits of the distinct elements problem is $\asymp 1/\varepsilon^2 + \log n$, when the input length is n and ε is the approximation parameter. Woodruff [40] initiated the study of average-case complexity of the distinct elements problem under certain distributions and, in particular, considered a stochastic setup in which we observe an n -length random sequence drawn from a universe of size m , where each element is drawn uniformly and independently from an unknown subset of some unknown size d . The main contribution of the paper is to show that if $n, d \asymp 1/\varepsilon^2$, then the problem requires $\Omega(1/\varepsilon^2)$ memory bits, thus, even in the random data model, distinct elements estimation can be as memory-complex as its data stream variant.

4.2 Inference Under Communication Constraints

The topic of inference under communication constraints is not covered in this survey. Nevertheless, as it is closely related to inference under memory constraints (as we will later see), we provide a short and non-exhaustive overview of this problem. The study of statistical inference under communication constraints has originally been considered in the information theoretic literature. Parameter estimation under communication constraints was initially studied by Zhang and Berger [41], who provided a rate-distortion upper bound on the quadratic error in distributed estimation of a scalar parameter using one-way communication under a rate constraint of bits per sample. There has been much follow-up work on this type of problems, especially in the discrete alphabet case, see, in particular [42, 43, 44, 45, 46] and references therein. The closely related problem of distributed hypothesis testing against independence under communication constraints was originally treated by Ahlswede and Csiszár [47], who provided an exact asymptotic characterization of the optimal tradeoff between the rate of one-way protocols, and the Type I error exponent attained under a vanishing Type II error probability. This result has been

extended in many interesting ways, see, e.g., [48, 49, 50, 51, 52, 53, 54]. Specifically, a generalization to the interactive setup with a finite number of rounds was addressed in [55], and the associated one-way communication complexity problem was recently studied via hypercontractivity in [56].

In a different line of work [57, 58], Hadar et al. assumed that the agents have access to an arbitrarily large number of local samples, and are only limited by the number of bits they can exchange. In this setup, it is assumed that the underlying distribution is unknown, and the goal is to estimate correlations in the minimax sense. They thus define the *communication complexity* of the inference task to be the minimal number of bits required to guarantee that the minimax risk is below some constant, in the limit of large number of local samples. This definition is different from the definition of communication complexity as presented in the machine learning literature, which is the least possible number of bits needed to be interactively exchanged by Alice and Bob, such that they are able to compute some given function of their local inputs, either exactly or with high probability, see, e.g., [59, 60, 61].

Finally, there has also been much contemporary interest in distributed inference problems with communication constraints in a different context, where a finite number of i.i.d. samples from an unknown distribution are observed by multiple remotely located parties, who communicate with a data center under a communication budget constraint (one-way or interactively), to obtain an estimate of some property of interest. In these type of problems, the samples observed by the agents are from the same distribution and the regime of interest is typically where the dimension of the problem is high compared to the number of local samples, see, e.g., [62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 3, 72]. This problem bears close resemblance to remote source coding under Gaussian noise, known as the CEO problem, which have been extensively studied in the information theory literature, see, for example, [73, 74, 75] and Chapter 33.11 of Polyanskiy and Wu [76].

5 Methodologies

In this section, we provide an overview of common methodologies and ideas used in deriving performance bounds in finite-memory inference. For lower bounds, we describe the method of reduction to a communication complexity problem (which originated from works on the streaming model), the information theoretic approach, and the approach of Markov chain analysis. We then delineate some useful guidelines for upper bound derivations.

5.1 Lower Bound Techniques

Reduction to Communication Complexity

It is often possible to reduce a finite memory inference problem to a communication complexity problem, essentially replacing the bounded memory constraint by a bounded communication constraint. The high level idea is as follows. Suppose we have a S -state algorithm that can solve the inference problem given n samples, within the desired level of risk. Now, we can partition the samples into two parts a and b , giving a to Alice and b to Bob. Alice can run the algorithm on a , and send the final state to Bob using $\log S$ bits. Bob can then start his machine at Alice's last state, and continue running it on b , thereby emulating the entire inference procedure. One can now try to find some function $f(a, b)$ such that solving the inference problem also allows Bob to compute $f(a, b)$ with good probability. The communication complexity of reliably computing $f(a, b)$, i.e., the minimal number of bits Alice and Bob need to exchange to that end, is therefore a lower bound on $\log S$, and can sometimes be evaluated directly. We note that this useful reduction originates from the streaming literature, particularly from the seminal work of Alon et al. [21] on the memory complexity of approximating the frequency moments of a finite sequence, a paper that initiated the research on the memory complexity of data streams statistics approximation.

Following the footsteps of [21], many subsequent works have applied the tools of communication complexity to derive lower bounds on different problems in the streaming model, see, e.g., [31, 38, 77, 78, 79]. Applications of this reduction approach in the memory-limited statistical inference literature are few and far between. Among those are the works of Woodruff [40] on counting distinct elements when the samples are drawn from a particular family of distributions; Crouch et al. [80] on approximating the collision probability and other memory limited problems; Dagan and Shamir [81] on detecting correlations with little memory and communication; and Dinur [82] on distinguishing between sampling with or without replacement.

The Information Theoretic Approach

Information-theoretic based lower bounds for the memory constrained inference problem typically rely on lower bounding the mutual information between π , the property we want to infer, and the memory state (or, similarly, bounding the statistical distance between the algorithm output distributions for different values of π). Capturing the tradeoff between how large this mutual information can be, and the number of available memory states, can be thought of as a variant of the information bottleneck problem [83, 84]. Though there are some examples that leverage this approach directly (see, e.g., [85, 86, 87]), in general it is quite difficult to quantify the information contraction or find the strong data processing coeffi-

cient for many memory-limited problems. Thus, a more simple (yet significantly looser) information theoretic approach seeks to derive bounds for a communication constrained problem known as the *blackboard* model: This is a broadcast model of communication in which each sample is observed by a different player, and each of the players is allowed to communicate at most some constant number of bits to all other players. The communication constraints do not allow for arbitrary communication, but rather only one-pass protocols in which all information is communicated from left to right. This corresponds to a much less restrictive memory-constrained algorithm whose output at time t is a function of all memory states up to time t (not just the current memory state, as in our model), so the lower bounds for the blackboard problem carry over to our model. In the statistical learning setup, this approach was used to obtain memory-sample trade-offs by Shamir [88], who gave lower bounds on the memory needed for a few learning and estimation problems by reducing them to a canonical 'hide and seek' problem, and was also used by Steinhardt et al. [89] to obtain lower bounds for memory constrained sparse linear regression.

Analysis of Finite Markov Chains

Another approach is to leverage the fact that any finite state machine driven by an i.i.d. sequence induces a homogeneous Markov chain behaviour. Thus, for time-invariant algorithms where the number of processed samples is sufficiently large w.r.t the number of states, one can expect the chain to converge to some limiting distribution, which allows us to obtain lower bounds by showing that the dependence of that limiting distribution on the property we want to infer is not too strong. We note that this approach breaks down for a time-varying algorithm, since the number of states grows with time and hence analyzing the limiting distribution becomes intractable.

An early example of this method appears in the seminal paper of Hellman and Cover [5] on memory-limited binary hypothesis testing. There, the authors showed that the state likelihood ratio between the stationary distributions under the two hypotheses cannot change too quickly between states, and leveraged this to obtain a lower bound on the achievable error probability. A more recent example is the work of Berg et al. [90] on binary hypothesis testing with memory limited deterministic machines. There, the authors build on probabilistic pigeon-hole arguments that capture the fact that different distributions have a non-negligible probability of traversing similar Markov trajectories. In particular, they showed that for any finite-state machine, there must exist a "bottleneck" state that results in large errors of both types, and is visited with non-negligible probability under both hypotheses.

5.2 Upper Bound Techniques

Loosely speaking, to construct good upper bound, we want the Markov chain induced by our finite state algorithm to have the following high-level properties:

- “Good” limiting distribution, e.g., strongly influenced by the input distribution, so that small changes in the property we want to infer would result in large changes in the limiting distribution.
- Fast convergence to the limiting / stationary distribution, that is, short mixing time. Such behavior will guarantee that a small number of samples suffice for approaching the limiting distribution, which in turn, results in a small sample complexity.

Though we are unaware of a general recipe for algorithms that satisfy both properties above, a couple of approaches that were used in the literature may be illuminating; First, when memory is large enough (say $\gtrsim \text{SC}(\delta) \cdot \log |\mathcal{X}|$ in bits), an appealing and simple approach to deriving upper bounds is to realize the existing “natural” algorithm from the unconstrained setup in a memory limited scheme. On the other hand, when samples are abundant, the approach of waiting for extreme events has been useful in formulating order-wise optimal (or close to optimal) algorithms, as we expand on in later sections.

It is worth mentioning that, in some particular problems, essentially only the “natural” or the rare events approach to algorithmic construction needs to be considered. For a concrete example, in the seminal work of Raz [91], the author addressed the problem of parity learning over $\{0, 1\}^d$ with finite-memory algorithms, and showed that any algorithm for parity learning requires either a quadratic memory size in bits, or an exponential number of samples. Consequently, if one has $O(d^2)$ memory bits, parity learning can be solved in polynomial time by Gaussian elimination, which is the natural approach. Conversely, if one has an exponential number of samples, parity learning can be solved by waiting for all the singletons, consuming only $d + o(d)$ memory bits, which would be a rare events approach.

Waiting for Rare Events

In settings where the number of available samples is arbitrarily large, waiting has no cost, thus it makes sense to focus on capturing suitable rare events in the sample space such that, conditioned on such events, the inference task becomes easier. The memory size is then related to the rarity of the event, since we need to have sufficiently large memory in order to detect it in the first place, but describing a few rare events usually does not take much memory. This type of approach has been originally suggested by Cover [4] for the binary hypothesis testing problem, and has been studied by others since then. In the construction of Hellman and Cover [5], the finite state machine only changes a state upon observing a maximal

or minimal likelihood ratio event. When it reaches an extreme state it also waits for an extreme event, but then exits the state with small probability, due to the use of artificial randomization. The deterministic binary hypothesis testing scheme of Berg et al. [90] is based on observing a long consecutive run of 0's or a long consecutive run of 1's, which have exponentially small probabilities, but conditioned on these events, one hypothesis becomes extremely more likely than its counterpart. Specifically, both waiting for rare events schemes described above attain a risk of $2^{-\Omega(S)}$. This should be juxtaposed with using the natural statistic for this problem, the number of 1's in the sample, which leads to poor performance; the quadratic risk cannot decay faster than $2^{-\Omega(\sqrt{S})}$. This follows since 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 $S = O(k^2)$. As it turns out, the deterministic binary hypothesis testing machine that seeks long repetitions can be used as a building block to construct optimal machines in other settings (see [92, 93]), thus it seems that the rare event approach is fundamental for finite memory problems where the number of samples is very large.

Reduction to Binary Problems

The approach of breaking down a complex problem into a collection of smaller simpler tasks has always been prominent in engineering and computer science. This approach turns out to be quite useful in the memory-limited context, where reducing the complexity of the problem often results in a major reduction in memory consumption. When the number of samples is not a bottleneck, one can sometimes break down the problem into a sequence of simpler, often even binary problems, using only a small number of memory states at each stage, resulting in savings to the total memory usage. The idea is, in each step, to partition the universe into two parts, and test whether we are likely to be in one or the other. It is somewhat reminiscent to noisy binary search (see, e.g., [94, 95, 96]), with the distinction that due to memory limitations one cannot deploy the optimal search strategy.

The earliest work adopting this approach in the memory-constrained framework seems to be that of Wagner [97], who constructed a time-varying algorithm for the estimation of the bias of a coin by decomposing the problem to a sequence of binary hypothesis testing sub-problems, leveraging the results of Cover [4]. A similar approach was recently taken by Berg et al. [92], who constructed a deterministic time-invariant bias estimation algorithm that breaks down the problem to deterministic time-invariant binary hypothesis testing sub-problems. In a recent work by the same authors [93], the problem of *uniformity testing* (in which we test whether the underlying distribution is uniform or far from it) was reduced to a sequence of binary hypothesis tests as well, such that a success in all tests is very indicative of the uniform distribution hypothesis. We expand on these results in the sequel.

6 Simple Hypothesis Testing

One of the earliest appearances of a memory constraint problem in the statistical inference literature is in the hypothesis testing context. Recall that in this problem, i.i.d. samples are drawn from an unknown distribution $P \in \mathcal{F} = \{P_0, \dots, P_{d-1}\}$, where under hypothesis \mathcal{H}_i we have $P = P_i$ with $\pi(P_i) = i$, and we want to guess the correct underlying hypothesis, under the 0/1 loss, $\ell(\pi, \hat{\pi}) = \mathbb{1}(\pi \neq \hat{\pi})$. For the majority of this section we will be interested in the *binary hypothesis testing* problem, in which $\mathcal{F} = \{P_0, P_1\}$. We show that even this simple version of maybe the most basic inference task is still difficult under memory limitations. In most of the section, we will be interested in the limit of the problem when $n \rightarrow \infty$, thus following Eq. (12), we define the asymptotic probability of error of an S -state algorithm, when the true hypothesis is \mathcal{H}_i , to be

$$P_e(P_i, \hat{\pi}) = \limsup_{n \rightarrow \infty} \mathbb{E}[\mathbb{1}(\hat{\pi}(X^n) \neq \pi(P_i))] = \limsup_{n \rightarrow \infty} \Pr(\hat{\pi}(X^n) \neq i). \quad (20)$$

Our main interest is thus to characterize the *asymptotic minimax error probability*

$$P_e^*(S) \triangleq \inf_{\hat{\pi} \in \text{FSM}(S)} \sup_{P \in \{P_0, P_1\}} P_e(P, \hat{\pi}). \quad (21)$$

Interest in the limited memory hypothesis testing problem seems to have been sparked by the work of Robbins [6] on the Two-Armed Bandit problem. In this setup, a player is given two coins, with Bernoulli parameters p_1 and p_2 respectively, whose values are unknown. The player is required to maximize the long-run proportions of “heads” obtained, by successively choosing which coin to flip at any moment. Robbins proposed an algorithm that works with a limited memory of S states and attains the optimal performance as S approaches infinity. Robbins’s algorithm was later improved by Isbell [98], Smith and Pike [99], and Samuels [100], where the latter improvement arises from a randomized algorithm.

Departing from the restriction to time-invariant algorithms imposed in the works mentioned above, Cover [18] discovered a time-varying finite memory algorithm that achieves the optimal proportion $\max\{p_1, p_2\}$, asymptotically in the number of tosses, with $S = 2$. Essentially, Cover suggested to modify the problem by allowing to keep a clock, such that at each stage we know how many tosses we have made so far. Cover’s algorithm interleaves test blocks and trial blocks, where a test block is a sequence of tosses that test the hypothesis $p_1 > p_2$ vs. $p_2 > p_1$, and a trial block is a sequence of exclusive tosses of the “favorite” coin resulting from the preceding test block. The choice of the sequence of test-trial block sizes then assures that the maximum is attained. The intuitions and tools developed for the Two-Armed Bandit problem were used by Cover in a subsequent paper addressing the binary hypothesis problem [4], in which he described a time-varying finite memory machine that solves the binary hypothesis testing problem with $S = 4$

states (this was later improved to $S = 3$ by Koplowitz [19], who more generally showed that for M -ary hypothesis testing with a time-varying finite memory, $M+1$ states are necessary and sufficient to resolve the correct hypothesis with an error probability that approaches zero asymptotically). Specifically, Cover describes an algorithm with four memory states that, given samples from a $\text{Bern}(p)$ distribution, can decide between the hypothesis $\mathcal{H}_0 : p = p_0$ and $\mathcal{H}_1 : p = p_1 > p_0$ with error probability that converges to zero almost surely.³

The algorithm works as follows. We keep two binary random variables (T, Q) , where T keeps the current guess of the correct hypothesis, and Q keeps the result of the current run test, to be immediately defined. The sequence of Bernoulli samples is then divided into blocks $S_1, R_1, S_2, R_2, \dots$ with corresponding lengths $s_1, r_1, s_2, r_2, \dots$, where in each S_i block we test for a run of s_i consecutive zeros, in which case the S_i block is considered a success, and in each R_i block we test for a run of r_i consecutive ones, in which case the R_i block is considered a success. The value of Q tracks after the current run, by being set to 1 as long as the run continues and being set to 0 for the rest of the block if it breaks. Only if the test succeeded, the value of the current hypothesis guess T is updated to that of the tested hypothesis at the end of the test block, otherwise it retains its previous value. The probability of success in an S_i block is thus p^{s_i} , and the probability of success in an R_i block is $(1-p)^{r_i}$. By setting $s_i \approx \log_{p_0}(1/i)$ and, similarly, $r_i \approx \log_{1-p_0}(1/i)$, we have that $p > p_0$ implies that both $\sum p^{s_i} = \infty$ and $\sum (1-p)^{r_i} < \infty$, so from Borel–Cantelli lemma we have that $T_n \rightarrow 1$ with probability 1, and, similarly, for $p < p_0$ we have that $T_n \rightarrow 0$ with probability 1.

In the paper itself, Cover states that “*However, the dependence of f_n on n requires external specification of n if the algorithm is to be considered to have truly finite memory.*” This point is important; without an external input, a time-varying algorithm is required to keep a clock, which results in unbounded memory under the asymptotic framework. Thus, it would seem natural that, to truly understand the memory limits of the hypothesis testing problem, we must restrict ourselves to time-invariant algorithms, thus taking into account any use of clock in the memory budget.

6.1 Randomized Binary Hypothesis testing

Hellman and Cover [5] addressed the problem of binary hypothesis testing within the class of randomized time-invariant finite-state machines, and they have found an *exact* characterization of the smallest attainable error probability for this prob-

³We note that the result in the paper is stronger, solving a composite hypothesis testing problem of the form $\mathcal{H}_0 : p > p_0$ vs. $\mathcal{H}_1 : p < p_0$.

lem.⁴ Specifically, define the maximum and minimum likelihood ratios as

$$\bar{L} = \max_{A \in \mathcal{X}} \frac{P_0(A)}{P_1(A)}; \quad \underline{L} = \min_{A \in \mathcal{X}} \frac{P_0(A)}{P_1(A)}, \quad (22)$$

and define $\gamma = \frac{\bar{L}}{\underline{L}}$. It holds that

$$P_e^*(S) = \frac{1}{\gamma^{\frac{1}{2}(S-1)} + 1}. \quad (23)$$

Hellman and Cover showed that the error probability of any algorithm cannot be smaller than (23), and that, for any $\varepsilon > 0$, there exists an S -states randomized estimation algorithm that achieves error probability of $P_e^*(S) + \varepsilon$. As the optimal error probability decreases exponentially with S , we can define the asymptotics of the error exponent with regards to S as

$$\lim_{S \rightarrow \infty} -\frac{1}{S} \log P_e^*(S) = \frac{1}{2} \cdot \log \gamma. \quad (24)$$

We call the above the *randomized error exponent* of the binary hypothesis testing problem with finite memory.

Hellman and Cover [5] - Lower bound

For the lower bound, first assume that the Markov process induced by the finite-state algorithm is irreducible. There is no loss of generality in that assumption, as it is proved that the error probability of any $(S+1)$ -state estimator that induces a reducible Markov process is lower bounded by $P_e^*(S)$. For such processes, there is a unique stationary distribution induced by the input distribution. Denote by μ_i^0 (resp. μ_i^1) the stationary probability of state i in the chain, under hypothesis \mathcal{H}_0 (resp. \mathcal{H}_1), and define the state likelihood ratios (SLR) $\lambda_i \triangleq \mu_i^0 / \mu_i^1$. A low error probability intuitively implies that $\lambda_i \ll 1$ or $\lambda_i \gg 1$ for most $i \in [S]$, so to get a bound that holds for any machine we first need to show that the state likelihood ratio cannot grow too much. Specifically, assuming w.l.o.g. that the SLR's are non-decreasing, it is shown that

$$1 \leq \frac{\lambda_{i+1}}{\lambda_i} \leq \gamma. \quad (25)$$

The lower bound follows from the assumption and the upper bound follows from a straightforward analysis of Markov chain stationary distributions, using the fact

⁴In the paper, the authors discuss the Bayesian setting, in which each hypothesis is assigned a prior probability. For consistency, we present the result for the minimax setting, which coincides with their result when the prior is equiprobable.

that the most distinguishing observation is the one that maximizes the likelihood ratio, thus our state likelihood ratio cannot grow by more than γ . Let α and β be the error probabilities under hypothesis \mathcal{H}_0 and \mathcal{H}_1 , respectively. The worst case error probability is thus $P_e(S) = \max\{\alpha, \beta\}$. If λ_{\min} is the minimal state likelihood ratio, then for all $i \in [S]$, Eq. (25) implies that

$$\lambda_{\min} \leq \frac{\mu_i^0}{\mu_i^1} \leq \lambda_{\min} \cdot \gamma^{S-1}. \quad (26)$$

Let \mathcal{S}_0 denote the decision region for hypothesis \mathcal{H}_0 , i.e., the set of states in $[S]$ that are mapped to the estimate $\hat{\pi} = 0$. Similarly, let \mathcal{S}_1 denote the decision region for hypothesis \mathcal{H}_1 . Now, writing the error probability for each of the hypotheses, we get the bounds

$$\begin{aligned} \alpha &= \sum_{i \in \mathcal{S}_1} \mu_i^0 \geq \lambda_{\min} \sum_{i \in \mathcal{S}_1} \mu_i^1 = \lambda_{\min}(1 - \beta) \\ \beta &= \sum_{i \in \mathcal{S}_0} \mu_i^1 \geq \frac{1}{\lambda_{\min} \cdot \gamma^{S-1}} \sum_{i \in \mathcal{S}_1} \mu_i^0 = \frac{1}{\lambda_{\min} \cdot \gamma^{S-1}}(1 - \alpha). \end{aligned}$$

Multiplying both equations we get $\alpha\beta \geq \gamma^{-(S-1)}(1 - \alpha)(1 - \beta)$, a constraint on the achievable error probabilities under both hypotheses, and minimizing $P_e(S)$ under the constraint yields the result.

Hellman and Cover [5] - Upper bound

The upper bound construction can be understood by the following intuition: Assume that the hypotheses testing problem is dealing with the Bernoulli symmetric case, where the samples arrive either from a $\text{Bern}(p)$ or a $\text{Bern}(1 - p)$ distribution. Consider a *saturable counter* machine, which is a machine with transition rule that moves up upon seeing 1, moves down upon seeing 0, and either stays in place or moves back when it arrives at the extreme states (called saturated states). Furthermore, let the decision rule map each state to the hypothesis with the larger stationary distribution. A simple math exercise shows that the error probability we obtain decreases exponentially in the number of states, but at a rate that does not attain the lower bound described above.

The ingenious modification of Hellman and Cover is to move more probability mass to the saturated states, where an error is less likely. To achieve this, the algorithm only stays back from saturation when, in addition to the input sample pointing the other way, an external $\text{Bern}(\delta)$ random coin falls on head. The parameter δ is typically very small, and its value depends on p . Extending it to the general binary hypothesis testing problem (not necessarily Bernoulli symmetric), the machine now moves up on high likelihood-ratio events, moves down on low

likelihood-ratio events, and stays put otherwise. In the extreme state corresponding to hypothesis \mathcal{H}_0 , the process moves back upon seeing a low likelihood-ratio event with probability δ , chosen according to γ . In the extreme state corresponding to hypothesis \mathcal{H}_1 , the process moves back upon seeing a high likelihood-ratio event with probability $k \cdot \delta$, where k is chosen to account for the asymmetry between the hypotheses (see Figure 1). This defines a class of ε -optimal automata, i.e., for any $\varepsilon > 0$, there exists an automaton in this class that achieves probability of error at most $P_e^*(S) + \varepsilon$. Thus, the optimal value can be approached arbitrarily closely.




Figure 1: Randomized binary hypothesis testing machine of [5] for $\text{Bern}(p)$ input

6.2 Deterministic Binary Hypothesis testing

As stated previously, the binary hypothesis testing algorithm that achieves the Hellman-Cover lower bound is randomized. There have been some focus on the limited-memory hypothesis testing problem with restriction to deterministic algorithms. For simplicity, assume we observe either the $\text{Bern}(p)$ distribution, under hypothesis \mathcal{H}_0 , or the $\text{Bern}(q)$ distribution, under hypothesis \mathcal{H}_1 , for $0 < q < p < 1$. Since the optimal error probability decreases exponentially in S , we can define the error exponent as

$$E(p, q) = - \lim_{S \rightarrow \infty} \frac{1}{S} \log P_e^*(S) = \frac{1}{2} \cdot \log \gamma = \frac{1}{2} \log \frac{p(1-q)}{q(1-p)}. \quad (27)$$

Now letting $P_{e,\text{det}}^*(S)$ be the optimal error probability that can be attained using deterministic algorithms, we want to similarly characterise the *deterministic error exponent*,

$$E_{\text{det}}(p, q) = - \lim_{S \rightarrow \infty} \frac{1}{S} \log P_{e,\text{det}}^*(S). \quad (28)$$

As deterministic machines are a special case of randomized machines, we clearly have $E_{\text{det}}(p, q) \leq E(p, q)$. To demonstrate the important role randomization plays

in approaching the optimal performance, Hellman and Cover showed in [101] that for any memory size $S < \infty$ and $\delta > 0$ there exists problems such that any S -state deterministic machine has probability of error $\mathbf{P}_e \geq \frac{1}{2} - \delta$, while the randomized machine from [5] has $\mathbf{P}_e \leq \delta$. In [102] (see also [103]) it was shown that $\mathbf{E}_{\text{det}}(p, q)$ is positive for all $p \neq q$, and for the symmetric setting, where $p = 1 - q$, Shubert et al. [104] have also derived a lower bound on $\mathbf{E}_{\text{det}}(p, q)$. This implies that whenever $\gamma < \infty$, i.e., for any $0 < p, q < 1$, there exists some integer $1 \leq C = C(p, q) < \infty$ such that $\mathbf{P}_{e, \text{det}}^*(S \cdot C) \leq \mathbf{P}_e^*(S)$, for all S .

The problem has been largely abandoned at that point, until fairly recently, when Berg et al. [90] analyzed the error probability of optimal deterministic machines, and gave upper and lower bounds on $\mathbf{E}_{\text{det}}(p, q)$ which coincide for some extreme values of the hypotheses. Specifically, they show that the optimal error probability of a deterministic machine has

$$r(p, q) \leq \mathbf{E}_{\text{det}}(p, q) \leq d(p, q), \quad (29)$$

where

$$r(p, q) = \frac{\log p \log(1 - q) - \log q \log(1 - p)}{\log p(1 - p) + \log q(1 - q)}, \quad (30)$$

$$d(p, q) = \frac{\log(\min\{p, 1 - p\}) \cdot \log(\min\{q, 1 - q\})}{\log(\min\{p, 1 - p\}) + \log(\min\{q, 1 - q\})}. \quad (31)$$

The lower bound on the error exponent is comprised of an algorithm that waits for events that very sharply distinguishes between the hypotheses, specifically, a long consecutive run of either zeros or ones, where the run length is chosen to optimize the error probability (see Figure 2). This N -state algorithm for testing between the hypotheses $\mathcal{H}_0 = p$ and $\mathcal{H}_1 = q$ is referred to as $\text{RUNS}(N, p, q)$, and it achieves error probability δ whenever $p = q + \varepsilon$ if the number of states is $N = O(1/\varepsilon \cdot \log(1/\delta))$.




Figure 2: $\text{RUNS}(N, p, q)$ - Deterministic Binary Hypothesis Testing Machine

The upper bound is based on the insight that, while in randomized schemes the transition probabilities between states can be as small as desired, in deterministic

machines the transition probabilities can only be as small as $\min\{p, 1 - p\}$ or $\min\{q, 1 - q\}$. More specifically, in order to derive a lower bound on $\mathbf{P}_{\mathbf{e}, \text{det}}^*(S)$ for ergodic machines, it is first shown that the error probability of any algorithm is lower bounded by $\max_{i \in [S]} \min\{\mu_i^p, \mu_i^q\}$. It is then left to show that there exists a state $i \in [S]$ for which both μ_i^p and μ_i^q are large. The proof follows the observation that, if state i is connected to state j , we must have $\mu_j^p \geq \mu_i^p \cdot \min\{p, 1 - p\}$. Using the argument recursively, and applying it to μ_i^q as well, gives the bound.

For the non ergodic case, first consider the two absorbing states case, one for each competing hypothesis. It is shown that in any finite state machine, there must exist some state that is highly likely under both hypotheses and is close to both absorbing states, and this forces the error probability to be large as well. The general non ergodic case is proven by a superposition of the ergodic and two absorbing states cases. Though in general not tight, the lower bound demonstrates the gap between the error exponent for deterministic machines and that of randomized ones, which was derived in [5]. In particular, for certain values of (p, q) , it was shown that the restriction to deterministic machines may arbitrarily degrade the error exponent. Specifically, for any fixed $q < \frac{1}{2}$, we have $\lim_{p \rightarrow 1} \mathbf{E}_{\text{det}}(p, q) = -\log q$ while $\lim_{p \rightarrow 1} \mathbf{E}(p, q) = \infty$, and, similarly, for any fixed $p > \frac{1}{2}$, we have $\lim_{q \rightarrow 0} \mathbf{E}_{\text{det}}(p, q) = -\log(1 - p)$ while $\lim_{q \rightarrow 0} \mathbf{E}(p, q) = \infty$.

6.3 Parity Learning

Here the family of distributions \mathcal{F} is characterized by some $\theta \in \{0, 1\}^d$ and is given as a pair of observations (X, Y) , where $X \sim \text{Bern}^{\otimes d}(1/2)$ and $Y = X^T \theta$, where $X^T \theta$ denotes inner product modulo 2. In particular, we observe samples $(X_1, Y_1), (X_2, Y_2) \dots$, where each X_t is uniformly distributed over $\{0, 1\}^d$ and for every t , $Y_t = X_t^T \theta$ and we want to estimate the vector θ under the 0 – 1 loss.

Following a conjecture of Steinhardt et al. [89], Raz proved in a seminal paper [91] that the problem of parity learning provides the following separation: Any algorithm for parity learning that uses less than $d^2/25$ bits of memory requires at least $2^{d/25}$ samples. To appreciate this result, assume that δ is some small constant. First note that for the parity learning problem we have $\text{SC}(\delta) \asymp d$ and $\text{MC}(\delta) \asymp d$ in bits. The sample-complexity follows since each observation gives at most 1 bit of information on θ , and $O(d)$ time can be achieved by using the natural approach of simply saving all the equations as if there are no memory limitations using $O(d^2)$ memory bits, and then performing Gaussian elimination. The memory-complexity follows since we need at least d bits to represent θ , and we can achieve $O(d)$ by the rare events approach of waiting for all the singletons, which takes an exponential number of samples in d . The results of [91] can be translated as a sharp decay of $n^*(S, \delta)$ on the sample-memory curve, around $S \asymp d^2$, from exponential in d to linear in d .

Building on this result, Kol et al. [105] considered the problem of learning

sparse parities. They have shown that the sparse parity learning problem, over the set of all vectors of Hamming weight exactly ℓ , requires either a memory of size super-linear in d or a number of samples super-polynomial in d , as long as $\ell \geq \omega(\log d / \log \log d)$. The work of Raz [106] expanded these previous results by showing that the same sharp drop-off of $n^*(S, \delta)$ in the sample memory curve is not only inherent to the parity learning problem, but is exhibited by a large class of problems. Specifically, let Θ, \mathcal{X} be two finite sets, and let $M : \Theta \times \mathcal{X} \rightarrow \{-1, 1\}$ be a matrix. The learning problem is the following: An unknown element $\theta \in \Theta$ is chosen uniformly at random. A learner then tries to learn θ from a stream of samples, $(X_1, Y_1), (X_2, Y_2), \dots$, where for every $i, X_i \in \mathcal{X}$ is chosen uniformly at random and $Y_i = M(X_i, \theta)$. Let σ_M be the largest singular value of M and note that $\sigma_M \leq |\mathcal{X}|^{\frac{1}{2}} \cdot |\Theta|^{\frac{1}{2}}$, since the spectral norm is upper bounded by the Frobenius norm. The paper shows that if $\sigma_M \leq |\mathcal{X}|^{\frac{1}{2}} \cdot |\Theta|^{\frac{1}{2}-\varepsilon}$, then any learning algorithm for the corresponding learning problem requires either a memory of at least $\Omega((\varepsilon d)^2)$ bits or at least $2^{\Omega(\varepsilon d)}$ samples, where $d = \log |\Theta|$.

In a recent work, Garg et al. [107] leveraged the tools developed in [91] to give memory-sample lower bounds for learning parity with noise. In particular, they showed that when a stream of random linear equations over \mathbb{F}_2 is of the form $(X_1, \theta^T X_1 \oplus Z_1), \dots, (X_n, \theta^T X_n \oplus Z_n)$ where $\{Z_i\}_{i=1}^n$ are distributed i.i.d $\text{Bern}(\varepsilon)$ for some $0 < \varepsilon < 1/2$, then the drop-off of $n^*(S, \delta)$ in the sample-memory curve occurs around $S \asymp d^2/\varepsilon$.

Lower Bound of Raz [91]

The works described above rely heavily on analyzing the evolution of the posterior on θ along the progress of a *branching program*. This technique was introduced in [91], and we will attempt to elucidate it now. A branching program of length n and width S is just a graphic representation of a time varying FSM with S states, operating on n samples. It is represented as a directed graph with vertices arranged in $n + 1$ layers containing at most S vertices each. Each layer represents a time step and each vertex represents a memory state of the learner. In the first layer, that we think of as layer 0, there is only one vertex, called the start vertex. A vertex of outdegree 0 is called a leaf. All vertices in the last layer are leaves (but there may be additional leaves). Every non-leaf vertex in the program has 2^{d+1} outgoing edges, one for each pair of values $(X, Y) = (x, y) \in \{0, 1\}^d \times \{0, 1\}$, with exactly one edge labeled by each such (x, y) , and all these edges going into vertices in the next layer. Each leaf v in the program is labeled by an element $\hat{\theta}(v) \in \{0, 1\}^d$, that we think of as the output of the program on that leaf. The samples $(X_1, Y_1), \dots, (X_n, Y_n) \in \{0, 1\}^d \times \{0, 1\}$ that are given as input define a computation-path in the branching program, by starting from the start vertex and following at step t the edge labeled by (x_t, y_t) , until reaching a leaf. The program outputs the label $\hat{\theta}(v)$ of the leaf v reached by the computation-path. The success

probability of the program is the probability that $\hat{\theta} = \theta$, where $\hat{\theta}$ is the label of the vertex reached by the path, and the probability is over $\theta, X_1, \dots, X_n, Y_1, \dots, Y_n$.

We give a sketch of the proof of [91], that shows that any branching program with width $S \leq 2^{d^2/25}$ and length $n \leq 2^{d/25}$ outputs the correct θ with exponentially small probability. We assume that each vertex v remembers a set of linear equations E_v such that if the computation path reaches v , then all equations in E_v are satisfied. This assumption restricts the branching program to be in a subclass called *affine branching programs*, but the author shows that any branching program for parity learning can be simulated by an affine branching program for parity learning (we skip the proof of this reduction). The output is then a random element $\hat{\theta} \in \{0, 1\}^d$ that satisfies E_v . The proof shows that the probability that the computation path reaches any v such that $\dim(E_v) \geq k$, for $k = \frac{1}{5}d$, is exponentially small. Why is that important? If we want to correctly estimate θ with high probability, we must arrive at a v^* with $\dim(E_{v^*}) \approx d$. For any edge (i, j) , we have $\dim(E_j) \leq \dim(E_i) + 1$ for all i , as we see only one new equation at each time step. Thus, in order to arrive at a v^* with $\dim(E_{v^*}) \approx d$, we must pass through some v with $\dim(E_v) \geq k$.

Assume that $\dim(E_v) \geq k$ and let Z_0, \dots, Z_m be the vertices on the computation path and $r_i = \dim(E_v \cap E_{Z_i})$. Note that $\dim(E_j) \leq \dim(E_i) + 1$ implies $r_{i+1} \leq r_i + 1$ for all i . Thus, if we reach a v with $\dim(E_v) \geq k$, then $r_{i+1} > r_i$ occurs at least k times, so it is enough to bound $\Pr(r_{i+1} > r_i)$. For this event to occur, the new sample (x_{i+1}, y_{i+1}) must lay in the space spanned by E_v , hence it can be shown that $\Pr(r_{i+1} > r_i) \leq 2^{-\Omega(d)}$. Then, by using the union bound over S^k vertices, the probability to reach any v with $\dim(E_v) \geq k$ is less than $S^k \cdot (2^{-\Omega(d)})^k = 2^{-\Omega(d^2)}$.

6.4 Additional Results

The finite-memory simple hypothesis testing problem had an impact on other problems in the memory constrained literature. In [108], Hellman and Cover solved the two armed bandit with time-invariant finite memory problem by using a similar construction and proof elements from [5]. Horos and Hellman [109] allowed a confidence to be associated with each decision, with errors being weighted according to the confidence with which they are made. They found the ε -optimal class of rules to be deterministic. Flower and Hellman [110] examined the finite sample problem for Bernoulli case. They found that randomization was needed on all transitions toward the center state, i.e., on transitions from states of low uncertainty to states with higher uncertainty. Continuing the exploration of the finite-sample hypothesis testing with finite memory problem, Cover et al. [20] established the existence of an optimal rule and found its structure for time-varying algorithms, and characterised the optimal error probability of the 2-state finite-sample problem that distinguishes between two Gaussian distributions with a different mean.

7 Parameter Estimation

Recall that in the parameter estimation setting, i.i.d. samples are drawn from some unknown distribution P belonging to a given parametric family $\mathcal{F} = \{P_\theta\}_{\theta \in \mathcal{Y}}$, where typically $\mathcal{Y} \subseteq \mathbb{R}^d$ for some d . The goal is to estimate the parameter θ itself under a prescribed loss function, hence in our notation we are interested in the property $\pi(P_\theta) = \theta$. A common choice of loss function in this setting is the squared error loss $\ell(\hat{\theta}, \theta) = \|\hat{\theta} - \theta\|^2$. Here also we will be interested in the limit of the problem when $n \rightarrow \infty$, thus following (12) we define the asymptotic quadratic risk attained by this estimator when the true value of the parameter is θ , to be

$$R(\theta, \hat{\pi}) = \limsup_{n \rightarrow \infty} \mathbb{E}(\hat{\pi}(X^n) - \theta)^2. \quad (32)$$

The main interest is thus to characterize the minimax risk $R^*(S)$ of the problem, which is the smallest asymptotic quadratic risk that can be uniformly guaranteed for all $\theta \in \mathcal{Y}$ using an FSM with S states.

The parameter estimation with finite-memory problem under the mean squared error loss function was first addressed by Roberts and Tooley [111], who tackled the problem of estimation under quadratic risk for a random variable with additive noise using time-varying machines. They restricted their rules to be of a special form which, although not optimal in general, does make sense. This restriction is that larger observations cause transitions to higher numbered states. In some problems (for example, the Cauchy distribution), very large observations yield very little information, and such a rule seems to be distinctly suboptimal. However, for Gaussian statistics, large observations are very informative, in which case their form does make sense. Koplowitz and Roberts [112] demonstrated necessary and sufficient conditions for the optimal state transition function.

Wagner [97] uses rules similar to Cover's [4] to estimate the mean of a distribution. For Bernoulli observations, Wagner's scheme is very close to optimal, since its maximum absolute error is at most ε , with $S = O(1/\varepsilon)$ memory states. He achieved this result by partitioning the observations into blocks $\{S_k^j, R_k^j\}_{k=1}^\infty$ sequentially testing whether $p > p_j$ or $p < p_j$, for $1 \leq j \leq r$, and $0 = p_0 < p_1 < \dots < p_r = 1, r = O(1/\varepsilon)$. By keeping the index j of the favourite hypothesis and choosing the block lengths to follow the Borrel-Cantelli convergence conditions, the correct block is identified with probability 1. This is an example of the methodology mentioned in Section 5.2, of breaking down a complex problem into multiple simple binary decision problems; we will see this again in the sequel.

Hellman [113] examined the problem of estimating the mean μ of a Gaussian Distribution $\mathcal{N}(\mu, \sigma^2)$ under the mean squared error loss, where μ is drawn from a known prior distribution $p(\mu)$. To obtain a lower bound, consider the simpler quantization problem, where we pick the best S -level scalar quantizer $Q(\mu)$ for μ , minimizing the mean squared error $\mathbb{E}(Q(\mu) - \mu)^2$. Clearly, this mean squared error

is a lower bound for the mean squared error attained by any S -state estimation algorithm. One might assume that this quantization lower bound is not tight, however, quite surprisingly, Hellman showed that the quantization lower bound can be approached arbitrarily closely by the following (deterministic) construction: Let $\mu_1^* < \mu_2^* < \dots < \mu_S^*$ be the optimal quantization points of an S -level quantizer. Consider the machine which transits from state i to state $i+1$ whenever $X_n - (\mu_i^* + \mu_{i+1}^*)/2 \geq M$, and from state i to state $i-1$ when $X_n - (\mu_{i-1}^* + \mu_i^*)/2 \leq -M$, and which estimates μ to be μ_i^* when it is in state i . Thus, the ratio between transition probabilities toward states with better μ estimates and transition probabilities toward states with worse μ estimates behaves exponentially with M , so that, as $M \rightarrow \infty$, the stationary distribution approaches a singleton on μ^* . Thus, in the Gaussian case, the mean estimation problem can be reduced to a quantization problem, that is, for any prior $p(\mu)$ on μ the minimum quadratic risk attained by an S -state estimator is equal to the quadratic risk attained by the optimal S -level quantizer.

The main reason we were able to attain the quantization lower bound in this case was that under the Gaussian distribution we sometimes, even though rarely, observe a sample that contains an arbitrarily large amount of information about the parameter. Thus, if we waited for such samples and only update our memory state upon observing them, we would converge to a state with the best estimator. Unlike the Gaussian distribution, the Bernoulli distribution has only two outcomes, 1 and 0, which contain very little information about the Bernoulli parameter. In this distribution, rare events can only occur w.r.t a sequence of samples, thus tracking these events incurs an additional memory cost. Indeed, we now turn our attention to the problem of estimating an unknown Bernoulli parameter, and show that in this case the quantization lower bound of $\Omega(1/S^2)$ *cannot* be achieved.

Samaniego [114] initiated the work on the problem of estimating the parameter of a coin using time-invariant finite-memory machines. He considered a Bayesian setup in which a prior is placed on the parameter distribution, and the loss function is a generalized quadratic loss function of the form $(\hat{\theta} - \theta)^2 \theta^{\beta-1} (1 - \theta)^{\gamma-1}$, for some $\beta, \gamma \geq 0$. Furthermore, he restricted to machines that can only make transitions between adjacent states, which he refers to as *tridiagonal* algorithms. Within the class of all tridiagonal algorithms, a particular rule was shown to be minimax as well as locally Bayes w.r.t the uniform prior for $S \leq 30$. Leighton and Rivest [115] followed up the work on estimating the parameter of a coin, and have found an exact characterization of the smallest attainable mean squared error for this problem. Specifically, they proved that

$$R^*(S) \asymp \frac{1}{S}, \tag{33}$$

or, equivalently, using the formulation of (19), they show that $\text{MC}(\delta) \asymp 1/\delta$.

Leighton and Rivest [115] - Upper bound

Samaniego's construction from [114] is analyzed and the mean squared error is recovered for any θ and S . Essentially, this finite-state algorithm uses S states to estimate the number of ones in a sliding window of length $S-1$, thus it is commonly referred to in the literature as "Imaginary Sliding Window" (See Figure 3). In this construction, if the machine is in state i and the fresh sample is "1", we move to state $i+1$ with probability $1 - \frac{i}{S-1}$ or stay in state i with probability $\frac{i}{S-1}$. If the fresh sample is "0", we move to state $i-1$ with probability $\frac{i}{S-1}$ or stay in state i with probability $1 - \frac{i}{S-1}$. The estimate function is $\hat{\theta}(i) = \frac{i}{S-1}$.




Figure 3: The bias estimation machine of [114] for $\text{Bern}(p)$ input and $q = 1 - p$

An analysis of the stationary distribution of the chain proves that this construction asymptotically induces a $\text{Binomial}(S-1, \theta)$ stationary distribution on the memory state space, thus achieving

$$\mathbb{E}(\hat{\theta}(i) - \theta)^2 = \text{Var}\left(\frac{\text{Binomial}(S-1, \theta)}{S-1}\right) = \frac{\theta(1-\theta)}{S-1}. \quad (34)$$

Leighton and Rivest [115] - Lower bound

The key ingredient in the proof is the use of the *Markov Chain Tree Theorem* (see the elegant proof of Anantharam and Tsoucas [116]), which implies that the limiting distribution $\pi(j)$ of an S -state Ergodic Markov chain (which is a unique stationary distribution) can be expressed as ratios of $S-1$ degree polynomials

with non-negative coefficients.⁵ Specifically, we have

$$\pi_j = \frac{\sum_{i=1}^S a_{ij} \theta^{i-1} (1-\theta)^{S-i}}{\sum_{i=1}^S a_i \theta^{i-1} (1-\theta)^{S-i}}, \quad (35)$$

where $a_{ij} \geq 0$ for $1 \leq i, j \leq S$, and $a_i = \sum_{j=1}^S a_{ij}$. Consequently, for any machine f and estimate function $\hat{\theta}$, the quadratic risk is

$$R_\theta(f, \hat{\theta}) = \frac{\sum_{j=1}^S \sum_{i=1}^S a_{ij} \theta^{i-1} (1-\theta)^{S-i} (\hat{\theta}(j) - \theta)^2}{\sum_{i=1}^S a_i \theta^{i-1} (1-\theta)^{S-i}}.$$

The remainder of the proof essentially shows that functions of this restricted form cannot accurately predict θ , by leveraging elementary polynomial analysis to show that there must be some θ^* for which $R_{\theta^*}(f, \hat{\theta}) = \Omega(1/S)$. Thus the limitations imposed by restricting the class of transition functions dominate the limitations imposed by quantization of the estimates (in contrary to the Gaussian case where quantization provided the main bottleneck).

7.1 Deterministic Algorithm for Bias Estimation

We consider now the bias estimation problem with restriction to deterministic machines. Following the results of Section 6, one might expect a degradation in performance when comparing to optimal randomized machines. In fact, in the same seminal work of [115], the authors further constructed a deterministic S -state estimation procedure by de-randomizing Samaniego's construction, and as a result showed that $R_{\text{det}}^*(S) = O(\log S/S)$. They then conjectured that this is indeed the deterministic minimax risk of the problem. A nice interpretation of their conjecture is the naturally appealing claim that an optimal deterministic algorithm can be obtained by derandomizing an optimal random algorithm. In their deterministic algorithm, randomness is extracted from the measurements via a Von Neumann like procedure [117], augmenting each state with $O(\log(S))$ additional states⁶.

⁵The theorem actually holds in the general case, for which the limiting distribution might not exist, with π_j replaced by \bar{p}_{1j} , where \bar{p}_{1j} is the long-run average transition probability that the chain will be in state j given that it started in state 1. Specifically, it is the $(1, j)$ element of the matrix $\bar{P} = \lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n P^{i-1}$, where P is the transition matrix of the chain, and \bar{P} always exists as P is stochastic.

⁶This derandomization is almost tight. We need to generate S Bernoulli r.v.s with parameters i/S for $i = 1, \dots, S$ for transition probabilities. Assume we use machines with N_t states respectively. The blowup factor N in the number of states is the average of N_t , and there must be at least $S/2$ machines with $N_t \leq 2N$. There are at most $(2N)^{4N+1}$ binary input FSM with at most $2N$ states, hence generating $S/2$ different Bernoulli parameters dictates that $(2N)^{4N+1} \geq S/2$, which implies $N = \Omega(\log(S)/\log \log(S))$.

After standing unresolved for a couple of decades, this conjecture was recently disproved by Berg et al. [92], who proved that

$$R_{\text{det}}^*(S) \asymp \frac{1}{S}, \quad (36)$$

which implies that the memory complexity of the problem is unaffected by the restriction to deterministic estimators.

This result is proved by an explicit construction of a deterministic S -state machine that attains quadratic risk of $O(1/S)$ uniformly for all $\theta \in [0, 1]$. The high-level idea is to break down the memory-constrained estimation task into a sequence of memory-constrained (composite) binary hypothesis testing sub-problems as in Wagner's [97] scheme, but using time invariant algorithms. In each such sub-problem, the goal is to decide whether the true underlying parameter θ satisfies $\{\theta < q\}$ or $\{\theta > p\}$, for some $0 < q < p < 1$. Those decisions are then used in order to traverse an induced Markov chain in a way that enables us to accurately estimate θ .

In particular, the state space is partitioned into K disjoint sets denoted by $\mathcal{S}_1, \dots, \mathcal{S}_K$, where the estimation function value is the same inside each \mathcal{S}_i . These sets are referred to as *mini-chains*, where the i th mini-chain consists of N_i states, and is designed to solve the composite binary hypothesis testing problem:

$$\mathcal{H}_0 : \left\{ \theta > \frac{i+1}{K} \right\} \text{ vs. } \mathcal{H}_1 : \left\{ \theta < \frac{i}{K} \right\}. \quad (37)$$

Each mini-chain \mathcal{S}_i is initialized in its entry state s_i , and eventually moves to the entry state s_{i+1} of mini-chain \mathcal{S}_{i+1} if it decided in favor of hypothesis \mathcal{H}_0 , or to the entry state s_{i-1} of mini-chain \mathcal{S}_{i-1} if it decided in favor of hypothesis \mathcal{H}_1 . The hypothesis testing mini-chains are taken to be the deterministic testers from [90], as can be seen in Figure 4.




Figure 4: Mini-chain \mathcal{S}_i for composite binary hypothesis testing

Recall that in section 6.2 we saw that $O(K)$ states are needed to resolve between two hypothesis that are $1/K$ apart with constant error probability. Thus, let-

ting $K = O(\sqrt{S})$, the described algorithm estimates θ to within $1/\sqrt{S}$ standard deviations using S states, which implies the correct minimax risk.

7.2 Linear Regression

The linear regression problem can be construed as the continuous variant of the parity learning problem. Here the family of distributions \mathcal{F} is characterized by some $\theta \in \mathbb{R}^d$ and is given as a stream of samples $(X_i, Y_i) \in \mathbb{R}^d \times \mathbb{R}$, where the X_i 's are drawn from a known i.i.d distribution $X \sim \mathcal{D}$, and $Y_i = \theta^T X_i + V_i$, where V_i is an i.i.d. noise drawn from a known distribution \mathcal{D}_V , independent of X and with zero mean. Our goal is then to estimate the d -dimensional vector θ under the ℓ_2 loss.

In [118], Steinhardt and Duchi considered the case of memory-bounded sparse linear regression, in which θ is k -sparse, meaning that it has at most k non-zero entries. Essentially, they prove upper and lower bound showing that

$$\Omega\left(\frac{kd}{b\varepsilon}\right) \leq n^*(2^b, \varepsilon) \leq \tilde{O}\left(\frac{kd}{b\varepsilon^2}\right).$$

The lower bound relies on adapting Assouad's method (see [119]) to reduce the analysis to bounding the mutual information between a single observation and the parameter vector θ . The limitation of their proof is in a particular assumption on X_i and letting the additive noise be the unnatural Laplace distribution. They also provide an algorithm that uses $\tilde{O}(b+k)$ bits and requires $\tilde{O}\left(\frac{kd}{b\varepsilon^2}\right)$ observations to achieve error ε whenever the covariance matrix of X is the identity matrix and X has a small ℓ_∞ norm. In the upper bound, they show how to implement ℓ_1 -regularized dual averaging using a count sketch data structure, which efficiently counts elements in a data stream. They then use the count sketch structure to maintain a coarse estimate of model parameters, while also exactly storing a small active set of at most k coordinates. These combined estimates allows to implement dual averaging when the ℓ_1 -regularization is sufficiently large; the necessary amount of ℓ_1 -regularization is inversely proportional to the amount of memory needed by the count sketch structure, leading to a tradeoff between memory and statistical efficiency.

In [120], Sharan et al. considered the case of linear regression where the X are drawn independently from a d -dimensional isotropic Gaussian and the noise is drawn uniformly from the interval $[-2^{-d/5}, 2^{-d/5}]$. They show that any algorithm with at most $d^2/4$ bits of memory requires at least $\Omega(d \log \log 1/\varepsilon)$ for ε sufficiently small as a function of d . In contrast, for such ε , θ can be recovered to error ε with probability $1 - o(1)$ with memory $O(d^2 \log(1/\varepsilon))$ bits using d samples. The uniform small error is crucial for the proof, as it reframes the problem as a quantization problem that allows the authors to leverage branching programs for the lower bound.

7.3 Additional Results

Recently, Jain and Tyagi [121] studied the shrinkage in memory between the hypothesis testing and the estimation problem, namely the interesting fact that a machine with S states can distinguish between two coins with biases that differ by $1/S$, whereas the best additive accuracy it can achieve in estimating the bias is only $1/\sqrt{S}$ (as a result of the previously works of [5, 115]). The authors explain the memory shrinkage phenomenon by showing that when estimating the unknown bias θ from an interval using an S -state machine, the effective number of states available is only $O(\sqrt{S})$. Specifically, they show that for any estimator with S states, there exist values p and q of the bias such that $|p - q| > 1/\sqrt{S}$, but the effective support sets of their equilibrium distributions are contained in a set of cardinality $O(\sqrt{S})$. In [122], Kontorovich defined a consistent estimator as an estimator that converges in probability to θ , and showed that it is impossible to achieve consistent estimation with bounded memory. However, relaxing the requirement of consistency to ε -consistency, he showed that $O(\log(1/\varepsilon))$ states suffice for a natural class of problems (the binary hypothesis testing problem being one of them, as can also be gleaned from Section 6).

In the problem of *Universal Prediction* initiated with the works of Merhav and Feder [123, 124], we observe either a worst case or a randomly generated sequence, and are interested in predicting the next value in the sequence. The performance is dominated by the choice of the loss function and is evaluated with respect to an adversary who can see the entire sequence but is limited to the constant predictors family. The figure of merit here is the coding *redundancy*, which is the difference between the optimal regret we can achieve and the one achieved by the adversary. It is known that optimal universal prediction under the log-loss measure is equivalent to optimal encoding of a sequence.

In [125], Rajwan and Feder considered universal finite-memory machines for coding binary sequences, and gave lower bounds for deterministic, randomized, time-invariant or time-variant machines. Meron and Feder [126] showed that for universal prediction of worst case sequences under the square error loss, the optimal asymptotic expected redundancy is $O(1/S)$ and it is achieved by the construction of [115]. For the log-loss case, using the same construction results in an upper bound on the redundancy of $O(\log S/S)$, while Pinsker's inequality gives a lower bound of $\Omega(1/S)$. In a follow up work [39], the same authors also consider universal prediction of worst case sequences using deterministic finite-state machines, proving a lower bound on the regret of $\Omega(S^{-4/5})$, and suggested the "Exponentially Decaying Memory" machine, which achieves regret of $O(S^{-2/3})$.

8 Distribution Property Testing

Distribution property testing is a special case of composite hypothesis testing, where we are interested in deciding whether the underlying distribution has a certain property or is ε -far in total variation from having this property. We focus here on a particular fundamental instance of the property testing problem, known as *uniformity testing*.

In the uniformity testing problem we observe a sequence of independent identically distributed random variables drawn from an unknown distribution P over $[k]$, which is guaranteed to be either the uniform distribution U_k , or ε -far from the uniform distribution under the total variation distance. In other words, there are two composite hypotheses in this problem, $\mathcal{F}_0 = \{U_k\}$, and \mathcal{F}_1 which contains all the distributions P that satisfy $d_{\text{TV}}(P, U_k) > \varepsilon$, and we want to identify the correct (composite) hypothesis with probability of success at least $2/3$. We note that this problem is a particular instance of the identity testing problem, in which $\mathcal{F}_0 = \{Q\}$, for some known distribution Q (not necessarily uniform). However, it was shown by Goldreich [127] that the problem of uniformity testing can be reduced to identity testing, and, in particular, that for every distribution Q over $[k]$ and every $\varepsilon > 0$, it holds that ε -testing identity to Q reduces to $\varepsilon/3$ -testing identity to U_{6k} , and that the same reduction can be used for all $\varepsilon > 0$.

The study of uniformity testing has been initiated by Goldreich and Ron [128], who proposed a simple and natural uniformity tester that relies on the *collision probability* of the unknown distribution, which is the probability that two samples drawn according to P are equal, and requires $O(\sqrt{k}/\varepsilon^4)$ samples. In subsequent work, Paninski [129] proved a lower bound of $\Omega(\sqrt{k}/\varepsilon^2)$ on the sample complexity and provided a matching upper bound that holds under some assumption on ε , which has been later shown to be unnecessary by Diakonikolas et al. [130]. The lower bound is established by replacing the composite class \mathcal{F}_1 with a random distribution, chosen according to the so called *Paninski prior* on \mathcal{F}_1 . In particular, for any vector $Z \in \{0, 1\}^{k/2}$, we define a distribution

$$P_{\text{pan},i}^\varepsilon(Z) = \begin{cases} \frac{1+2\varepsilon(-1)^{Z_{i/2}}}{k}, & i \text{ even,} \\ \frac{1-2\varepsilon(-1)^{Z_{(i+1)/2}}}{k}, & i \text{ odd,} \end{cases} \quad (38)$$

for which the value of each coordinate $i \in [k]$ is determined by the vector $Z \in \{0, 1\}^{k/2}$. Then, under \mathcal{H}_1 , we draw $Z \sim \text{Unif}\{0, 1\}^{k/2}$ and set $P = P_{\text{pan}}^\varepsilon(Z)$. The upper bound is based on collision estimators, since \sqrt{k} samples will result in a single collision on average under the uniform distribution, and $1 + \Omega(\varepsilon^2)$ collisions on average under the alternative, and the variances are small enough to distinguish between the hypotheses.

The memory complexity of estimating the second moment of the empirical distribution (often referred to as the empirical collision probability) is well-studied

for worst case data streams of a given length, dating back to the seminal work of Alon et al. [21]. In the stochastic setup, Crouch et al. [80] studied the trade-off between sample complexity and memory complexity for the problem of estimating the collision probability. The trade-off between sample complexity and memory / communication complexity for the distribution testing problem has recently been addressed by Diakonikolas et al. [85], who gave upper and lower bounds on the sample complexity under the constraint of using at most b memory bits.

Their upper bound is comprised of a modified uniformity tester that can be implemented in both the memory and communication restricted settings. In particular, a bipartite collision tester is suggested, that works as follows: Draw two independent sets of samples S_1 and S_2 from the unknown distribution P , of sizes $N_1 = \frac{b}{2 \log k}$ and $N_2 = \Theta\left(\frac{k \log k}{b \varepsilon^4}\right)$ respectively, and count the number of pairs of an element of S_1 and an element of S_2 that collide. Importantly, they have $N_1 \cdot N_2 \gg k/\varepsilon^4$ which implies $N_1 \ll N_2$, and $\min\{N_1, N_2\} \gg 1/\varepsilon^2$ which implies $b \gg \log k/\varepsilon^2$. The algorithm first stores S_1 exactly, and then proceeds to test whether there is some element of S_1 that repeats itself more than 10 times. As the probability of such recurrence is exceedingly small under the uniform distribution, if some symbol in S_1 does appear more than 10 times, the algorithm confidently rejects the uniform hypothesis. This is done to save in the amount of memory needed for the second stage, in which, for each element of the set S_2 , the algorithm counts the number of collisions this element has with elements of S_1 . To show that the number of collisions is much larger in the non-uniform case than in the uniform case, note that the expected number of collisions is just $N_1 \cdot N_2 \cdot \|P\|_2^2$, and by standard concentration bounds one can show that it will likely be close to this value. The amount of memory bits necessary for the algorithm in the first stage is $N_1 \cdot \log k$, as we save the entire set S_1 . The amount of memory bits necessary in the second stage is at most $\log(10 \cdot N_2)$, as the multiplicity of a symbol in S_1 is at most 10 after the first stage, and that each symbol collides with at most N_2 elements of S_2 . Thus, with at most $N_1 \cdot \log k + \log N_2 + 4 < b$ bits and $\Theta\left(\frac{k \log k}{b \varepsilon^4}\right)$ samples, one can solve the uniformity testing problem.

The lower bound of [85] is proved in a Bayesian setting using information theoretic tools. Typically, for the purpose of lower bounds, the minimax setting is replaced by a Bayesian one, where the alternative hypothesis is the so-called Paninski prior or *Paninski mixture*, a uniform mixture over all Paninski distributions of the form (38) (the specific distribution from the mixture is picked only once, uniformly at random, and then we observe i.i.d. samples from it). A distribution from this mixture is constructed as follows: Let $X = 2J - (Z_J \oplus B_J)$, where $J \sim \text{Unif}([k/2])$, $B \sim \text{Bern}^{\otimes k/2}(1/2 - \varepsilon)$, and B is independent of I . Note that for $\varepsilon = 0$, this results in the uniform distribution regardless of Z . On the other hand, for any choice of $z \in \{0, 1\}^{k/2}$, each of these distributions is exactly $P_{\text{pan}}^\varepsilon(z)$, which satisfies the total variation condition $d_{\text{TV}}(P_{\text{pan}}^\varepsilon(z), U_k) = \varepsilon$, namely is a valid alternative to the uniform hypothesis.

The lower bound is established for the case where Z is first drawn uniformly on $\{0, 1\}^{k/2}$ and then X_1, \dots, X_n are generated i.i.d. either from $P_{\text{pan}}^\varepsilon(Z)$ or from U_k . A high-level interpretation of the argument is as follows: For any algorithm, the mutual information between the memory content and Z is at most b , so that, from standard rate-distortion results [14], the best estimate we can get for any Z_i has error $h_2^{-1}(1 - \frac{b}{k}) = \frac{1}{2} - \Omega\left(\sqrt{\frac{b}{k}}\right)$ on average, where $h_2(\cdot)$ is the binary entropy function, and h_2^{-1} is its inverse restricted to $[0, 1/2]$. In each measurement we observe, the random variable I entails no information on whether or not B_J is biased, and hence the informative part of the observation is $Z_J \oplus B_J$, which is unbiased if $B_J \sim \text{Bern}(1/2)$, corresponding to $P = U_K$, and has bias $\Omega\left(\varepsilon \cdot \sqrt{\frac{b}{k}}\right)$ on average, under the alternative hypothesis. Thus, assuming both hypotheses are equally likely a-priori, the mutual information for determining whether the bias of B_J is 0 or ε , is $O(\varepsilon^2 \cdot \frac{b}{k})$ per measurement, and $O(n\varepsilon^2 \cdot \frac{b}{k})$ in total. Thus we need $n = \Omega(k/(b\varepsilon^2))$ samples to succeed.

Using our definition of the sample-memory curve at $\delta = 1/3$, the above results can be written as

$$\Omega\left(\frac{k}{b\varepsilon^2}\right) \leq n^*(2^b, 1/3) \leq \tilde{O}\left(\frac{k}{b\varepsilon^4}\right),$$

where the upper bound holds for $b \gg \frac{\log k}{\varepsilon^2}$ bits. A natural question that arises from [85] is to what extent this memory size can be reduced if we are allowed to process more samples. Another open question posed by the authors is whether the problem can be solved with $o(\log k)$ memory bits. The first question was addressed by Meir [131], who showed that one can achieve the sample complexity of [85] whenever the memory size in bits is $\Omega(\log(k/\varepsilon))$. This is done via comparison graphs-based algorithm, which is a general model for batch collision estimators, whose memory need only store the samples used for collision testing, and a global counter. The second open question was resolved in a recent work of Berg et al. [93] on finding the minimax memory complexity of uniformity testing. In the paper, the authors gave a lower bound of $\Omega(\log k)$ on the memory complexity, thus answering the question from [85] in the negative. They further proposed an algorithm that succeeds with $\log(k/\varepsilon) + o(\log \log k)$ memory bits. Next, we outline the derivation of these bounds.

Upper Bound on the Memory complexity of [93]

We now describe an algorithm that succeeds with $O\left(\frac{k \log k}{\varepsilon}\right)$ states. In a nutshell, the suggested algorithm reduces uniformity testing to a sequence of binary hypothesis tests. The underlying idea is that if $d_{\text{TV}}(P, U_k) > \varepsilon$, then there must be some symbol $i \in [k]$ such that either $\Pr(X = 1 | X \in \{1, i\})$ or $\Pr(X = i | X \in \{1, i\})$ are ε -biased, where by ε -bias we mean equals $1/2 + \Omega(\varepsilon)$. On the other hand, if

$P = U_k$, none of these probabilities is ε -biased, in fact, they are exactly $\text{Bern}(1/2)$. Following this observation, we can construct a finite state machine whose state space is partitioned into $2k$ disjoint sets (or mini-chains) denoted by \mathcal{S}_i , which are identical binary hypothesis testing $\text{RUNS}(N, 1/2 + \Omega(\varepsilon), 1/2)$ machines with $N = O(1/\varepsilon \cdot \log k)$ states. These mini-chains test, for all $i > 1$, whether $\Pr(X = 1|X \in \{1, i\})$ or $\Pr(X = i|X \in \{1, i\})$ are ε -biased. If some \mathcal{S}_i tests positive for ε -bias, the algorithm decides that P is ε -far from uniform and terminates, otherwise it moves to the next mini-chain, whereas if no test is positive for ε -bias, the machine decides that P is uniform. Due to the properties of the $\text{RUNS}(N, p, q)$ algorithm, the fact that each machine has $N = O(1/\varepsilon \cdot \log k)$ states guarantees that Bernoulli parameters that have $O(\varepsilon)$ bias between them can be distinguished with probability of error $\ll 1/k$, thus all tests are successful with some constant error probability (1/3, for example). We note that the effective run time of the above algorithm is $O(k^2 e^{O(1/\varepsilon)})$. This follows since there are $O(k)$ mini-chains that we traverse in sequence, each is activated every $O(k)$ samples on average, and the mixing time of each one is $e^{O(1/\varepsilon)}$.

Lower Bound on the Memory complexity of [93]

The authors show that $\Omega(k + 1/\varepsilon)$ states are necessary for uniformity testing. For the $\Omega(1/\varepsilon)$ lower bound, they use the fact that a good uniformity tester is also a good binary hypothesis tester for hypotheses that are ε -far. This is done by choosing the hard setting to be testing between the uniform distribution and a fixed member of the Paninski family (38). The problem is then reduced to deciding between $\text{Bern}(1/2)$ and $\text{Bern}((1 - \varepsilon)/2)$, which, from the lower bound of Hellman and Cover [5], must have $\Omega(1/\varepsilon)$ states.

To obtain the $\Omega(k)$ lower bound, we focus first on the case of ergodic Markov chains. The proof is based on analysis of stationary distributions and properties of Markov chain transition matrices null space, and it shows that an input distribution can be perturbed in a way that leaves the stationary distribution of the chain unchanged. Thus, there exists a perturbation of the uniform distribution that results in a distribution ε far from uniform, but induces the same stationary distribution on the state space of the chain. For a given FSM, denote the transition matrix and stationary vector induced by a distribution p as $\mathbf{P}(p)$ and π_p , respectively. We can write the stationary equations under the uniform distribution, $\pi_{U_k} = \pi_{U_k} \mathbf{P}(U_k)$, as $\pi_{U_k} = M A U_k$, where $M \in [0, 1]^{S \times S^2}$ is of the form

$$M = \begin{pmatrix} \pi_{U_k} & 0 & \dots & 0 \\ 0 & \pi_{U_k} & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \pi_{U_k} \end{pmatrix}, \quad (39)$$

and $A \in [0, 1]^{S^2 \times k}$ is a matrix whose rows are indexed by $(i, j) \in [S] \times [S]$, such that the entry $A_{(i,j),m}$ specifies the probability of moving to state j from state i when the input is $m \in [k]$. To simplify, we can write $\pi_{U_k} = MAU_k = BU_k$, where $B = MA \in [0, 1]^{S \times k}$. If $S < k$, the matrix B must have a nontrivial kernel of rank at least $k - S$. Letting V be the subspace spanned by $\ker(B) \cap \ker([1 \dots 1])$, we have that $\dim(V) \geq k - 1 - S \triangleq K$. Using properties of norm contraction, one can show that there exists some vector $x \in V$ with $\|x\|_\infty = 1/k$ and $\|x\|_1 \geq K/k$. For this x , we can define $Q = U_k - x$. Recalling that V is orthogonal to $[1 \dots 1]$, we have that Q is a valid probability distribution and $d_{\text{TV}}(U_k, Q) = \frac{1}{2}\|x\|_1 \geq K/2k$. Since $x \in \ker(B)$, we thus have that $BQ = BU_k = \pi_{U_k}$. But $BQ = \pi_{U_k}$ are simply the stationary equations for the chain under Q , thus $\pi_Q = \pi_{U_k}$. This implies that for any finite state machine with $S = (1 - 2\varepsilon)k - 1$ states, there exists a distribution Q with $d_{\text{TV}}(U_k, Q) \geq \varepsilon$ such that $\pi_{U_k} = \pi_Q$, so the machine cannot distinguish between U_k and Q . Thus, an ergodic finite state machine that attains non-trivial error must have $\Omega(k)$ states.

The $\Omega(k)$ lower bound bound for the non-ergodic case is based on the observation that any (S, δ) non-ergodic uniformity tester can be reduced to an $(S, \delta + \gamma)$ ergodic uniformity tester (for any $\gamma > 0$) by connecting all recurrent states to the initial state with some small positive probability. The result then follows by appealing to the lower bound of the ergodic case above.

8.1 Additional Results

In a recent work, Canonne and Yang [132] purposed simpler algorithms for uniformity testing with n samples that use b bits of memory. In their setup, the restriction on the memory size are $b \geq \max\{\log k, \log n\}$ and $b \leq \min\{n \log k, k \log n\}$. The first upper bound follows since straightforward lossless encoding takes $n \log k$ bits, and the second follows since only keeping the number of times each symbol is seen among the n samples, which is a sufficient statistic for the problem, takes $k \log n$ bits. Their first (deterministic) algorithm achieves the same upper bound as [85], and relies on a uniformity testing algorithm purposed by Diakonikolas et al. [133], which is the total variation distance between the empirical distribution and the uniform distribution. Letting N_i denote the frequency of the symbol i , and assuming that $n \leq k$, this statistic is just the normalized number of unseen symbols, that is

$$Z = \frac{1}{2} \sum_{i=1}^k \left| \frac{N_i}{n} - \frac{1}{k} \right| = \frac{1}{k} \sum_{i=1}^k \mathbb{1}_{N_i=0}. \quad (40)$$

Specifically, the algorithm divides the stream of n independent samples into T batches of s samples, which we keep in memory during each batch. At the end of the batch we compute the statistic Z_t and only keep it in memory, while reusing

the memory of the s batch samples. This implies $b = \Theta(s \log k)$, and a variance analysis of the estimator shows that the algorithm works as long as $s^2 T = \Omega(k/\varepsilon^4)$. Since $n = s \cdot T$, this results in the upper bound of $n = O\left(\frac{k \log k}{b \varepsilon^4}\right)$.

The second (randomized) algorithm relies on the primitive of *domain compression* introduced by Acharya et al. [71], a variant of hashing tailored to distribution testing and learning, where one can transform an instance of testing over domain size k and distance parameter ε to a new instance over domain size k' and (smaller) distance parameter $\varepsilon' \asymp \varepsilon \cdot \sqrt{\frac{k'}{k}}$. Letting $k' \asymp b/\log n$, the memory can now save all the 'new' samples, and due to known results on uniformity testing, the resulting algorithm works as long as the number of samples n satisfies

$$n = O\left(\frac{\sqrt{k'}}{\varepsilon'^2}\right) \implies \frac{n}{\sqrt{\log n}} = O\left(\frac{k}{\varepsilon^2 \sqrt{b}}\right). \quad (41)$$

A recent paper by Roy and Vasudev [134] considers distribution testing of a range of properties in the streaming model. They rely on the work of [85] on uniformity testing to obtain streaming algorithms for shape-restricted properties. They also consider streaming distribution testing in other models than the standard i.i.d. sampling one, specifically the conditional sampling model (see [135, 136]).

9 Distribution Property Estimation

The family of distributions is again parameterized by $\mathcal{F} = \{P_\theta\}_{\theta \in \mathbb{R}^d}$, but here we are trying to learn some simple scalar function $h(P_\theta)$ instead of θ itself. This is usually some continuous (but not one to one) function of the parameter, unlike in the property testing in which we were interested in a discrete function of the parameter. We discuss exclusively the problem of entropy estimation in this section.

9.1 Entropy Estimation

In this problem, the parametric family is $\mathcal{F} = \Delta_k$, where Δ_k denotes the collection of all discrete distributions P over $[k]$, and we want to estimate the Shannon entropy of P , given by $H(P) = -\sum_{x \in [k]} P(x) \log P(x)$, thus $\pi(P) = H(P)$. The loss function is taken to be the 0-1 loss w.r.t. the ε -error event, and it gives rise to the risk

$$P_\varepsilon(P, \hat{\pi}) = \limsup_{n \rightarrow \infty} \Pr(|\hat{\pi}(T_n) - H(P)| > \varepsilon), \quad (42)$$

for $\varepsilon > 0$. The minimax risk is thus defined as

$$P_e^*(S) = \inf_{\hat{\pi} \in \text{FSM}(S)} \sup_{P \in \Delta_k} P_e(P, \hat{\pi}). \quad (43)$$

We are generally interested in $\text{MC}(1/3)$, i.e., the smallest number of states needed to guarantee ε -approximate error probability of at most $1/3$ for all $P \in \Delta_k$.

The problem of estimating the entropy of a distribution from i.i.d. samples was addressed by Basharin [137], who suggested the simple and natural *plug-in estimator*, which simply computes the entropy of the empirical distribution of the samples, and requires $\asymp \frac{k}{\varepsilon} + \frac{\log^2 k}{\varepsilon^2}$ samples, thus giving an upper bound on the sample complexity of the problem. In the last two decades, many efforts were made to improve the bounds on the sample complexity of the entropy estimation problem. Paninski [138] was the first to prove that it is possible to consistently estimate the entropy using a sample size sublinear in the alphabet size k . The scaling of the sample complexity was shown to be $\frac{k}{\log k}$ in the seminal results of Valiant and Valiant [139, 140]. The optimal dependence of the sample size on both k and ε was not completely resolved until recently where, following the works of [141, 12, 13], the sharp sample complexity was shown to be

$$\text{SC} \asymp \frac{k}{\varepsilon \log k} + \frac{\log^2 k}{\varepsilon^2}. \quad (44)$$

The problem of estimating the entropy of a distribution under memory constraints is the subject of a more recent line of work. In [142], Chien et al. addressed the problem of deciding if the entropy of a distribution is above or below some predefined threshold, using algorithms with limited memory. The sample-memory curve of entropy estimation was investigated by Acharya et al. [36], who gave an upper bound on the sample complexity under the constraint that the algorithm only uses $O(\log(\frac{k}{\varepsilon}))$ memory bits, which has the same dependence on k as the plug-in estimator with no memory constraints. Specifically, he constructed an algorithm which is guaranteed to work with $O(k/\varepsilon^3 \cdot \text{polylog}(1/\varepsilon))$ samples and any memory size $b \geq 20 \log(\frac{k}{\varepsilon})$ bits (the constant 20 in the memory size can be improved by a careful analysis, but cannot be reduced to 1). Their upper bound was later improved by Aliakbarpour et al. [37] to $O(k/\varepsilon^2 \cdot \text{polylog}(1/\varepsilon))$ with memory complexity of $O(\log(\frac{k}{\varepsilon}))$ bits.

In a very recent result [143], Berg et al. have proved that $\log k + 2 \log(1/\varepsilon) + o(\log k)$ bits suffice for entropy estimation when $\varepsilon > 10^{-5}$, thus showing that the constant number of $\log k$ bits needed can be reduced all the way down to 1 (barring that ε is not too small). Furthermore, the proposed algorithm obtaining the upper bound was shown to converge within $\tilde{O}(k^c)$ samples, for any $c > 1$. The algorithm approximates the logarithm of $P(x)$, for a given $x \in [k]$, using a *Morris counter* [144], a concept we now briefly explain: Suppose one wishes to implement a counter

that counts up to m . Maintaining this counter exactly can be accomplished using $\log m$ bits. 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. The Morris Counter was later analyzed in more detail by Flajolet [29], who showed that $O(\log \log m + \log(1/\varepsilon) + \log(1/\delta))$ bits of memory are sufficient to return a $(1 \pm \varepsilon)$ approximation with success probability $1 - \delta$ (this was later improved by Nelson and Yu [145], who showed that $O(\log \log m + \log(1/\varepsilon) + \log \log(1/\delta))$ bits suffice).

The inherent structure of the Morris counter is particularly suited for constructing a nearly-unbiased estimator for $\log P(x)$, which makes it a judicious choice as a component in memory efficient entropy estimation. The algorithm of [143] uses two Morris counters, one that approximates a clock, and one that approximates $\log P(X)$. In order to compute the mean of these estimators, $\mathbb{E}[\log P(\bar{X})]$, in a memory efficient manner, a finite-memory bias estimation machine (e.g., [115, 92]) is leveraged for simulating the expectation operator. The bias estimation machine is incremented whenever a count is concluded in the Morris counter approximating the clock, thus essentially averaging our estimates over all x values. With $\tilde{O}(k^c)$ samples, this averaging converges (approximately) to the mean of $-\log P(X)$, and thus outputs an approximation to the true underlying entropy.

10 Conclusions, Outlook, and Open Problems

We have discussed the problem of statistical inference under memory constraints, and surveyed the state of the art for several problems including binary hypothesis testing, bias estimation, uniformity testing and entropy estimation. We focused our attention on these particular problems, since we see them as canonical problems in statistics. Nevertheless, about any other problem in statistical inference can, and should, be studied under limited memory constraints. We therefore refrain from specifying particular examples of important open problems, with only one notable exception - the linear regression problem. While this is arguably one of the most extensively studied problems in statistics, the understanding of how memory constraints affect it is quite limited. In Section 7.2 we describe some known results in this context, and point out their limitations. Ideally, one would like to characterize the optimal dependence of the ℓ_2 loss (or any other loss) on the number of samples n , the number of memory states S , the dimension d and the noise variance σ^2 , in the setup $Y_i = \theta^T X_i + V_i$, $i = 1, \dots, n$, where $\theta \in \mathbb{R}^d$ belongs to a bounded set or is sparse (or is drawn from some prior), $X_i \stackrel{i.i.d.}{\sim} \mathcal{D}$ for some standard distribution \mathcal{D} , say $\mathcal{D} = \mathcal{N}(0, I)$, and $V_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2 I)$. First order methods, such as stochastic gradient descent, often give a reasonable tradeoff between the memory and sample cost in this model, but the optimal loss remains unknown (see [120]).

We have seen that Markov chain theory plays a major role in characterizing the minimax risk in limited memory setting, often in deriving upper bounds by analyzing the induced limiting distributions for certain carefully designed finite-state machines. But this methodology was also used in [93] to obtain a lower bounds in the uniformity testing problem, by showing that if a machine has a small number of states, then there must exist some far from uniform distribution that induces the same stationary distribution on the state space as the uniform does. This observation suggests that the problem of statistical inference under memory constraints, in the limit of multiple samples, is in fact equivalent to question of Markov chain *instability*, namely, to characterizing how sensitive a given chain topology can be to its driving distribution. Let us make this concrete. For a state space with S states and a sample alphabet size k , a chain topology is a collection $\mathcal{M} = \{M_s \in \{0, 1\}^{k \times S}\}_{s \in [S]}$ of binary matrices, each satisfying $M_s \cdot \mathbf{1} = \mathbf{1}$. We interpret $(M_s)_{i, s'} = 1$ as an edge connecting state s to state s' , when the input symbol is i . This formulation corresponds to deterministic S -state machines, while allowing $M_s \in \{[0, 1]^{k \times S}\}_{s \in [S]}$ corresponds to randomized machines. Now, let P be a probability distribution of the samples, represented as a column vector of length k . Define the function

$$\xi_{S, \mathcal{M}}(P) = \frac{1}{S} \mathbf{1}^T \cdot \lim_{n \rightarrow \infty} ((I_S \otimes P^T) \cdot \text{diag}(\{M_s\}_{s \in [S]}))^n, \quad (45)$$

whenever the limit exists, where \otimes is the Kronecker product, and I_S is the S -dimensional identity matrix. In that case $\xi_{S, \mathcal{M}}(P)$ is the stationary distribution of an S -state chain with topology \mathcal{M} under a driving distribution P . Now, we can phrase questions about the feasibility of limited-memory inference in terms of $\xi_{S, \mathcal{M}}(P)$. For example, in a simple binary hypothesis testing between P and Q , we see that

$$\sup_{\mathcal{M}} d_{\text{TV}}(\xi_{S, \mathcal{M}}(P), \xi_{S, \mathcal{M}}(Q)) > 1/2 \quad (46)$$

is a necessary condition for S states to suffice. Or, for estimation in a parametric family $\{P_\theta\}$ under some metric loss function $\ell(\hat{\theta}, \theta)$, we have that

$$\sup_{\mathcal{M}} \inf_{\ell(\theta, \theta') \geq \epsilon} d_{\text{TV}}(\xi_{S, \mathcal{M}}(P_\theta), \xi_{S, \mathcal{M}}(P_{\theta'})) > 1/2 \quad (47)$$

is a necessary condition to attain a risk ϵ . Hence, it is interesting to study the joint range

$$\{(d_{\text{TV}}(\xi_{S, \mathcal{M}}(P_\theta), \xi_{S, \mathcal{M}}(P_{\theta'})), \ell(\theta, \theta')) : \theta, \theta' \in \Theta, \mathcal{M}\} \quad (48)$$

consisting of all pairs of stationary total variation distance and loss that can be attained by any (deterministic) S -state machine.

This survey has focused mainly on characterizing the memory complexity of the problems that were considered, while assuming we have access to an unlimited number of samples. Characterizing the optimal scaling of the risk as a function of both the number of samples n and the number of memory states S is a significantly harder problem, which remains open even for the simplest models one can think of, e.g., the bias estimation model. Information theoretic tools, mainly variants of the information bottleneck concept, were found useful for deriving lower bounds in these scenarios, as described, e.g., in Section 5. The idea is to show that with a small number of memory states, we cannot predict well the next sample X_{t+1} from the memory state U that depends only on X^t (which implies that we do not learn much about θ from U). To that end, one analyzes the information bottleneck function

$$\text{IB}(\log S) = \max_{P_{U|X^t}: I(X^t; U) \leq \log S} I(X_{t+1}; U). \quad (49)$$

In some cases, the restriction to the Markov chain $U - X^t - X_{t+1}$ is too weak to capture the sequential nature of the problem, where U is actually obtained by t consecutive updates. We believe that developing stronger techniques for upper bounding $I(U; X_{t+1})$ is a key step in improving the lower bounds on memory constrained inference.

Finally, we mention a problem in communication with memory constraints, in the same spirit as the problems surveyed in this article, that is worth studying. Consider the problem of communicating over n channel uses of a discrete memoryless channel (DMC) $P_{Y|X}$, when the decoder has limited memory of S states. The question is what is the largest achievable rate. Specifically, an (n, R, S, ε) -code for the DMC $P_{Y|X}$ is specified by an encoder $f: \llbracket 2^{nR} \rrbracket \rightarrow \mathcal{X}^n$, a decoder update function $g: \mathcal{Y} \times [S] \rightarrow [S]$ and a decision function $d: [S] \rightarrow \llbracket 2^{nR} \rrbracket$. For a message $W \sim \text{Uniform} \llbracket 2^{nR} \rrbracket$ the channel input is $X^n = f(W)$. The channel output Y^n is generated by passing X^n through $P_{Y|X}^{\otimes n}$. The update function is then applied recursively with $U_0 = 1$ and $U_t = g(Y_t, U_{t-1})$, $t = 1, \dots, n$. Finally, the estimate for W is $\hat{W} = d(U_n)$, and the error probability is $\varepsilon = \Pr(W \neq \hat{W})$. We denote

$$R^*(n, P_{Y|X}, S, \varepsilon) = \max\{R : \exists(n, R, S, \varepsilon) \text{ - code}\} \quad (50)$$

$$C_\varepsilon(P_{Y|X}, S) = \sup_{n \in \mathbb{N}} R^*(n, P_{Y|X}, S, \varepsilon) \quad (51)$$

and the goal is to characterize this quantity. Note that for finite output alphabet channels $C_\varepsilon(P_{Y|X}, S)$ is lower bounded by the finite-blocklength fundamental limit (without memory constraints) $R^*(\log S / \log |\mathcal{Y}|, P_{Y|X}, \infty, \varepsilon)$, since we can use the S memory states to save $n_S = \log S / \log |\mathcal{Y}|$ channel outputs. The interesting question here is how much better we can do, if at all, by using $n > n_S$.

Acknowledgements

This work was supported by the ISF under Grants 1641/21 and 1766/22.

References

- [1] C. Dwork, “Differential privacy: A survey of results,” in *International conference on theory and applications of models of computation*, pp. 1–19, Springer, 2008.
- [2] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” in *2013 IEEE 54th annual symposium on foundations of computer science*, pp. 429–438, IEEE, 2013.
- [3] J. Acharya, C. L. Canonne, C. Freitag, Z. Sun, and H. Tyagi, “Inference under information constraints iii: Local privacy constraints,” *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 253–267, 2021.
- [4] T. M. Cover, “Hypothesis testing with finite statistics,” *The Annals of Mathematical Statistics*, vol. 40, no. 3, pp. 828–835, 1969.
- [5] M. E. Hellman and T. M. Cover, “Learning with finite memory,” *The Annals of Mathematical Statistics*, pp. 765–782, 1970.
- [6] H. Robbins, “A sequential decision problem with a finite memory,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 42, no. 12, p. 920, 1956.
- [7] A. B. Tsybakov, “On nonparametric estimation of density level sets,” *The Annals of Statistics*, vol. 25, no. 3, pp. 948–969, 1997.
- [8] L. Devroye and G. Lugosi, *Combinatorial methods in density estimation*. Springer Science & Business Media, 2001.
- [9] E. L. Lehmann and G. Casella, *Theory of point estimation*. Springer Science & Business Media, 2006.
- [10] O. Goldreich, “Property testing,” *Lecture Notes in Comput. Sci.*, vol. 6390, 2010.
- [11] T. T. Cai and M. G. Low, “Testing composite hypotheses, hermite polynomials and optimal estimation of a nonsmooth functional,” 2011.
- [12] J. Jiao, K. Venkat, Y. Han, and T. Weissman, “Minimax estimation of functionals of discrete distributions,” *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2835–2885, 2015.

- [13] Y. Wu and P. Yang, “Minimax rates of entropy estimation on large alphabets via best polynomial approximation,” *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3702–3720, 2016.
- [14] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [15] J. Neyman and E. S. Pearson, “On the problem of the most efficient tests of statistical hypotheses,” *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 231, no. 694-706, pp. 289–337, 1933.
- [16] E. Tuncel, “On error exponents in hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 51, no. 8, pp. 2945–2950, 2005.
- [17] L. G. Valiant, “A theory of the learnable,” *Communications of the ACM*, vol. 27, no. 11, pp. 1134–1142, 1984.
- [18] T. M. Cover, “A note on the two-armed bandit problem with finite memory,” *Information and Control*, vol. 12, no. 5, pp. 371–377, 1968.
- [19] J. Koplowitz, “Necessary and sufficient memory size for m-hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 21, no. 1, pp. 44–46, 1975.
- [20] T. M. Cover, M. A. Freedman, and M. E. Hellman, “Optimal finite memory learning algorithms for the finite sample problem,” *Information and Control*, vol. 30, no. 1, pp. 49–85, 1976.
- [21] N. Alon, Y. Matias, and M. Szegedy, “The space complexity of approximating the frequency moments,” *Journal of Computer and system sciences*, vol. 58, no. 1, pp. 137–147, 1999.
- [22] P. Indyk and D. Woodruff, “Optimal approximations of the frequency moments of data streams,” in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pp. 202–208, 2005.
- [23] A. Lall, V. Sekar, M. Ogihara, J. Xu, and H. Zhang, “Data streaming algorithms for estimating entropy of network traffic,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 34, no. 1, pp. 145–156, 2006.
- [24] A. Chakrabarti, K. Do Ba, and S. Muthukrishnan, “Estimating entropy and entropy norm on data streams,” *Internet Mathematics*, vol. 3, no. 1, pp. 63–78, 2006.

- [25] A. Chakrabarti, G. Cormode, and A. McGregor, “A near-optimal algorithm for estimating the entropy of a stream,” *ACM Transactions on Algorithms (TALG)*, vol. 6, no. 3, pp. 1–21, 2010.
- [26] M. Charikar, K. Chen, and M. Farach-Colton, “Finding frequent items in data streams,” in *International Colloquium on Automata, Languages, and Programming*, pp. 693–703, Springer, 2002.
- [27] G. Cormode and S. Muthukrishnan, “An improved data stream summary: the count-min sketch and its applications,” *Journal of Algorithms*, vol. 55, no. 1, pp. 58–75, 2005.
- [28] A. Metwally, D. Agrawal, and A. El Abbadi, “Efficient computation of frequent and top-k elements in data streams,” in *International conference on database theory*, pp. 398–412, Springer, 2005.
- [29] P. Flajolet, “Approximate counting: a detailed analysis,” *BIT Numerical Mathematics*, vol. 25, no. 1, pp. 113–134, 1985.
- [30] M. Durand and P. Flajolet, “Loglog counting of large cardinalities,” in *Algorithms-ESA 2003: 11th Annual European Symposium, Budapest, Hungary, September 16-19, 2003. Proceedings 11*, pp. 605–617, Springer, 2003.
- [31] P. Indyk and D. Woodruff, “Tight lower bounds for the distinct elements problem,” in *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pp. 283–288, IEEE, 2003.
- [32] P. Flajolet, É. Fusy, O. Gandouet, and F. Meunier, “Hyperloglog: the analysis of a near-optimal cardinality estimation algorithm,” *Discrete mathematics & theoretical computer science*, no. Proceedings, 2007.
- [33] D. M. Kane, J. Nelson, and D. P. Woodruff, “An optimal algorithm for the distinct elements problem,” in *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 41–52, 2010.
- [34] S. Guha, A. McGregor, and S. Venkatasubramanian, “Sublinear estimation of entropy and information distances,” *ACM Transactions on Algorithms (TALG)*, vol. 5, no. 4, pp. 1–16, 2009.
- [35] A. Chakrabarti, T. Jayram, and M. Pătraşcu, “Tight lower bounds for selection in randomly ordered streams,” in *Proceedings of the nineteenth annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 720–729, 2008.
- [36] J. Acharya, S. Bhadane, P. Indyk, and Z. Sun, “Estimating entropy of distributions in constant space,” *arXiv preprint arXiv:1911.07976*, 2019.

- [37] M. Aliakbarpour, A. McGregor, J. Nelson, and E. Waingarten, “Estimation of entropy in constant space with improved sample complexity,” *arXiv preprint arXiv:2205.09804*, 2022.
- [38] D. P. Woodruff, “Optimal space lower bounds for all frequency moments,” in *SODA*, vol. 4, pp. 167–175, Citeseer, 2004.
- [39] E. Meron and M. Feder, “Finite-memory universal prediction of individual sequences,” *IEEE Transactions on Information Theory*, vol. 50, no. 7, pp. 1506–1523, 2004.
- [40] D. P. Woodruff, “The average-case complexity of counting distinct elements,” in *Proceedings of the 12th International Conference on Database Theory*, pp. 284–295, 2009.
- [41] Z. Zhang and T. Berger, “Estimation via compressed information,” *IEEE transactions on Information theory*, vol. 34, no. 2, pp. 198–211, 1988.
- [42] S.-i. Amari, “Fisher information under restriction of shannon information in multi-terminal situations,” *Annals of the Institute of Statistical Mathematics*, vol. 41, pp. 623–648, 1989.
- [43] R. Ahlswede and M. V. Burnashev, “On minimax estimation in the presence of side information about remote data,” *The Annals of Statistics*, pp. 141–171, 1990.
- [44] S. Amari *et al.*, “Parameter estimation with multiterminal data compression,” *IEEE transactions on Information Theory*, vol. 41, no. 6, pp. 1802–1833, 1995.
- [45] S. Amari *et al.*, “Statistical inference under multiterminal data compression,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2300–2324, 1998.
- [46] S.-i. Amari, “On optimal data compression in multiterminal statistical inference,” *IEEE Transactions on Information Theory*, vol. 57, no. 9, pp. 5577–5587, 2011.
- [47] R. Ahlswede and I. Csiszár, “Hypothesis testing with communication constraints,” *IEEE transactions on information theory*, vol. 32, no. 4, pp. 533–542, 1986.
- [48] S.-I. Amari and T. S. Han, “Statistical inference under multiterminal rate restrictions: A differential geometric approach,” *IEEE Transactions on Information Theory*, vol. 35, no. 2, pp. 217–227, 1989.

- [49] H. M. Shalaby and A. Papamarcou, “Multiterminal detection with zero-rate data compression,” *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 254–267, 1992.
- [50] M. Wigger and R. Timo, “Testing against independence with multiple decision centers,” in *2016 International Conference on Signal Processing and Communications (SPCOM)*, pp. 1–5, IEEE, 2016.
- [51] S. Watanabe, “Neyman–pearson test for zero-rate multiterminal hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 64, no. 7, pp. 4923–4939, 2017.
- [52] N. Weinberger and Y. Kochman, “On the reliability function of distributed hypothesis testing under optimal detection,” *IEEE Transactions on Information Theory*, vol. 65, no. 8, pp. 4940–4965, 2019.
- [53] S. Sreekumar and D. Gündüz, “Distributed hypothesis testing over discrete memoryless channels,” *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 2044–2066, 2019.
- [54] S. Salehkalaibar and M. Wigger, “Distributed hypothesis testing with variable-length coding,” *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 3, pp. 681–694, 2020.
- [55] Y. Xiang and Y.-H. Kim, “Interactive hypothesis testing against independence,” in *2013 IEEE International Symposium on Information Theory*, pp. 2840–2844, IEEE, 2013.
- [56] K. Sahasranand and H. Tyagi, “Extra samples can reduce the communication for independence testing,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 2316–2320, IEEE, 2018.
- [57] U. Hadar, J. Liu, Y. Polyanskiy, and O. Shayevitz, “Communication complexity of estimating correlations,” in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 792–803, 2019.
- [58] U. Hadar and O. Shayevitz, “Distributed estimation of gaussian correlations,” *IEEE Transactions on Information Theory*, vol. 65, no. 9, pp. 5323–5338, 2019.
- [59] A. C.-C. Yao, “Some complexity questions related to distributive computing (preliminary report),” in *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pp. 209–213, 1979.
- [60] E. Kushilevitz, “Communication complexity,” in *Advances in Computers*, vol. 44, pp. 331–360, Elsevier, 1997.

- [61] A. Rao and A. Yehudayoff, *Communication Complexity: and Applications*. Cambridge University Press, 2020.
- [62] Y. Zhang, J. Duchi, M. I. Jordan, and M. J. Wainwright, “Information-theoretic lower bounds for distributed statistical estimation with communication constraints,” *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [63] A. Garg, T. Ma, and H. Nguyen, “On communication cost of distributed statistical estimation and dimensionality,” *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [64] J. Acharya, C. L. Canonne, Y. Han, Z. Sun, and H. Tyagi, “Domain compression and its application to randomness-optimal distributed goodness-of-fit,” in *Conference on Learning Theory*, pp. 3–40, PMLR, 2020.
- [65] M. Braverman, A. Garg, T. Ma, H. L. Nguyen, and D. P. Woodruff, “Communication lower bounds for statistical estimation problems via a distributed data processing inequality,” in *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 1011–1020, 2016.
- [66] A. Xu and M. Raginsky, “Information-theoretic lower bounds on bayes risk in decentralized estimation,” *IEEE Transactions on Information Theory*, vol. 63, no. 3, pp. 1580–1600, 2016.
- [67] M. I. Jordan, J. D. Lee, and Y. Yang, “Communication-efficient distributed statistical inference,” *Journal of the American Statistical Association*, 2018.
- [68] Y. Han, A. Özgür, and T. Weissman, “Geometric lower bounds for distributed parameter estimation under communication constraints,” in *Conference On Learning Theory*, pp. 3163–3188, PMLR, 2018.
- [69] O. Fischer, U. Meir, and R. Oshman, “Distributed uniformity testing,” in *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pp. 455–464, 2018.
- [70] J. Acharya, C. L. Canonne, and H. Tyagi, “Inference under information constraints: Lower bounds from chi-square contraction,” in *Conference on Learning Theory*, pp. 3–17, PMLR, 2019.
- [71] J. Acharya, C. L. Canonne, and H. Tyagi, “Inference under information constraints ii: Communication constraints and shared randomness,” *IEEE Transactions on Information Theory*, vol. 66, no. 12, pp. 7856–7877, 2020.

- [72] A. Pensia, V. Jog, and P.-L. Loh, “Communication-constrained hypothesis testing: Optimality, robustness, and reverse data processing inequalities,” *arXiv preprint arXiv:2206.02765*, 2022.
- [73] T. Berger, Z. Zhang, and H. Viswanathan, “The ceo problem [multiterminal source coding],” *IEEE Transactions on Information Theory*, vol. 42, no. 3, pp. 887–902, 1996.
- [74] H. Viswanathan and T. Berger, “The quadratic gaussian ceo problem,” *IEEE Transactions on Information Theory*, vol. 43, no. 5, pp. 1549–1559, 1997.
- [75] K. Eswaran and M. Gastpar, “Remote source coding under gaussian noise: Dueling roles of power and entropy power,” *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4486–4498, 2019.
- [76] Y. Polyanskiy and Y. Wu, “Information theory: From coding to learning,” *Book draft*, 2022.
- [77] D. P. Woodruff, *Efficient and private distance approximation in the communication and streaming models*. PhD thesis, Massachusetts Institute of Technology, 2007.
- [78] D. M. Kane, J. Nelson, and D. P. Woodruff, “On the exact space complexity of sketching and streaming small norms,” in *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pp. 1161–1178, SIAM, 2010.
- [79] T. S. Jayram and D. P. Woodruff, “Optimal bounds for johnson-lindenstrauss transforms and streaming problems with subconstant error,” *ACM Transactions on Algorithms (TALG)*, vol. 9, no. 3, pp. 1–17, 2013.
- [80] M. Crouch, A. McGregor, G. Valiant, and D. P. Woodruff, “Stochastic streams: Sample complexity vs. space complexity,” in *24th Annual European Symposium on Algorithms (ESA 2016)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [81] Y. Dagan and O. Shamir, “Detecting correlations with little memory and communication,” in *Conference On Learning Theory*, pp. 1145–1198, PMLR, 2018.
- [82] I. Dinur, “On the streaming indistinguishability of a random permutation and a random function,” in *Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II 30*, pp. 433–460, Springer, 2020.

- [83] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” *arXiv preprint physics/0004057*, 2000.
- [84] Z. Goldfeld and Y. Polyanskiy, “The information bottleneck problem and its applications in machine learning,” *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 19–38, 2020.
- [85] I. Diakonikolas, T. Gouleakis, D. M. Kane, and S. Rao, “Communication and memory efficient testing of discrete distributions,” in *Conference on Learning Theory*, pp. 1070–1106, PMLR, 2019.
- [86] I. Shahaf, O. Ordentlich, and G. Segev, “An information-theoretic proof of the streaming switching lemma for symmetric encryption,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 858–863, IEEE, 2020.
- [87] J. Jaeger and S. Tessaro, “Tight time-memory trade-offs for symmetric encryption,” in *Advances in Cryptology—EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I 38*, pp. 467–497, Springer, 2019.
- [88] O. Shamir, “Fundamental limits of online and distributed algorithms for statistical learning and estimation,” *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [89] J. Steinhardt, G. Valiant, and S. Wager, “Memory, communication, and statistical queries,” in *Conference on Learning Theory*, pp. 1490–1516, PMLR, 2016.
- [90] T. Berg, O. Ordentlich, and O. Shayevitz, “Binary hypothesis testing with deterministic finite-memory decision rules,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 1259–1264, IEEE, 2020.
- [91] R. Raz, “Fast learning requires good memory: A time-space lower bound for parity learning,” 2016.
- [92] T. Berg, O. Ordentlich, and O. Shayevitz, “Deterministic finite-memory bias estimation,” in *Conference on Learning Theory*, pp. 566–585, PMLR, 2021.
- [93] T. Berg, O. Ordentlich, and O. Shayevitz, “On the memory complexity of uniformity testing,” in *Conference on Learning Theory*, pp. 3506–3523, PMLR, 2022.
- [94] M. Horstein, “Sequential transmission using noiseless feedback,” *IEEE Transactions on Information Theory*, vol. 9, no. 3, pp. 136–143, 1963.

- [95] R. Nowak, “Generalized binary search,” in *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pp. 568–574, IEEE, 2008.
- [96] O. Shayevitz and M. Feder, “Optimal feedback communication via posterior matching,” *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1186–1222, 2011.
- [97] T. Wagner, “Estimation of the mean with time-varying finite memory (corresp.),” *IEEE Transactions on Information Theory*, vol. 18, no. 4, pp. 523–525, 1972.
- [98] J. Isbell, “On a problem of robbins,” *The Annals of Mathematical Statistics*, vol. 30, no. 2, pp. 606–610, 1959.
- [99] C. V. Smith and R. Pyke, “The robbins-isbell two-armed-bandit problem with finite memory,” *The Annals of Mathematical Statistics*, pp. 1375–1386, 1965.
- [100] S. M. Samuels, “Randomized rules for the two-armed-bandit with finite memory,” *The Annals of Mathematical Statistics*, vol. 39, no. 6, pp. 2103–2107, 1968.
- [101] M. E. Hellman and T. M. Cover, “On memory saved by randomization,” *The Annals of Mathematical Statistics*, vol. 42, no. 3, pp. 1075–1078, 1971.
- [102] M. Hellman, “The effects of randomization on finite-memory decision schemes,” *IEEE Transactions on Information Theory*, vol. 18, no. 4, pp. 499–502, 1972.
- [103] M. Hellman and T. Cover, “A review of recent results on learning with finite memory,” in *2nd International Symposium on Information Theory*, pp. 289–294, 1973.
- [104] B. Shubert and C. Anderson, “Testing a simple symmetric hypothesis by a finite-memory deterministic algorithm,” *IEEE Transactions on Information Theory*, vol. 19, no. 5, pp. 644–647, 1973.
- [105] G. Kol, R. Raz, and A. Tal, “Time-space hardness of learning sparse parities,” in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1067–1080, 2017.
- [106] R. Raz, “A time-space lower bound for a large class of learning problems,” in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 732–742, IEEE, 2017.

- [107] S. Garg, P. K. Kothari, P. Liu, and R. Raz, “Memory-sample lower bounds for learning parity with noise,” *arXiv preprint arXiv:2107.02320*, 2021.
- [108] T. Cover and M. Hellman, “The two-armed-bandit problem with time-invariant finite memory,” *IEEE Transactions on Information Theory*, vol. 16, no. 2, pp. 185–195, 1970.
- [109] J. Horos and M. Hellman, “A confidence model for finite-memory learning systems (corresp.),” *IEEE Transactions on Information Theory*, vol. 18, no. 6, pp. 811–813, 1972.
- [110] R. Flower and M. Hellman, “Hypothesis testing with finite memory in finite time (corresp.),” *IEEE Transactions on Information Theory*, vol. 18, no. 3, pp. 429–431, 1972.
- [111] R. Roberts and J. Tooley, “Estimation with finite memory,” *IEEE Transactions on Information Theory*, vol. 16, no. 6, pp. 685–691, 1970.
- [112] J. Koplowitz and R. Roberts, “Sequential estimation with a finite statistic,” *IEEE Transactions on Information Theory*, vol. 19, no. 5, pp. 631–635, 1973.
- [113] M. E. Hellman, “Finite-memory algorithms for estimating the mean of a gaussian distribution (corresp.),” *IEEE Transactions on Information Theory*, vol. 20, no. 3, pp. 382–384, 1974.
- [114] F. J. Samaniego, “Estimating a binomial parameter with finite memory,” *IEEE Transactions on Information Theory*, vol. 19, no. 5, pp. 636–643, 1973.
- [115] F. Leighton and R. Rivest, “Estimating a probability using finite memory,” *IEEE Transactions on Information Theory*, vol. 32, no. 6, pp. 733–742, 1986.
- [116] V. Anantharam and P. Tsoucas, “A proof of the markov chain tree theorem,” *Statistics & Probability Letters*, vol. 8, no. 2, pp. 189–192, 1989.
- [117] J. Von Neumann, “13. various techniques used in connection with random digits,” *Appl. Math Ser*, vol. 12, no. 36-38, p. 5, 1951.
- [118] J. Steinhardt and J. Duchi, “Minimax rates for memory-bounded sparse linear regression,” in *Conference on Learning Theory*, pp. 1564–1587, PMLR, 2015.
- [119] B. Yu, “Assouad, fano, and le cam,” in *Festschrift for Lucien Le Cam: research papers in probability and statistics*, pp. 423–435, Springer, 1997.
- [120] V. Sharan, A. Sidford, and G. Valiant, “Memory-sample tradeoffs for linear regression with small error,” in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 890–901, 2019.

- [121] A. Jain and H. Tyagi, “Effective memory shrinkage in estimation,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1071–1075, IEEE, 2018.
- [122] L. A. Kontorovich, “Statistical estimation with bounded memory,” *Statistics and Computing*, vol. 22, no. 5, pp. 1155–1164, 2012.
- [123] N. Merhav and M. Feder, “Universal schemes for sequential decision from individual data sequences,” *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1280–1292, 1993.
- [124] N. Merhav, M. Feder, and M. Gutman, “Some properties of sequential predictors for binary markov sources,” *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 887–892, 1993.
- [125] D. Rajwan and M. Feder, “Universal finite memory machines for coding binary sequences,” in *Proceedings DCC 2000. Data Compression Conference*, pp. 113–122, IEEE, 2000.
- [126] E. Meron and M. Feder, “Optimal finite state universal coding of individual sequences,” in *Data Compression Conference, 2004. Proceedings. DCC 2004*, pp. 332–341, IEEE, 2004.
- [127] O. Goldreich, “The uniform distribution is complete with respect to testing identity to a fixed distribution,” in *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 23, p. 1, 2016.
- [128] O. Goldreich and D. Ron, “On testing expansion in bounded-degree graphs,” in *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 20, 2000.
- [129] L. Paninski, “A coincidence-based test for uniformity given very sparsely sampled discrete data,” *IEEE Transactions on Information Theory*, vol. 54, no. 10, pp. 4750–4755, 2008.
- [130] I. Diakonikolas, D. M. Kane, and V. Nikishkin, “Testing identity of structured distributions,” in *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1841–1854, SIAM, 2014.
- [131] U. Meir, “Comparison Graphs: A Unified Method for Uniformity Testing,” in *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, vol. 185 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 17:1–17:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

- [132] C. L. Canonne and J. Q. Yang, “Simpler distribution testing with little memory,” in *2024 Symposium on Simplicity in Algorithms (SOSA)*, pp. 406–416, SIAM, 2024.
- [133] I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price, “Sample-optimal identity testing with high probability,” in *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [134] S. Roy and Y. Vasudev, “Testing properties of distributions in the streaming model,” *arXiv preprint arXiv:2309.03245*, 2023.
- [135] S. Chakraborty, E. Fischer, Y. Goldhirsh, and A. Matsliah, “On the power of conditional samples in distribution testing,” in *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 561–580, 2013.
- [136] C. L. Canonne, D. Ron, and R. A. Servedio, “Testing probability distributions using conditional samples,” *SIAM Journal on Computing*, vol. 44, no. 3, pp. 540–616, 2015.
- [137] G. P. Basharin, “On a statistical estimate for the entropy of a sequence of independent random variables,” *Theory of Probability & Its Applications*, vol. 4, no. 3, pp. 333–336, 1959.
- [138] L. Paninski, “Estimating entropy on m bins given fewer than m samples,” *IEEE Transactions on Information Theory*, vol. 50, no. 9, pp. 2200–2203, 2004.
- [139] G. Valiant and P. Valiant, “A clt and tight lower bounds for estimating entropy,” in *Electron. Colloquium Comput. Complex.*, vol. 17, p. 179, 2010.
- [140] 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 *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 685–694, 2011.
- [141] G. Valiant and P. Valiant, “The power of linear estimators,” in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 403–412, IEEE, 2011.
- [142] S. Chien, K. Ligett, and A. McGregor, “Space-efficient estimation of robust statistics and distribution testing,” in *ICS*, pp. 251–265, Citeseer, 2010.
- [143] T. Berg, O. Ordentlich, and O. Shayevitz, “Entropy estimation with less memory,” in *Draft - to be submitted*, 2023.

- [144] R. Morris, “Counting large numbers of events in small registers,” *Communications of the ACM*, vol. 21, no. 10, pp. 840–842, 1978.
- [145] J. Nelson and H. Yu, “Optimal bounds for approximate counting,” *arXiv preprint arXiv:2010.02116*, 2020.