
Supplementary material: The Broad Optimality of Profile Maximum Likelihood

Yi Hao

Dept. of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093
yih179@ucsd.edu

Alon Orlitsky

Dept. of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093
alon@ucsd.edu

Abstract

We study three fundamental statistical learning problems: distribution estimation, property estimation, and property testing. We establish the profile maximum likelihood (PML) estimator as the first unified sample-optimal approach to a wide range of learning tasks. In particular, for every alphabet size k and desired accuracy ε :

Distribution estimation Under ℓ_1 distance, PML yields optimal $\Theta(k/(\varepsilon^2 \log k))$ sample complexity for sorted distribution estimation, and a PML-based estimator empirically outperforms the Good-Turing estimator on the actual distribution;

Additive property estimation For a broad class of additive properties, the PML plug-in estimator uses just four times the sample size required by the best estimator to achieve roughly twice its error, with exponentially higher confidence;

α -Rényi entropy estimation For an integer $\alpha > 1$, the PML plug-in estimator has optimal $k^{1-1/\alpha}$ sample complexity; for non-integer $\alpha > 3/4$, the PML plug-in estimator has sample complexity lower than the state of the art;

Identity testing In testing whether an unknown distribution is equal to or at least ε far from a given distribution in ℓ_1 distance, a PML-based tester achieves the optimal sample complexity up to logarithmic factors of k .

With minor modifications, most of these results also hold for a near-linear-time computable variant of PML.

Contents

1	Introduction	3
1.1	Problems of interest	3
1.2	Profile maximum likelihood	4
2	New results and implications	5
2.1	New results	5
2.2	Implications	6
3	Related work and comparisons	7
3.1	Additive property estimation	7
3.2	Comparison I: Theorem 1 and related property-estimation work	8
3.3	Rényi entropy estimation	8
3.4	Comparison II: Theorem 2 to 4 and related Rényi-entropy-estimation work	8
3.5	(Sorted) distribution estimation	9
3.6	Comparison III: Theorem 5 and related distribution-estimation work	9
3.7	Identity testing	9
4	Numerical experiments	10
4.1	MCMC-EM algorithm variant	10
4.2	Experiment distributions	11
4.3	Experiment results and details	11
5	Lipschitz-property estimation	15
5.1	Proof outline of Theorem 1	15
5.2	Technical details	15
6	α-Rényi entropy estimation	19
6.1	Proof of Theorem 2: $\alpha \in (3/4, 1)$	19
6.2	Proof of Theorem 3: Non-integer $\alpha > 1$	21
6.3	Proof of Theorem 4: Integer $\alpha > 1$	23
7	Sorted distribution estimation	25
7.1	Sorted ℓ_1 distance and Wasserstein duality	25
7.2	Proof of Theorem 5	26
8	Uniformity testing	28
8.1	PML-based tester	28
8.2	Proof of Theorem 6	29
A	Proof of Lemma 3	30

1 Introduction

A distribution p over a discrete alphabet \mathcal{X} of size k corresponds to an element of the simplex

$$\Delta_{\mathcal{X}} := \left\{ p \in \mathbb{R}_{\geq 0}^k : \sum_{x \in \mathcal{X}} p(x) = 1 \right\}.$$

A distribution *property* is a mapping $f : \Delta_{\mathcal{X}} \rightarrow \mathbb{R}$ associating a real value with each distribution. A distribution property f is *symmetric* if it is invariant under domain-symbol permutations. A symmetric property is *additive*, i.e., additively separable, if it can be written as $f(p) := \sum_x f(p(x))$, where for simplicity we use f to denote both the property and the corresponding real function.

Many important symmetric properties are additive. For example,

- **Support size** $S(p) := \sum_x \mathbb{1}_{p(x) > 0}$, a fundamental quantity arising in the study of vocabulary size [33, 61, 77], population estimation [39, 59], and database studies [42].
- **Support coverage** $C_m(p) := \sum_x (1 - (1 - p(x))^m)$, where m is a given parameter, the expected number of distinct elements observed in a sample of size m , arising in biological [20, 56] and ecological [20–22, 27] research;
- **Shannon entropy** $H(p) := -\sum_x p(x) \log p(x)$, the primary measure of information [28, 76] with numerous applications to machine learning [15, 26, 73] and neuroscience [34, 58];
- **Distance to uniformity** $D(p) := \|p - p_u\|_1$, where p_u is the uniform distribution over $\Delta_{\mathcal{X}}$, a property being central to the field of distribution property testing [11, 13, 16, 75].

Besides being additive and symmetric, these four properties have yet another attribute in common. Under the appropriate interpretation, they are also all 1-Lipschitz. Specifically, for two distributions $p, q \in \Delta_{\mathcal{X}}$, let $\Gamma_{p,q}$ be the collection of distributions over $\mathcal{X} \times \mathcal{X}$ with marginals p and q on the first and second factors respectively. The *relative earth-mover distance* [80], between p and q is

$$R(p, q) := \inf_{\gamma \in \Gamma_{p,q}} \mathbb{E}_{(X,Y) \sim \gamma} \left| \log \frac{p(X)}{q(Y)} \right|.$$

One can verify [80, 81] that H , D , and $\tilde{C}_m := C_m/m$ are all 1-Lipschitz on the metric space $(\Delta_{\mathcal{X}}, R)$, and $\tilde{S} := S/k$ is 1-Lipschitz over $(\Delta_{\geq 1/k}, R)$, the set of distributions in $\Delta_{\mathcal{X}}$ whose nonzero probabilities are at least $1/k$. We will study all such Lipschitz properties in later sections.

An important symmetric non-additive property is *Rényi entropy*, a well-known measure of randomness with numerous applications to unsupervised learning [50, 89] and image registration [57, 62]. For a distribution $p \in \Delta_{\mathcal{X}}$ and a non-negative real parameter $\alpha \neq 1$, the α -*Rényi entropy* [74] of p is $H_{\alpha}(p) := (1 - \alpha)^{-1} \log (\sum_x p_x^{\alpha})$. In particular, denoted by $H_1(p) := \lim_{\alpha \rightarrow 1} H_{\alpha}(p)$, the 1-*Rényi entropy* is exactly Shannon entropy [74].

1.1 Problems of interest

In this work, we consider three fundamental statistical learning problems concerning the estimation and testing of distributions and their properties.

(Sorted) distribution estimation

A natural learning problem is to estimate an unknown distribution $p \in \Delta_{\mathcal{X}}$ from an i.i.d. sample $X^n \sim p$. For any two distributions $p, q \in \Delta_{\mathcal{X}}$, let $\ell(p, q)$ be the *loss* when we approximate p by q . A *distribution estimator* $\hat{p} : \mathcal{X}^* \rightarrow \Delta_{\mathcal{X}}$ associates every sequence $x^n \in \mathcal{X}^*$ with a distribution $\hat{p}(x^n)$. We measure the performance of an estimator by its *sample complexity*

$$n(\hat{p}, \varepsilon, \delta) := \min\{n : \forall p \in \Delta_{\mathcal{X}}, \Pr_{X^n \sim p} (\ell(p, \hat{p}(X^n)) \geq \varepsilon) \leq \delta\},$$

the smallest sample size that \hat{p} requires to estimate all distributions in $\Delta_{\mathcal{X}}$ to a desired accuracy $\varepsilon > 0$, with error probability $\delta \in (0, 1)$. The sample complexity of distribution estimation over $\Delta_{\mathcal{X}}$ is

$$n(\varepsilon, \delta) := \min\{n(\hat{p}, \varepsilon, \delta) : \hat{p} : \mathcal{X}^* \rightarrow \Delta_{\mathcal{X}}\},$$

the lowest sample complexity of any estimator. For simplicity, we will omit δ when $\delta = 1/3$.

For a distribution $p \in \Delta_{\mathcal{X}}$, we denote by $\{p\}$ the multiset of its probabilities. The *sorted ℓ_1 distance* between two distributions $p, q \in \Delta_{\mathcal{X}}$ is

$$\ell_1^{\leq}(p, q) := \min_{p' \in \Delta_{\mathcal{X}}: \{p'\} = \{p\}} \|p' - q\|_1,$$

the smallest ℓ_1 distance between q and any sorted version of p . As illustrated in Section 7.1, this is essentially the 1-Wasserstein distance between uniform measures on the probability multisets $\{p\}$ and $\{q\}$. We will consider both the sorted and unsorted ℓ_1 distances.

Property estimation

Often we would like to estimate a given property f of an unknown distribution $p \in \Delta_{\mathcal{X}}$ based on a sample $X^n \sim p$. A *property estimator* is a mapping $\hat{f} : \mathcal{X}^* \rightarrow \mathbb{R}$. Analogously, the *sample complexity* of \hat{f} in estimating f over a set $\mathcal{P} \subset \Delta_{\mathcal{X}}$ is

$$n_f(\hat{f}, \mathcal{P}, \varepsilon, \delta) := \min\{n: \forall p \in \mathcal{P}, \Pr_{X^n \sim p} (|\hat{f}(X^n) - f(p)| \geq \varepsilon) \leq \delta\},$$

the smallest sample size that \hat{f} requires to estimate f with accuracy ε and confidence $1 - \delta$, for all distributions in \mathcal{P} . The sample complexity of estimating f over \mathcal{P} is

$$n_f(\mathcal{P}, \varepsilon, \delta) := \min\{n_f(\hat{f}, \mathcal{P}, \varepsilon, \delta): \hat{f}: \mathcal{X}^* \rightarrow \mathbb{R}\},$$

the lowest sample complexity of any estimator. For simplicity, we will omit \mathcal{P} when $\mathcal{P} = \Delta_{\mathcal{X}}$, and omit δ when $\delta = 1/3$. The standard ‘‘median trick’’ shows that $\log(1/\delta) \cdot n_f(\mathcal{P}, \varepsilon) \geq \Omega(n_f(\mathcal{P}, \varepsilon, \delta))$. By convention, we say an estimator \hat{f} is *sample-optimal* if $n_f(\hat{f}, \mathcal{P}, \varepsilon) = \Theta(n_f(\mathcal{P}, \varepsilon))$.

Property testing: Identity testing

A closely related problem is distribution property testing, of which identity testing is the most fundamental and well-studied [16, 36]. Given an error parameter ε , a distribution q , and a sample X^n from an unknown distribution p , *identity testing* aims to distinguish between the null hypothesis

$$H_0 : p = q$$

and the alternative hypothesis

$$H_1 : \|p - q\|_1 \geq \varepsilon.$$

A *property tester* is a mapping $\hat{t} : \mathcal{X}^* \rightarrow \{0, 1\}$, indicating whether H_0 or H_1 is accepted. Analogous to the two formulations above, the *sample complexity* of \hat{t} is

$$n_q(\hat{t}, \varepsilon, \delta) := \min\{n: \forall i \in \{0, 1\} \text{ and } \forall p \in H_i, \Pr_{X^n \sim p} (\hat{t}(X^n) \neq i) \leq \delta\},$$

and the sample complexity of identity testing with respect to q is

$$n_q(\varepsilon, \delta) := \min\{n(\hat{t}, \varepsilon, \delta): \hat{t}: \mathcal{X}^* \rightarrow \{0, 1\}\}.$$

Again, when $\delta = 1/3$, we will omit δ . For $q = p_u$, the problem is also known as *uniformity testing*.

1.2 Profile maximum likelihood

The *multiplicity* of a symbol $x \in \mathcal{X}$ in a sequence $x^n := x_1, \dots, x_n \in \mathcal{X}^*$ is $\mu_x(x^n) := |\{j : x_j = x, 1 \leq j \leq n\}|$, the number of times x appears in x^n . These multiplicities induce an *empirical distribution* $p_\mu(x^n)$ that associates a probability $\mu_x(x^n)/n$ with each symbol $x \in \mathcal{X}$.

The *prevalence* of an integer $i \geq 0$ in x^n is the number $\varphi_i(x^n)$ of symbols appearing i times in x^n . For known \mathcal{X} , the value of φ_0 can be deduced from the remaining multiplicities, hence we define the *profile* of x^n to be $\varphi(x^n) = (\varphi_1(x^n), \dots, \varphi_n(x^n))$, the vector of all positive prevalences. For example, $\varphi(\text{alfalfa}) = (0, 2, 1, 0, 0, 0, 0)$. Note that the profile of x^n also corresponds to the multiset of multiplicities of distinct symbols in x^n .

For a distribution $p \in \Delta_{\mathcal{X}}$, let

$$p(x^n) := \Pr_{X^n \sim p} (X^n = x^n)$$

be the probability of observing a sequence x^n under i.i.d. sampling from p , and let

$$p(\varphi) := \sum_{y^n: \varphi(y^n)=\varphi} p(y^n)$$

be the probability of observing a profile φ . While the sequence maximum likelihood estimator maps a sequence to its empirical distribution, which maximizes the sequence probability $p(x^n)$, the *profile maximum likelihood (PML)* estimator [66] over a set $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$ maps each profile φ to a distribution

$$p_\varphi := \arg \max_p p(\varphi)$$

that maximizes the profile probability. Relaxing the optimization objective, for any $\beta \in (0, 1)$, a β -*approximate PML* estimator [5] maps each profile φ to a distribution p_φ^β such that $p_\varphi^\beta(\varphi) \geq \beta \cdot p_\varphi(\varphi)$.

Originating from the principle of maximum likelihood, PML was proved [2, 5, 7, 8, 29, 66] to possess a number of useful attributes, such as existence over finite discrete domains, majorization by empirical distributions, consistency for distribution estimation under both sorted and unsorted ℓ_1 distances, and competitiveness to other profile-based estimators.

Let ε be an error parameter and f be one of the four properties in the introduction. Set $n := n_f(\varepsilon)$. Recently, Acharya et al. [5] showed that for some absolute constant $c' > 0$, if $c < c'$ and $\varepsilon \geq n^{-c}$, then a plug-in estimator for f , using an $\exp(-n^{1-\Theta(c)})$ -approximate PML, is sample-optimal. Motivated by this result, Charikar et al. [23] constructed an explicit $\exp(-\mathcal{O}(n^{2/3} \log^3 n))$ -approximate PML (APML) whose computation time is near-linear in n . Combined, these results provide a unified, sample-optimal, and near-linear-time computable plug-in estimator for the four properties.

2 New results and implications

2.1 New results

Additive property estimation

Recall that for any property f , the expression $n_f(\varepsilon)$ denotes the smallest sample size required by any estimator to achieve accuracy ε with confidence $2/3$, for all distributions in $\Delta_{\mathcal{X}}$. Let f be an additive symmetric property that is 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$. Let $\varepsilon > 0$ and $n \geq n_f(\varepsilon)$ be error and sampling parameters. For an absolute constant $c \in (10^{-2}, 10^{-1})$, if $\varepsilon \geq n^{-c}$,

Theorem 1. *The PML plug-in estimator, when given a sample of size $4n$ from any distribution $p \in \Delta_{\mathcal{X}}$, will estimate $f(p)$ up to an error of $(2 + o(1))\varepsilon$, with probability at least $1 - \exp(-4\sqrt{n})$.*

For a different $c > 0$, Theorem 1 also holds for APML, which is near-linear-time computable [23].

Rényi entropy estimation

For \mathcal{X} of finite size k and any $p \in \Delta_{\mathcal{X}}$, it is well-known that $H_\alpha(p) \in [0, \log k]$. The following theorems characterize the performance of the PML plug-in estimator in estimating Rényi entropy.

For any distribution $p \in \Delta_{\mathcal{X}}$, error parameter $\varepsilon \in (0, 1)$, absolute constant $\lambda \in (0, 0.1)$, and sampling parameter n , draw a sample $X^n \sim p$ and denote its profile by φ . Then for sufficiently large k ,

Theorem 2. *For $\alpha \in (3/4, 1)$, if $n = \Omega_\alpha(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$,*

$$\Pr(|H_\alpha(p_\varphi) - H_\alpha(p)| \geq \varepsilon) \leq \exp(-\sqrt{n}).$$

Theorem 3. *For non-integer $\alpha > 1$, if $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$,*

$$\Pr(|H_\alpha(p_\varphi) - H_\alpha(p)| \geq \varepsilon) \leq \exp(-n^{1-\lambda}).$$

Theorem 4. *For integer $\alpha > 1$, if $n = \Omega_\alpha(k^{1-1/\alpha}(\varepsilon^2 \log(1/\varepsilon))^{-(1+\alpha)})$ and $H_\alpha(p) \leq (\log n)/4$,*

$$\Pr(|H_\alpha(p_\varphi) - H_\alpha(p)| \geq \varepsilon) \leq 1/3.$$

Replacing $3/4$ by $5/6$, Theorem 2 also holds for APML with a better probability bound $\exp(-n^{2/3})$. In addition, Theorem 3 holds for APML without any modifications.

Sorted distribution estimation

Let c be the absolute constant defined just prior to Theorem 1. For any distribution $p \in \Delta_{\mathcal{X}}$, error parameter $\varepsilon \in (0, 1)$, and sampling parameter n , draw a sample $X^n \sim p$ and denote its profile by φ .

Theorem 5. *If $n = \Omega(n(\varepsilon)) = \Omega(k/(\varepsilon^2 \log k))$ and $\varepsilon \geq n^{-c}$,*

$$\Pr(\ell_1^{\leq}(p_{\varphi}, p) \geq \varepsilon) \leq \exp(-\Omega(\sqrt{n})).$$

For a *different* $c > 0$, Theorem 5 also holds for APMML with a better probability bound $\exp(-n^{2/3})$.

Identity testing

The recent works of Diakonikolas and Kane [30] and Goldreich [35] provided a procedure reducing identity testing to uniformity testing, while modifying the desired accuracy and alphabet size by only absolute constant factors. Hence below we consider uniformity testing.

The uniformity tester T_{PML} shown in Figure 1 is purely based on PML and satisfies

Theorem 6. *If $\varepsilon = \tilde{\Omega}(k^{-1/4})$ and $n = \tilde{\Omega}(\sqrt{k}/\varepsilon^2)$, then the tester $T_{\text{PML}}(X^n)$ will be correct with probability at least $1 - k^{-2}$. The tester also distinguishes between $p = p_u$ and $\|p - p_u\|_2 \geq \varepsilon/\sqrt{k}$.*

The $\tilde{\Omega}(\cdot)$ notation only hides logarithmic factors of k . The tester T_{PML} is near-optimal as for uniform distribution p_u , the results in [32] yield an $\Omega(\sqrt{k} \log k / \varepsilon^2)$ lower bound on $n_{p_u}(\varepsilon, k^{-2})$.

The rest of the paper is organized as follows. Section 2.2 presents several immediate implications. Section 3 and Section 4 illustrate PML’s theoretical and practical advantages by comparing it to existing methods for a variety of learning tasks. Section 5 to 8 present the proofs of Theorem 1 to 6.

Input: parameters k, ε , and a sample $X^n \sim p$ with profile φ .
if $\max_x \mu_x(X^n) \geq 3 \max\{1, n/k\} \log k$ **then return 1;**
elif $\|p_{\varphi} - p_u\|_2 \geq 3\varepsilon/(4\sqrt{k})$ **then return 1;**
else return 0.

Figure 1: Uniformity tester T_{PML}

2.2 Implications

Several immediate implications are in order.

We say that a plug-in estimator is *universally sample-optimal* for estimating symmetric properties if there exist absolute positive constants c_1, c_2 and c_3 , such that for any 1-Lipschitz property on $(\Delta_{\mathcal{X}}, R)$, with probability $\geq 9/10$, the plug-in estimator uses just c_1 times the sample size n required by the minimax estimator to achieve c_2 times its error, whenever this error is at least n^{-c_3} .

Note that the “1-Lipschitz property” class can be replaced by other general property classes, but not by those containing only a few specific properties, since “universal” means “applicable to all cases”.

Theorem 1 makes PML the *first* plug-in estimator that is *universally sample-optimal* for a broad class of distribution properties. In particular, Theorem 1 also covers the four properties considered in [5]. To see this, as mentioned in the introduction, \tilde{C}_m, H , and D are 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$; as for \tilde{S} , the following result [5] relates it to \tilde{C}_m for distributions in $\Delta_{\geq 1/k}$, and proves PML’s optimality.

Lemma 1. *For any $\varepsilon > 0$, $m = k \log(1/\varepsilon)$, and $p \in \Delta_{\geq 1/k}$,*

$$|\tilde{S}(p) - \tilde{C}_m(p) \log(1/\varepsilon)| \leq \varepsilon.$$

The theorem also applies to many other properties. As an example [80], given an integer $s > 0$, let $f_s(x) := \min\{x, |x - 1/s|\}$. Then to within a factor of two, $f_s(p) := \sum_x f_s(p_x)$ approximates the ℓ_1 distance between any distribution p and the closest uniform distribution in $\Delta_{\mathcal{X}}$ of support size s .

In Section 3.2 we compare Theorem 1 with existing results and present more of its implications.

Theorem 2 and 3 imply that for all non-integer $\alpha > 3/4$ (resp. $\alpha > 5/6$), the PML (resp. APML) plug-in estimator achieves a sample complexity better than the best currently known [6]. This makes both the PML and APML plug-in estimators the state-of-the-art algorithms for estimating non-integer order Rényi entropy. See Section 3.3 for an introduction of known results, and see Section 3.4 for a detailed comparison between existing methods and ours.

Theorem 4 shows that for all integer $\alpha > 1$, the sample complexity of the PML plug-in estimator has optimal $k^{1-1/\alpha}$ dependence [6, 63] on the alphabet size.

Theorem 5 makes APML the first distribution estimator under sorted ℓ_1 distance that is both near-linear-time computable and sample-optimal for a range of desired accuracy ε beyond inverse polylogarithmic of n . In comparison, existing algorithms [2, 43, 82] either run in polynomial time in the sample sizes, or are only known to achieve optimal sample complexity for $\varepsilon = \Omega(1/\sqrt{\log n})$, which is essentially different from the applicable range of $\varepsilon \geq n^{-\Theta(1)}$ in Theorem 5. We provide a more detailed comparison in Section 3.6.

Theorem 6 provides the first PML-based uniformity tester with near-optimal sample complexity. As stated, the tester also distinguishes between $p = p_u$ and $\|p - p_u\|_2 \geq \varepsilon/\sqrt{k}$. This is a stronger guarantee since by the Cauchy-Schwarz inequality, $\|p - p_u\|_1 \geq \varepsilon$ implies $\|p - p_u\|_2 \geq \varepsilon/\sqrt{k}$.

Note that several other uniformity testers in the literature (see Section 3.7) also provide the same ℓ_2 testing guarantee, since all of them are essentially counting sample collisions, i.e., the number of location pairs such that the sample points at those locations are equal.

3 Related work and comparisons

3.1 Additive property estimation

The study of additive property estimation dates back at least half a century [17, 39, 40] and has steadily grown over the years. For any additive symmetric property f and sequence x^n , the simplest and most widely-used approach uses the *empirical (plug-in)* estimator $\hat{f}^E(x^n) := f(p_\mu(x^n))$ that evaluates f at the empirical distribution. While the empirical estimator performs well in the large-sample regime, modern data science applications often concern high-dimensional data, for which more involved methods have yielded property estimators that are more sample-efficient. For example, for relatively large k and for f being \tilde{S} , \tilde{C}_m , H , or D , recent research [51, 68, 79, 80, 87, 88] showed that the empirical estimator is optimal up to logarithmic factors, namely $n_f(\mathcal{P}, \varepsilon) = \Theta_\varepsilon(n_f(\hat{f}^E, \mathcal{P}, \varepsilon)/\log n_f(\hat{f}^E, \mathcal{P}, \varepsilon))$, where \mathcal{P} is $\Delta_{\geq 1/k}$ for \tilde{S} , and is $\Delta_{\mathcal{X}}$ for the other properties.

Below we classify the methods for deriving the corresponding sample-optimal estimators into two categories: plug-in and approximation, and provide a high-level description. For simplicity of illustration, we assume that $\varepsilon \in (0, 1]$.

The *plug-in* approach essentially estimates the unknown distribution multiset, which suffices for computing any symmetric properties. Besides the empirical and PML estimators, Efron and Thisted [33] proposed a linear-programming approach that finds a multiset estimate consistent with the sample's profile. This approach was then adapted and analyzed by Valiant and Valiant [79, 82], yielding plug-in estimators that achieve near-optimal sample complexities for H and \tilde{S} , and optimal sample complexity for D , when ε is relatively large.

The *approximation* approach modifies non-smooth segments of the probability function to correct the bias of empirical estimators. A popular modification is to replace those non-smooth segments by their low-degree polynomial approximations and then estimate the modified function. For several properties including the above four and *power sum* $P_\alpha(p) := \sum_x p_x^\alpha$, where α is a given parameter, this approach yields property-dependent estimators [51, 68, 87, 88] that are sample-optimal for all ε .

More recently, Acharya et al. [5] proved the aforementioned results on PML estimator and made it the first unified, sample-optimal plug-in estimator for \tilde{S} , \tilde{C}_m , H and D and relatively large ε . Following these advances, Han et al. [43] refined the linear-programming approach and designed a plug-in estimator that implicitly performs polynomial approximation and is sample-optimal for H , \tilde{S} , and P_α with $\alpha < 1$, when ε is relatively large.

3.2 Comparison I: Theorem 1 and related property-estimation work

In terms of the estimator’s theoretical guarantee, Theorem 1 is essentially the same as Valiant and Valiant [80]. However, for each property, k , and n , [80] solves a different linear program and constructs a new estimator, which takes polynomial time. On the other hand, both the PML estimator and its near-linear-time computable variant, once computed, can be used to accurately estimate exponentially many properties that are 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$. A similar comparison holds between the PML method and the approximation approach, while the latter is provably sample-optimal for only a few properties. In addition, Theorem 1 shows that the PML estimator often achieves the optimal sample complexity up to a small constant factor, which is a desired estimator attribute shared by some, but not all approximation-based estimators [51, 68, 87, 88].

In term of the method and proof technique, Theorem 1 is closest to Acharya et al. [5]. On the other hand, [5] establishes the optimality of PML for only four properties, while our result covers a much broader property class. In addition, both the above mentioned “small constant factor” attribute, and the confidence boost from $2/3$ to $1 - \exp(-4\sqrt{n})$ are unique contributions of this work. The PML plug-in approach is also close in flavor to the plug-in estimators in Valiant and Valiant [79, 82] and their refinement in Han et al. [43]. On the other hand, as pointed out previously, these plug-in estimators are provably sample-optimal for only a few properties. More specifically, for estimating H , \tilde{S} , and \tilde{C}_m , the plug-in estimators in [79, 82] achieve sub-optimal sample complexities with regard to the desired accuracy ε ; and the estimation guarantee in [43] is provided in terms of the approximation errors of $\tilde{\mathcal{O}}(\sqrt{n})$ polynomials that are not directly related to the optimal sample complexities.

3.3 Rényi entropy estimation

Motivated by the wide applications of Rényi entropy, heuristic estimators were proposed and studied in the physics literature following [41], and asymptotically consistent estimators were presented and analyzed in the statistical learning literature [52, 90]. For the special case of 1-Rényi (or Shannon) entropy, the works of [79, 80] determined the sample complexity to be $n_f(\varepsilon) = \Theta(k/(\varepsilon \log k))$.

For general α -Rényi entropy, the best-known results in Acharya et al. [6] state that for integer and non-integer α values, the corresponding sample complexities $n_f(\varepsilon, \delta)$ are $\mathcal{O}_\alpha(k^{1-1/\alpha} \log(1/\delta)/\varepsilon^2)$ and $\mathcal{O}_\alpha(k^{\min\{1/\alpha, 1\}} \log(1/\delta)/(\varepsilon^{1/\alpha} \log k))$, respectively. The upper bounds for integer α are achieved by an estimator that corrects the bias of the empirical plug-in estimator. To achieve the upper bounds for non-integer α values, one needs to compute some best polynomial approximation of z^α , whose degree and domain both depend on n , and construct a more involved estimator using the approximation approach [51, 87] mentioned in Section 3.1.

3.4 Comparison II: Theorem 2 to 4 and related Rényi-entropy-estimation work

Our result shows that a single PML estimate suffices to estimate the Rényi entropy of different orders α . Such adaptiveness to the order parameter is a significant advantage of PML over existing methods. For example, by Theorem 3 and the union bound, one can use a *single* APML or PML to accurately approximate exponentially many non-integer order Rényi entropy values, yet still maintains an overall confidence of $1 - \exp(-k^{0.9})$. By comparison, the estimation heuristic in [6] requires different polynomial-based estimators for different α values. In particular, to construct each estimator, one needs to compute some best polynomial approximation of z^α , which is not known to admit a closed-form formula for $\alpha \notin \mathbb{Z}$. Furthermore, even for a single α and with a sample size \sqrt{k} times larger, such estimator is not known to achieve the same level of confidence as PML or APML.

As for the theoretical guarantees, the sample-complexity upper bounds in both Theorem 2 and 3 are better than those mentioned in the previous section. More specifically, for any $\alpha \in (3/4, 1)$ and $\delta \geq \exp(-k^{-0.5})$, Theorem 2 shows that $n_f(\varepsilon, \delta) = \mathcal{O}_\alpha(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$. Analogously, for any non-integer $\alpha > 1$ and $\delta \geq \exp(-k^{-0.9})$, Theorem 3 shows that $n_f(\varepsilon, \delta) = \mathcal{O}_\alpha(k/(\varepsilon^{1/\alpha} \log k))$. Both bounds are better than the best currently known by a $\log(1/\delta)$ factor.

3.5 (Sorted) distribution estimation

Estimating large-alphabet distributions from their samples is a fundamental statistical learning tenet. Over the past few decades, distribution estimation has found numerous applications, ranging from natural language modeling [24] to biological research [9], and has been studied extensively. Under the classical ℓ_1 and KL losses, existing research [14, 53] showed that the corresponding sample complexities $n(\varepsilon)$ are $\Theta(k/\varepsilon^2)$ and $\Theta(k/\varepsilon)$, respectively. Several recent works have investigated the analogous formulation under sorted ℓ_1 distance, and revealed a lower sample complexity of $n(\varepsilon) = \Theta(k/(\varepsilon^2 \log k))$. Specifically, under certain conditions, Valiant and Valiant [82], Han et al. [43] derived sample-optimal estimators using linear programming, and Acharya et al. [2], Das. [29] showed that PML achieves a sub-optimal $\mathcal{O}(k/(\varepsilon^{2.1} \log k))$ sample complexity for relatively large ε .

3.6 Comparison III: Theorem 5 and related distribution-estimation work

We compare our results with existing ones from three different perspectives.

Applicable parameter ranges: As shown by [43], for $\varepsilon \ll n^{-1/3}$, the simple empirical estimator is already sample-optimal. Hence we consider the parameter range $\varepsilon = \Omega(n^{-1/3})$. For the results in [2, 29] and [82] to hold, we would need ε to be at least $\Omega(1/\sqrt{\log n})$. On the other hand, Theorem 5 shows that PML and APML are sample-optimal for ε larger than $n^{-\Theta(1)}$. Here, the gap is exponentially large. The result in [43] applies to the whole range $\varepsilon = \Omega(n^{-1/3})$, which is larger than the applicable range of our results.

Time complexity: Both the APML and the estimator in [82] are near-linear-time computable in the sample sizes, while the estimator in [43] would require polynomial time to be computed.

Statistical confidence: The PML and APML achieve the desired accuracy with an error probability at most $\exp(-\Omega(\sqrt{n}))$. On the contrary, the estimator in [43] is known to achieve an error probability that decreases only as $\mathcal{O}(n^{-3})$. The gap is again exponentially large. The estimator in [82] admits a better error probability bound of $\exp(-n^{0.02})$, which is still far from ours.

3.7 Identity testing

Initiated by the work of [37], identity testing is arguably one of the most important and widely-studied problem in distribution property testing. Over the past two decades, a sequence of works [4, 12, 30–32, 37, 71, 78] have addressed the sample complexity of this problem and proposed testers with a variety of guarantees. In particular, applying a coincidence-based tester, Paninski [71] determined the sample complexity of uniformity testing up to constant factors; utilizing a variant of the Pearson chi-squared statistic, Valiant and Valiant [78] resolved the general identity testing problem. For an overview of related results, we refer interested readers to [16] and [36]. The contribution of this work is mainly showing that PML, is a unified sample-optimal approach for several related problems, and as shown in Theorem 6, also provides a near-optimal tester for this important testing problem.

4 Numerical experiments

A number of different approaches have been taken to computing the PML and its approximations. Among the existing works, Acharya et al. [1] considered exact algebraic computation, Orlitsky et al. [65, 66] designed an EM algorithm with MCMC acceleration, Vontobel [85, 86] proposed a Bethe approximation heuristic, Anevski et al. [8] introduced a sieved PML estimator and a stochastic approximation of the associated EM algorithm, and Pavlichin et al. [72] derived a dynamic programming approach. Notably and recently, for a sample size n , Charikar et al. [23] constructed an explicit $\exp(-\mathcal{O}(n^{2/3} \log^3 n))$ -approximate PML whose computation time is near-linear in n .

In this section, we first introduce a variant of the MCMC-EM algorithm in [65, 66, 69] and then demonstrate the efficacy of PML on a variety of learning tasks through experiments.

4.1 MCMC-EM algorithm variant

To approximate PML, the work [65] proposed an MCMC-EM algorithm, where MCMC and EM stand for Markov chain Monte Carlo and expectation maximization, respectively. A sketch of the original MCMC-EM algorithm can be found in [65], and a detailed description is available in Chapter 6 of [69]. The EM part uses a simple iteration procedure to update the distribution estimates. One can show [69] that it is equivalent to the conventional *generalized gradient ascent method*. The MCMC part exploits local properties of the update process and accelerates the EM computation. Below we present a variant of this algorithm that often runs faster and is more accurate.

Step 1: We separate the large and small multiplicities. Define a threshold parameter $\tau := 1.5 \log^2 n$ and suppress X^n in $p_\mu(X^n)$ for simplicity. For symbols x with $\mu_x(X^n) \geq \tau$, estimate their probabilities by $p_\mu(x) = \mu_x(X^n)/n$ and remove them from the sample. Denote the collection of removed symbols by R and the remaining sample sequence by X^r . In the subsequent steps, we apply the EM-MCMC algorithm to X^r .

The idea is simple: By the Chernoff-type bound for binomial random variables, with high probability, the empirical frequency $\mu_x(X^n)/n$ of a large-multiplicity symbol x is very close to its mean value $p(x)$. Hence for large-multiplicity symbols we can simply use the empirical estimates and focus on estimating the probabilities of small-multiplicity symbols. This is similar to initializing the EM algorithm by the empirical distribution and fixing the large probability estimates through the iterations. However, the approach described here is more efficient.

Step 2: We determine a proper alphabet size for the output distribution of the EM algorithm. If the true value k is provided, then we simply use $k - |R|$. Otherwise, we apply the following support size estimator [5] to X^r :

$$\hat{S}(X^r) := \sum_{j \geq 1} (1 - (-(t-1))^j \Pr(L \geq j)) \cdot \varphi_j(X^r),$$

where $t = \log r$ and L is an independent binomial random variable with support size $\lceil \frac{1}{2} \log_2(\frac{rt^2}{t-1}) \rceil$ and success probability $(t+1)^{-1}$. For any ε larger than an absolute constant, estimator \hat{S} achieves the optimal sample complexity $n_f(\Delta_{\geq 1/k}, \varepsilon)$ in estimating support size, up to constant factors [5].

Step 3: Apply the MCMC-EM algorithm in [65, 69] to $\varphi(X^r)$ with the output alphabet size determined in the previous step, and denote the resulting distribution estimate by p_r . (In the experiments, we perform the EM iteration for 30 times.) Intuitively, this estimate corresponds to the conditional distribution given that the next observation is a symbol with small probability.

Step 4: Let $T_\mu := \sum_{x \in R} p_\mu(x)$ be the total probability of the large-multiplicity symbols. Treat p_r as a vector and let $p'_r := (1 - T_\mu) \cdot p_r$. For every symbol $x \in R$, append $p_\mu(x)$ to p'_r , and return the resulting vector. Note that this vector corresponds to a valid discrete distribution.

Algorithm code

The implementation of our algorithm is also available at <https://github.com/ucsdyi/PML>.

For computational efficiency, the program code for the original MCMC-EM algorithm in [65, 69] is written in C++, with a file name “MCMCEM.cpp”. The program code for other functions is written in Python3. Note that to execute the program, one should have a 64-bit Windows/Linux system with Python3 installed (64-bit version). In addition, we also use functions provided by “NumPy” and “SciPy”, while the latter is not crucial and can be removed by modifying the code slightly.

Our implementation also makes use of “ctypes”, a *built-in* foreign language library for Python that allows us to call C++ functions directly. Note that before calling C++ functions in Python, we need to compile the corresponding C++ source files into DLLs or shared libraries. We have compiled and included two such files, one is “MCMCEM.so”, the other is “MCMCEM.dll”.

Functions in “MCMCEM.cpp” can be used separately. To compute a PML estimate, simply call the function “int PML(int MAXSZ=10000, int maximum_EM=20, int EM_n=100)”, where the first parameter specifies an upper bound on the support size of the output distribution, the second provides the maximum number of EM iteration, and the last corresponds to the sample size n . This function takes as input a local file called “proFile”, which contains the profile vector $\varphi(X^n)$ in the format of “1 4 7 10 . . .”. Specifically, the file “proFile” consists of only space-separated non-negative integers, and the i -th integer represents the value of $\varphi_i(X^n)$. The output is a vector of length at most MAXSZ, and is stored in another local file called “PMLFile”. Each line of the file “PMLFile” contains a non-negative real number, corresponding to a probability estimate.

To perform experiments and save the plots to the directory containing the code, simply execute the file “Main.py”. To avoid further complication, the code compares our estimator with only three other estimators: empirical, empirical with a larger $n \log n$ sample size, and improved Good-Turing [64] (for distribution estimation under unsorted ℓ_1 distance). The implementation covers all the distributions described in the next section. One can test any of these distributions by including it in “D_List” of the “main()” function. The implementation also covers a variety of learning tasks, such as distribution estimation under sorted and unsorted ℓ_1 distances, and property estimation for Shannon entropy, α -Rényi entropy, support coverage, and support size.

Finally, functions related to distribution and sample generation are available in file “Samples.py”. Others including the property computation functions, the sorted and unsorted ℓ_1 distance functions, and the previously-described support size estimator, are contained in file “Functions.py”.

4.2 Experiment distributions

In the following experiments, samples are generated according to six distributions with the same support size $k = 5,000$.

Three of them have finite support by definition: uniform distribution, two-step distribution with half the symbols having probability $2/(5k)$ and the other half have probability $8/(5k)$, and a three-step distribution with one third the symbols having probability $3/(13k)$, another third having probability $9/(13k)$, and the remaining having probability $27/(13k)$.

The other three distributions are over $\{i \in \mathbb{Z} : i \geq 1\}$, and are truncated at $i = 5,000$ and re-normalized: geometric distribution with parameter $g = 1/k$ satisfying $p_i \propto (1 - g)^i$, Zipf distribution with parameter $1/2$ satisfying $p_i \propto i^{-1/2}$, and log-series distribution with parameter $\gamma = 2/k$ satisfying $p_i \propto (1 - \gamma)^i/i$.

4.3 Experiment results and details

As shown below, the proposed PML approximation algorithm has exceptional performance.

Distribution estimation under ℓ_1 distance

We derive a new distribution estimator under the (unsorted) ℓ_1 distance by combining the proposed PML computation algorithm with the denoising procedure in [81] and a missing mass estimator [64].

First we describe this distribution estimator, which takes a sample X^n from some unknown distribution p . An optional input is \mathcal{X} , the underlying alphabet.

Step 1: Apply the PML computation algorithm described in Section 4.1 to X^n , and denote the returned vector, consisting of non-negative real numbers that sum to 1, by V .

Step 2: Employ the following variant of the denoising procedure in [81]. Arbitrarily remove a total probability mass of $\log^{-2} n$ from entries of the vector V without making any entry negative. Then for each $j \leq \log^2 n$, augment the vector by $n/(j \log^4 n)$ entries of probability j/n . For every multiplicity $\mu \geq 1$ appearing in the sample, assign to all symbols appearing μ times the following probability value. If $\mu \geq \log^2 n$, simply assign to each of these symbols the empirical estimate μ/n ; otherwise, temporally associate a weight of $\text{bin}(n, v, \mu) := \binom{n}{\mu} (1-v)^{n-\mu} v^\mu$ with each entry v in V , and assign to each of these symbols the current weighted median of V .

Step 3: If \mathcal{X} is available, we can estimate the total probability mass $M(X^n) := \sum_{x \in \mathcal{X}} \mathbb{1}_{x \notin X^n}$ of the unseen symbols (a.k.a., the *missing mass*) by the following estimator:

$$\hat{M}(X^n) := \frac{\varphi_1(X^n)}{\sum_j (j\varphi_j(X^n)\mathbb{1}_{j > \varphi_{j+1}} + (j+1)\varphi_{j+1}(X^n)\mathbb{1}_{j \leq \varphi_{j+1}})}.$$

We equally distribute this probability mass estimate among symbols that do not appear in the sample.

As shown below, the proposed distribution estimator achieves the state-of-the-art performance.

In Figures 2, the horizontal axis reflects the sample size n , ranging from 10,000 to 100,000, and the vertical axis reflects the (unsorted) ℓ_1 distance between the true distribution and the estimates, averaged over 30 independent trials. We compare our estimator with three others: the improved Good-Turing estimator [64, 46], the empirical estimator, serving as a baseline, and the empirical estimator with a larger $n \log n$ sample size. Note that $\log n$ is roughly 11. As shown in [64], the improved Good-Turing estimator is provably instance-by-instance near-optimal and substantially outperforms other estimators such as the Laplace (add-1) estimator, the Braess-Sauer estimator [14], and the Krichevsky-Trofimov estimator [55]. Hence we do not include those estimators below.

As the following plots show, our proposed estimator outperformed the improved Good-Turing estimator in all experiments.

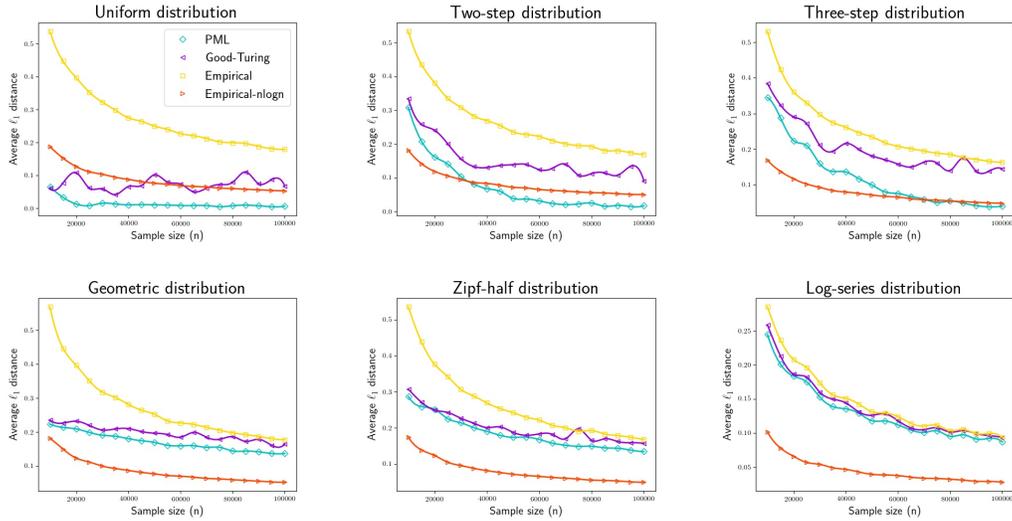


Figure 2: Distribution estimation under ℓ_1 distance

Distribution estimation under sorted ℓ_1 distance

In Figure 3, the sample size n ranges from 2,000 to 20,000, and the vertical axis reflects the sorted ℓ_1 distance between the true distribution and the estimates, averaged over 30 independent trials. We compare our estimator with that proposed by Valiant and Valiant [82] that utilizes linear programming, with the empirical estimator, and with the empirical estimator with a larger $n \log n$ sample size.

We do not include the estimator in [43] since there is no implementation available, and as pointed out by the recent work of [84] (page 7), the approach in [43] “is quite unwieldy. It involves significant parameter tuning and special treatment for the edge cases.” and “Some techniques . . . are quite crude and likely lose large constant factors both in theory and in practice.”

As shown in Figure 3, with the exception of uniform distribution, where the estimator in Valiant and Valiant [82] (VV-LP) is the best and PML is the closest second, the PML estimator outperforms VV-LP for all other tested distributions. As the underlying distribution becomes more skewed, the improvement of PML over VV-LP grows. For the log-series distribution, the performance of VV-LP is even worse than the empirical estimator.

Additionally, the plots also demonstrate that PML has a more stable performance than VV-LP.

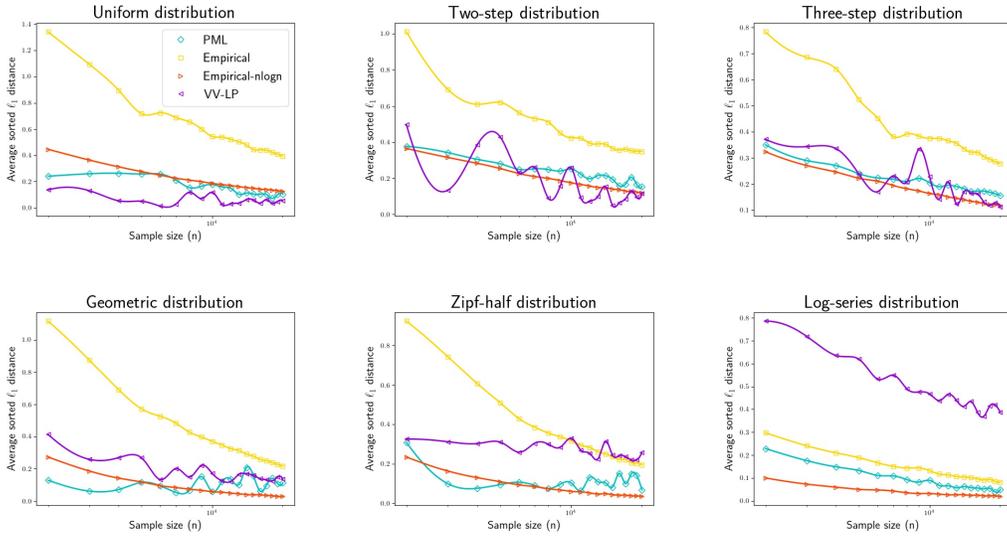


Figure 3: Distribution estimation under sorted ℓ_1 distance

Shannon entropy estimation under absolute error

In Figure 4, the sample size n ranges from 1,000 to 1,000,000, and the vertical axis reflects the absolute difference between the true entropy values and the estimates, averaged over 30 independent trials. We compare our estimator with two state-of-the-art estimators, WY in [87], and JVHW in [51], as well as the empirical estimator, and the empirical estimator with a larger $n \log n$ sample size. Additional entropy estimators such as the Miller-Mallow estimator [18], the best upper bound (BUB) estimator [70], and the Valiant-Valiant estimator [82] were compared in [87, 51] and found to perform similarly to or worse than the two estimators that we compared with, therefore we do not include them here. Also, considering [82], page 50 in [91] notes that “the performance of linear programming estimator starts to deteriorate when the sample size is very large.”

Note that the alphabet size k is a crucial input to WY, but is not required by either JVHW or our PML algorithm. In the experiments, we provide WY with the true value of $k = 5,000$.

As shown in the plots, our estimator performs as well as these state-of-the-art estimators.

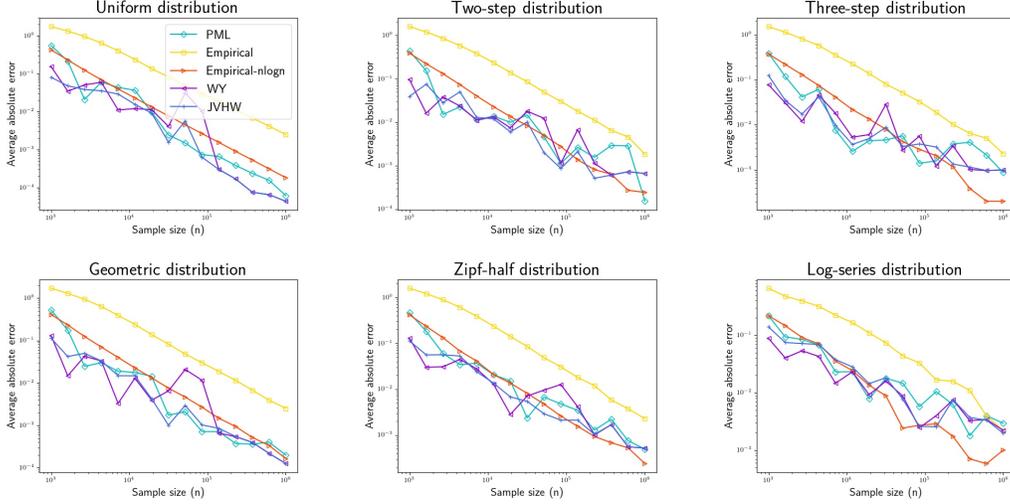


Figure 4: Shannon entropy estimation under absolute error

α -Rényi entropy estimation under absolute error

For a distribution $p \in \Delta_{\mathcal{X}}$, recall that the α -power sum of p is $P_{\alpha}(p) = \sum_x p(x)^{\alpha}$, implying $H_{\alpha}(p) = (1 - \alpha)^{-1} \log(P_{\alpha}(p))$. To establish the sample-complexity upper bounds mentioned in Section 3.3 for non-integer α values, Acharya et al. [6] first estimate the $P_{\alpha}(p)$ using the α -power-sum estimator proposed in [51], and then substitute the estimate into the previous equation. The authors of [51] have implemented this two-step Rényi entropy estimation algorithm. In the experiments, we take a sample of size n , ranging from 10,000 to 100,000, and compare our estimator with this implementation, referred to as *JVHW*, the empirical estimator, and the empirical estimator with a larger $n \log n$ sample size. Note that $\log n$ ranges from 9.2 to 11.5. According to the results in [6], the sample complexities for estimating α -Rényi entropy are quite different for $\alpha < 1$ and $\alpha > 1$, hence we consider two cases: $\alpha = 0.5$ and $\alpha = 1.5$.

As shown in Figure 5 and 6, our estimator clearly outperformed the one proposed by [6, 51].

We further note that for small sample sizes and several distributions, the estimator in [6, 51] performs significantly worse than ours. Also, for large sample sizes, the estimators in [6, 51] degenerates to the simple empirical plug-in estimator. In comparison, our proposed estimator tracks the performance of the empirical estimator with a larger $n \log n$ sample size for nearly all the tested distributions.

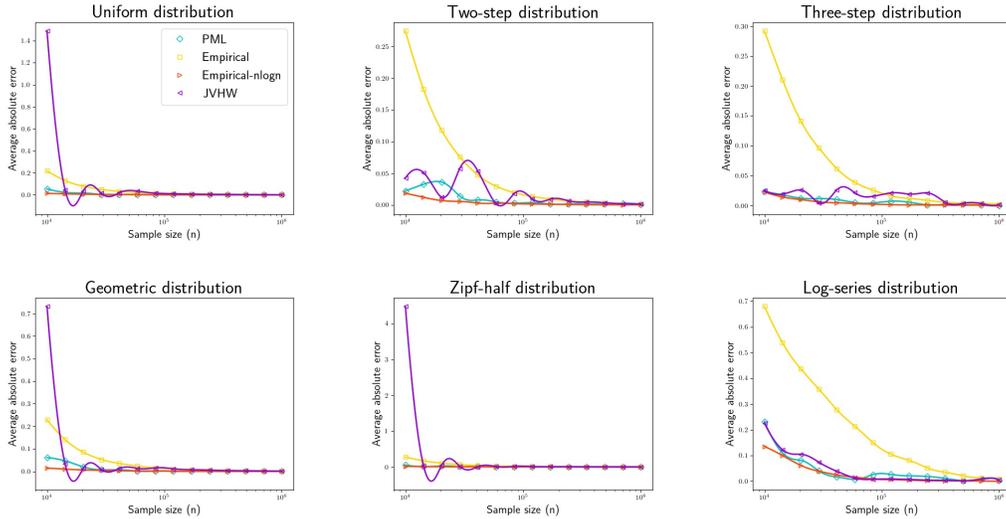


Figure 5: 0.5-Rényi entropy estimation under absolute error

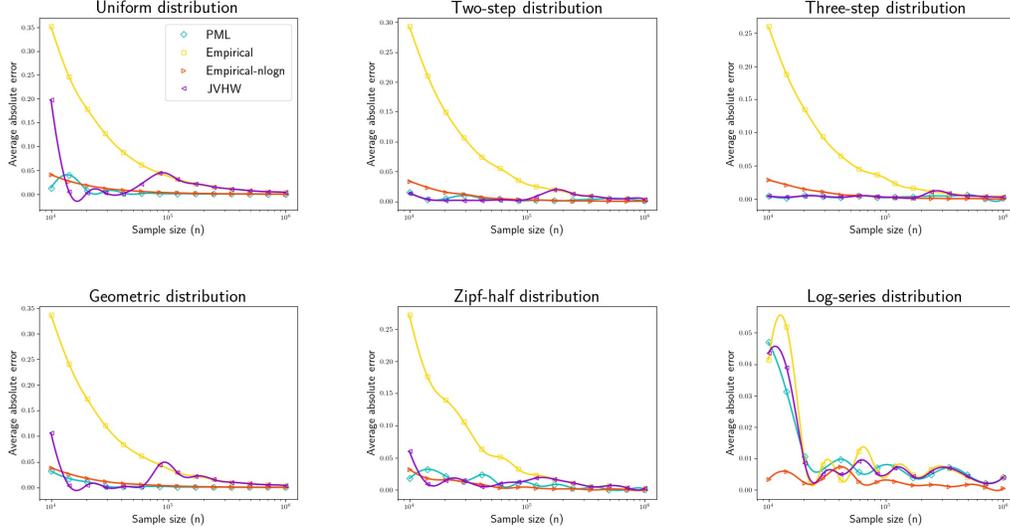


Figure 6: 1.5-Rényi entropy estimation under absolute error

5 Lipschitz-property estimation

5.1 Proof outline of Theorem 1

The proof proceeds as follows. First, fixing n , \mathcal{X} , and a symmetric additive property f that is 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$, we consider a related linear program defined in [83], and lower bound the worst-case error of any estimators using the linear program's objective value, say v . Second, following the construction in [83], we find an explicit estimator \hat{f}^* that is *linear*, i.e., can be expressed as a linear combination of φ_i 's, and show optimality by upper bounding its worst-case error in terms of v . Third, we study the concentration of a general linear estimator, and through the McDiarmid's inequality [60], relate the tail probability of its estimate to the estimator's sensitivity to the input changes. Fourth, we bound the sensitivity of \hat{f}^* by the maximum difference between its consecutive coefficients, and further bound this difference by a function of n , showing that the estimate induced by \hat{f}^* highly concentrates around its expectation. Finally, we invoke the result in [5] that the PML-plug-in estimator is competitive to all profile-based estimators whose estimates are highly concentrated, concluding that PML shares the optimality of \hat{f}^* , thereby establishing Theorem 1.

5.2 Technical details

Let f be a symmetric additive property that is 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$. Without loss of generality, we assume that $f(p) = 0$ if $p(x) = 1$ for some $x \in \mathcal{X}$.

Lower bound First, fixing n , \mathcal{X} , and f , we lower bound the worst-case error of any estimators.

Let $u \in (0, 1/2)$ be a small absolute constant. If there is an estimator \hat{f} that, when given a length- n sample from any distribution $p \in \Delta_{\mathcal{X}}$, will estimate $f(p)$ up to an error of ε with probability at least $1/2 + u$. Then for any two distributions $p_1, p_2 \in \Delta_{\mathcal{X}}$ satisfying $|f(p_1) - f(p_2)| > \varepsilon$, we can use \hat{f} to distinguish $X^n \sim p_1$ from $X^n \sim p_2$, and will be correct with probability at least $1/2 + u$.

On the other hand, for any parameter $c_1 \in (1/100, 1/25]$ and $c_2 = 1/2 + 6c_1$, consider the corresponding linear program defined in Linear Program 6.7 in [83], and denote by v the objective value of any of its solutions. Then, Proposition 6.8 in [83] implies that we can find two distributions $p_1, p_2 \in \Delta_{\mathcal{X}}$ such that $|f(p_1) - f(p_2)| > v \cdot (1 - o(1)) - O(n^{-c_1} \log n)$, and no algorithm can use $\text{Poi}(n)$ sample points to distinguish these two distributions with probability at least $1/2 + u$.

The previous reasoning yields that $v < (1 + o(1))\varepsilon + O(n^{-c_1} \log n)$. By construction, v is a function of \mathcal{X} , n , and f , and essentially serves as a lower bound for ε .

Upper Bound Second, fixing n , \mathcal{X} , and f , we construct an explicit estimator based on the previously mentioned linear program, and show optimality by upper bounding its worst-case error in terms of v , the linear program's objective value.

A property estimator \hat{f} is *linear* if there exist real coefficients $\{\ell_i\}_{i \geq 1}$ such that the identity $\hat{f}(x^n) = \sum_{i \geq 1} \ell_i \cdot \varphi_i(x^n)$ holds for all x^n . The following lemma (Proposition 6.10 in [83]) bounds the worst-case error of a linear estimator when its coefficients satisfy certain conditions.

Lemma 2. *Given any positive integer m , and real coefficients $\{\beta_i\}_{i \geq 0}$, define $\varepsilon(y) := f(y)/y - e^{-my} \sum_{i \geq 0} \beta_i \cdot (my)^i / i!$. Let $\beta_i^* := \beta_{i-1} \cdot i/m, \forall i \geq 1$, and $\beta_0^* := 0$. If for some $a', b', c' > 0$,*

1. $|\varepsilon(y)| \leq a' + b'/y$,
2. $|\beta_j^* - \beta_\ell^*| \leq c' \sqrt{j/m}$ for any j and ℓ such that $|j - \ell| \leq \sqrt{j} \log m$,

then given a sample X^m from any $p \in D_{\mathcal{X}}$, the estimator defined by $\sum_{i \geq 1} \beta_i^ \cdot \varphi_i$ will estimate $f(p)$ with an accuracy of $a' + b' \cdot k + c' \cdot \log m$ and a failure probability at most $o(1/\text{poly}(m))$.*

Following the construction in [83] (page 124), let $z := (z_0, z_1, \dots)$ be the vector of coefficients induced by any solution of the dual program of the previously mentioned linear program. For our purpose, the way in which these coefficients are derived is largely irrelevant. One can show that $|z_\ell| \leq v \cdot n^{c_2}, \forall \ell \geq 0$. Let $t_n := 2n^{-c_1} \log n$ and $\alpha \in (0, 1)$, and define

$$\beta_i := (1 - e^{-t_n \alpha^i}) f\left(\frac{(i+1)\alpha}{n}\right) \frac{n}{(i+1)\alpha} + \sum_{\ell=0}^i z_\ell (1 - t_n)^\ell \alpha^\ell (1 - \alpha)^{i-\ell} \binom{i}{\ell}.$$

for any $i \leq n$, and $\beta_i := \beta_n$ for $i > n$. The next lemma shows that we can find proper parameters a, b , and c to apply Lemma 2 to the above construction. Specifically,

Lemma 3. *For any $\alpha \in [1/100, 1)$ and some $a'', b'' \geq 0$ such that $a'' + b''k \leq v$, if $v \leq \log^2 n$ and c_1, c_2 satisfy $\alpha c_2 + (3/2 - \alpha)c_1 \leq 1/4$, the two conditions in Lemma 2 hold for the above construction with $m = n/\alpha$, $a' = a'' + \mathcal{O}(n^{-c_1/2} \log^2 n)$, $b' = b''(1 + \mathcal{O}(t_n))$, and $c' = \mathcal{O}(n^{-1/4} \log^3 n)$. Furthermore, for any $i \geq 0$, we have $|\beta_i| \leq \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \log^3 n)$.*

This lemma differs from the results established in the proof of Proposition 6.19 in [83] only in the applicable range of α , where the latter assumes that $\alpha \in [1/2, 1)$. For completeness, we will present a proof of Lemma 3 in Appendix A.

By Lemma 2 and 3, if $v \leq \log^2 n$, given a sample $X^{n/\alpha}$ from any $p \in \Delta_{\mathcal{X}}$, the linear estimator $\sum_{i \geq 1} \beta_i^* \cdot \varphi_i$ will estimate $f(p)$ with an accuracy of $a' + b'k + c' \log(n/\alpha) = a'' + \mathcal{O}(n^{-c_1/2} \log^2 n) + b''k(1 + \mathcal{O}(t_n)) + \mathcal{O}(n^{-1/4} \log^4 n) \leq v(1 + \mathcal{O}(t_n)) + \mathcal{O}(n^{-c_1/2} \log^2 n)$ and a failure probability at most $o(1/\text{poly}(n))$. Recall that for fixed \mathcal{X} , n , and f , the value of v is a constant, thus can be computed without samples. Furthermore according to the last claim in Proposition 6.19 in [83], for $v > \log^2 n$, the estimator that always returns 0 has an error of at most $(1 + o(1))v$. Hence with high probability, the estimator $\hat{f}^* := \sum_{i \geq 1} (\beta_i^* \cdot \mathbb{1}_{v \leq \log^2 n}) \cdot \varphi_i$ will estimate $f(p)$ up to an error of $v(1 + o(1)) + \mathcal{O}(t_n \log n)$, for any possible values of v .

Concentration of linear estimators Third, we slightly diverge from the previous discussion and study the concentration of general linear estimators.

The *sensitivity* of a property estimator $\hat{f}: \mathcal{X}^* \rightarrow \mathbb{R}$ for a given input size n is

$$s_n(\hat{f}) := \max \{f(x^n) - f(y^n) : x^n \text{ and } y^n \text{ differ in one element}\},$$

the maximum change in its value when the input sequence is modified at exactly one location. For any $p \in \Delta_{\mathcal{X}}$ and $X^n \sim p$, the following corollary of the McDiarmid's inequality [60] relates the two-side tail probability of $\hat{f}(X^n)$ to $s_n(\hat{f})$.

Lemma 4. For all $t \geq 0$, we have $\Pr\left(|\hat{f}(X^n) - \mathbb{E}[\hat{f}(X^n)]| \geq t\right) \leq 2 \exp(-2t^2 \cdot (\sqrt{n}s_n(\hat{f}))^{-2})$.

Define $\ell_0 := 0$. The next lemma bounds the sensitivity of a linear estimator $\hat{f} := \sum_{i \geq 1} \ell_i \cdot \varphi_i$ in terms of $\max_{i \geq 1} |\ell_i - \ell_{i-1}|$, the maximum absolute difference between its consecutive coefficients.

Lemma 5. For any n and linear estimator $\hat{f} := \sum_{i \geq 1} \ell_i \cdot \varphi_i$, we have $s_n(\hat{f}) \leq 2 \max_{i \geq 1} |\ell_i - \ell_{i-1}|$.

Proof. Let x^n and y^n be two arbitrary sequences over \mathcal{X} that differ in one element. Let i be the index where $x_i \neq y_i$. Then by definition, the following multiplicity equalities hold: $\mu_{x_i}(x^n) = \mu_{x_i}(y^n) + 1$, $\mu_{y_i}(y^n) = \mu_{y_i}(x^n) + 1$, and $\mu_x(x^n) = \mu_x(y^n)$ for $x \in \mathcal{X}$ satisfying $x \neq x_i, y_i$. For simplicity of notation, let $\mu_0 := \mu_{x_i}(x^n)$, $\mu_1 := \mu_{y_i}(y^n)$, and for any $i \geq 1$, let $\hat{f}_i := \ell_{i-1} \cdot \varphi_{i-1} + \ell_i \cdot \varphi_i$.

The first multiplicity equality implies $\varphi_{\mu_0}(x^n) = \varphi_{\mu_0}(y^n) + 1$ and $\varphi_{\mu_0-1}(x^n) = \varphi_{\mu_0-1}(y^n) - 1$. Therefore, we have $\hat{f}_{\mu_0}(x^n) - \hat{f}_{\mu_0}(y^n) = \ell_{\mu_0} - \ell_{\mu_0-1}$. Similarly, the second equality implies $\hat{f}_{\mu_1}(x^n) - \hat{f}_{\mu_1}(y^n) = -\ell_{\mu_1} + \ell_{\mu_1-1}$. The third equality combines these two results and yields

$$\hat{f}(x^n) - \hat{f}(y^n) = \ell_{\mu_0} - \ell_{\mu_0-1} + (-\ell_{\mu_1} + \ell_{\mu_1-1}).$$

Applying the triangle inequality to the right-hand side completes the proof. \square

By these two lemmas, we have the following result for the concentration of linear estimators.

Corollary 1. For any $t \geq 0$, $p \in \Delta_{\mathcal{X}}$, and $\hat{f} := \sum_{i \geq 1} \ell_i \cdot \varphi_i$, if $X^n \sim p$, then

$$\Pr\left(|\hat{f}(X^n) - \mathbb{E}[\hat{f}(X^n)]| \geq t\right) \leq 2 \min_{i \geq 1} \exp(-t^2 \cdot (\sqrt{2n}(\ell_i - \ell_{i-1}))^{-2}).$$

Sensitivity bound Fourth, we bound the sensitivity of $\hat{f}^* = \sum_{i \geq 1} (\beta_i^* \cdot \mathbb{1}_{v \leq \log^2 n}) \cdot \varphi_i$. By Lemma 5, it suffices to consider the absolute difference between consecutive β_i^* 's. We assume $v \leq \log^2 n$ and $\alpha \in [1/100, 1)$, and analyze two cases below, depending on whether i is greater than $400n^{c_1}$ or not.

By Lemma 3, for $i \leq 400n^{c_1}$, we have $|\beta_i| \leq \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \log^3 n)$. Define $\beta_{-1} := 0$. Then,

$$|\beta_{i+1}^* - \beta_i^*| = \left| \frac{i+1}{n/\alpha} \beta_i - \frac{i}{n/\alpha} \beta_{i-1} \right| \leq \left| \frac{400n^{c_1} + 1}{n/\alpha} \beta_i \right| + \left| \frac{400n^{c_1}}{n/\alpha} \beta_{i-1} \right| \leq \mathcal{O}\left(n^{\alpha c_2 + (2-\alpha)c_1 - 1} \log^3 n\right).$$

For $i > 400n^{c_1}$, we only need to consider $i < n$ since $\beta_{i+1}^* = \beta_i^*$ for all $i \geq n$. Then,

$$\begin{aligned} |\beta_{i+1}^* - \beta_i^*| &\stackrel{(a)}{\leq} \left| \sum_{\ell=0}^i z_\ell (1-t_n)^\ell \alpha^\ell (1-\alpha)^{i-\ell} \binom{i}{\ell} \frac{(i+1)\alpha}{n} \right| + \left| \sum_{\ell=0}^{i-1} z_\ell (1-t_n)^\ell \alpha^\ell (1-\alpha)^{i-1-\ell} \binom{i-1}{\ell} \frac{i\alpha}{n} \right| \\ &\quad + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| + \left| e^{-t_n \alpha i} f\left(\frac{(i+1)\alpha}{n}\right) \right| + \left| e^{-t_n \alpha (i-1)} f\left(\frac{i\alpha}{n}\right) \right| \\ &\stackrel{(b)}{\leq} (n^{c_2} \log^2 n) \left(\left| \sum_{\ell=0}^i (1-t_n)^\ell \alpha^\ell (1-\alpha)^{i-\ell} \binom{i}{\ell} \right| + \left| \sum_{\ell=0}^{i-1} (1-t_n)^\ell \alpha^\ell (1-\alpha)^{i-1-\ell} \binom{i-1}{\ell} \right| \right) \\ &\quad + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| + \left| e^{-t_n \alpha i} f\left(\frac{(i+1)\alpha}{n}\right) \right| + \left| e^{-t_n \alpha (i-1)} f\left(\frac{i\alpha}{n}\right) \right| \\ &\stackrel{(c)}{\leq} (n^{c_2} \log^2 n) ((1-t_n \alpha)^i + (1-t_n \alpha)^{i-1}) + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| + 2e^{-t_n \alpha (i-1)}/e \\ &\stackrel{(d)}{\leq} 2(n^{c_2} \log^2 n) \left(1 - \frac{\log n}{50n^{c_1}}\right)^{400n^{c_1}} + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| + 2n^{-2}/e \\ &\stackrel{(e)}{=} 2(n^{c_2} \log^2 n) \left(\left(1 - \frac{\log n}{50n^{c_1}}\right)^{\frac{50n^{c_1}}{\log n}} \right)^{8 \log n} + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| + 2n^{-2}/e \\ &\stackrel{(f)}{\leq} 2n^{-2} + \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right|, \end{aligned}$$

where (a) follows from the triangle inequality; (b) follows from $i \leq n$, $v \leq \log^2 n$, and $|z_\ell| \leq v \cdot n^{c_2}$ for all $\ell \geq 0$; (c) follows from the binomial theorem and $|f(x)| \leq x |\log x| \leq 1/e$ for $x \in (0, 1]$; (d) follows from $\alpha \geq 1/100$, $i > 400n^{c_1}$, and $t_n = 2n^{-c_1} \log n$; (e) follows from simple algebra; and (f) follows from $c_2 = 1/2 + 6c_1 < 1$ and $(1 - 1/x)^x \leq e^{-1}$ for $x > 1$.

It remains to analyze the second term on the right-hand side.

$$\begin{aligned} \left| f\left(\frac{(i+1)\alpha}{n}\right) - f\left(\frac{i\alpha}{n}\right) \right| &\stackrel{(a)}{=} \frac{(i+1)\alpha}{n} \left| f\left(\frac{(i+1)\alpha}{n}\right) \frac{n}{(i+1)\alpha} - f\left(\frac{i\alpha}{n}\right) \frac{n}{(i+1)\alpha} \right| \\ &\stackrel{(b)}{=} \frac{(i+1)\alpha}{n} \left| f\left(\frac{(i+1)\alpha}{n}\right) \frac{n}{(i+1)\alpha} - f\left(\frac{i\alpha}{n}\right) \frac{n}{i\alpha} + f\left(\frac{i\alpha}{n}\right) \frac{n}{i(i+1)\alpha} \right| \\ &\stackrel{(c)}{\leq} \frac{(i+1)\alpha}{n} \left| \log \frac{i+1}{i} \right| + \frac{(i+1)\alpha}{n} \left| \frac{i\alpha}{n} \left(\log \left(\frac{i\alpha}{n} \right) \right) \frac{n}{i(i+1)\alpha} \right| \\ &\stackrel{(d)}{\leq} \frac{(i+1)\alpha}{n} \frac{1}{i} + \mathcal{O}\left(\frac{\log n}{n}\right) \stackrel{(e)}{\leq} \mathcal{O}\left(\frac{\log n}{n}\right), \end{aligned}$$

where (a), (b) and (e) follows from simple algebra; (c) follows from $|f(x)/x - f(y)/y| \leq |\log(x/y)|$ for all $x, y \in (0, 1]$; (d) follows from $\log(1+x) \leq x$ for $x \geq 0$ and $x |\log x| \leq 1/e$ for $x \in (0, 1]$.

Consolidating the above inequalities and applying Lemma 5, we get the sensitivity bound

$$s_n(f^*) \leq \mathcal{O}\left(n^{\alpha c_2 + (2-\alpha)c_1 - 1} \log^3 n\right).$$

Competitiveness of PML A property estimator \hat{f} is *profile-based* if there exists a mapping \hat{g} such that $\hat{f}(x^n) = \hat{g}(\varphi(x^n))$ for all $x^n \in \mathcal{X}^*$. The following lemma [2, 5, 29] states that the PML estimator is competitive to other profile-based estimators.

Lemma 6. *For any positive real numbers ε and δ , additive symmetric property f , and profile-based estimator \hat{f} , the PML-plug-in estimator $f(p_\varphi)$ satisfies*

$$n_f(f(p_\varphi), 2\varepsilon, \delta \cdot \exp(3\sqrt{n})) \leq n_f(\hat{f}, \varepsilon, \delta).$$

For any β -approximate PML, a similar result holds with $\delta \cdot \exp(3\sqrt{n})$ replaced by $\delta \cdot \exp(3\sqrt{n})/\beta$.

The factor $\exp(3\sqrt{n})$ directly comes from the well-known result of Hardy and Ramanujan [49] on integer partitions, since there is a bijective mapping from profiles of size n to partitions of integer n .

Final analysis Finally, we combine the above results and establish Theorem 1.

Denote by $\tau(n)$ the previous upper bound on $s_n(f^*)$. Let p be a distribution in $\Delta_{\mathcal{X}}$ and $X^n \sim p$. Let γ be an absolute constant in $(0, 1/4)$. Then by Lemma 4,

$$\Pr\left(|\hat{f}^*(X^n) - \mathbb{E}[\hat{f}^*(X^n)]| \geq 2n^{1-\gamma}\tau(n)\right) \leq 2\exp(-8n^{1-2\gamma}).$$

Let $\varepsilon > 0$ be an error parameter. Assume there exists an estimator \hat{f} that, when given a length- αn sample from any distribution $p' \in \Delta_{\mathcal{X}}$, estimates $f(p')$ up to an absolute error ε with probability at least $2/3$. Then according to the results in the upper- and lower-bound sections, with probability at most $o(1/\text{poly}(n))$, the estimate $\hat{f}^*(X^n)$ will differ from $f(p)$ by more than $v(1 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n) \leq \varepsilon(1 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n)$. In addition, by the equality $\sum_{i \geq 1} i \cdot \varphi_i(X^n) = n$ and Lemma 3, we surely have $|\hat{f}^*(X^n)| \leq |\sum_{i \geq 1} (i/m)\beta_{i-1} \cdot \varphi_i(X^n)| \leq \max_{i \geq 0} |\beta_i| \leq \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \log^3 n)$. Multiplying this bound by $o(1/\text{poly}(n))$ yields a quantity that is negligible comparing to $\mathcal{O}(n^{-c_1/2} \log^2 n)$. Therefore, the absolute bias $|\mathbb{E}[\hat{f}^*(X^n)] - f(p)|$ is at most $\varepsilon(1 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n)$. The triangle inequality combines this with the tail bound above:

$$\Pr\left(|\hat{f}(X^n) - f(p)| \geq \varepsilon(1 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n) + 2n^{1-\gamma}\tau(n)\right) \leq 2\exp(-8n^{1-2\gamma}).$$

Let $\alpha = 1/4$. For PML and APML estimators, set (γ, c_1) to be $(1/4, 1/31)$ and $(0.166, 1/91)$, respectively. Combined, the last inequality and Lemma 6 imply Theorem 1. There is a simple trade-off between α and c_1 induced by our proof technique. Specifically, if we increase the value of c_1 to achieve a better lower bound on ε , the value of α may need to be reduced accordingly, which enlarges the sample complexity gap between our estimators and the optimal one. For example, reducing α to $1/12$ and $1/22$, we can improve c_1 to $1/25$ and $1/20$, respectively, for both PML and APML.

6 α -Rényi entropy estimation

For any $p \in \Delta_{\mathcal{X}}$ and non-negative $\alpha \neq 1$, the α -Rényi entropy [74] of p is

$$H_{\alpha}(p) := \frac{1}{1-\alpha} \log P_{\alpha}(p) = \frac{1}{1-\alpha} \log \left(\sum_x p(x)^{\alpha} \right).$$

For \mathcal{X} of finite size k and any $p \in \Delta_{\mathcal{X}}$, it is well-known that $H_{\alpha}(p) \in [0, \log k]$.

6.1 Proof of Theorem 2: $\alpha \in (3/4, 1)$

For $\alpha \in (3/4, 1)$, the following theorem characterizes the performance of the PML-plug-in estimator.

For any distribution $p \in \Delta_{\mathcal{X}}$, error parameter $\varepsilon \in (0, 1)$, and sampling parameter n , draw a sample $X^n \sim p$ and denote its profile by φ . Then for sufficiently large k ,

Theorem 2. For an $\alpha \in (3/4, 1)$, if $n = \Omega_{\alpha}(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$,

$$\Pr(|H_{\alpha}(p_{\varphi}) - H_{\alpha}(p)| \geq \varepsilon) \leq \exp(-\sqrt{n}).$$

We establish both this theorem and an analogous result for APML in the remaining section. Let n be a sampling parameter and $p \in \Delta_{\mathcal{X}}$ be an unknown distribution. For some α -dependent positive constants $c_{\alpha,1}$ and $c_{\alpha,2}$ to be determined later, let $\tau := c_{\alpha,1} \log n$ and $d := c_{\alpha,2} \log n$ be threshold and degree parameters, respectively. Let N, N' be independent Poisson random variables with mean n . Consider Poisson sampling with two samples drawn from p , first of size N and the second N' . Suppressing the sample representations, for each $x \in \mathcal{X}$, we denote by μ_x and μ'_x the multiplicities of symbol x in the first and second samples, respectively. Denote by $q(z) := \sum_{m=0}^d a_m z^m$ be the degree- d min-max polynomial approximation of z^a over $[0, 1]$. We consider the following variant of the polynomial-based estimator proposed in [6].

$$\hat{P}_{\alpha} := \sum_x \left(\sum_{m=0}^d \frac{a_m (2\tau)^{\alpha-m} \mu_x^m}{n^{\alpha}} \right) \mathbb{1}_{\mu_x \leq 4\tau} \cdot \mathbb{1}_{\mu'_x \leq \tau} + \sum_x \left(\frac{\mu_x}{n} \right)^{\alpha} \mathbb{1}_{\mu'_x > \tau}.$$

The smaller the value of μ'_x is, the smaller we expect the value of $p(x)$ to be. In view of this, we denote the first and second components of \hat{P}_{α} by $\hat{P}_{\alpha}^{(s)}$ and $\hat{P}_{\alpha}^{(\ell)}$, and refer to them as small- and large-probability estimators, respectively. Note that our estimator differs from that in [6] only by the additional $\mathbb{1}_{\mu_x \leq 4\tau}$ term, which for sufficiently large $c_{\alpha,1}$, only modifies $\mathbb{E}[\hat{P}_{\alpha}^{(s)}]$ by at most $n^{-2\alpha}$.

Note that μ' naturally induces a partition over \mathcal{X} . For symbols x with $\mu_x \leq 4\tau$, we denote by

$$P_{\alpha, \mu'}^{(s)}(p) := \sum_{x: \mu_x \leq 4\tau} p(x)^{\alpha}$$

the small-probability power sum. Analogously, for symbols x with $\mu_x > 4\tau$, we denote by

$$P_{\alpha, \mu'}^{(\ell)}(p) := \sum_{x: \mu_x > 4\tau} p(x)^{\alpha}$$

the large-probability power sum. These are random properties with non-trivial variances and are hard to be analyzed. To address this, we apply an “expectation trick” and denote by $P_{\alpha}^{(s)}(p) := \mathbb{E}[P_{\alpha, \mu'}^{(s)}(p)]$ and $P_{\alpha}^{(\ell)}(p) := \mathbb{E}[P_{\alpha, \mu'}^{(\ell)}(p)]$ their expected values, both of which are additive symmetric properties.

Let ε be a given error parameter and $n = \Omega_{\alpha}(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$ be a sampling parameter. First we consider the small probability estimator. By the results in [6], for sufficiently large $c_{\alpha,1}$, the bias of $\hat{P}_{\alpha}^{(s)}$ in estimating $P_{\alpha}^{(s)}(p)$ satisfies

$$|\mathbb{E}[\hat{P}_{\alpha}^{(s)}] - P_{\alpha}^{(s)}(p)| \leq \mathcal{O}_{\alpha}(1) \cdot P_{\alpha}(p) \left(\frac{k}{n \log n} \right)^{\alpha} + n^{-\alpha} \leq \varepsilon P_{\alpha}(p),$$

where we have used $n^{-\alpha} = \mathcal{O}_{\alpha}(\varepsilon k^{-1} (\log k)^{\alpha}) \leq \varepsilon P_{\alpha}(p)$. To show concentration, we bound the sensitivity of estimator $\hat{P}_{\alpha}^{(s)}$. For $m \geq 0$, we can bound the coefficients of $q(x)$ as follows.

$$|a_m| = \mathcal{O}_{\alpha}((\sqrt{2} + 1)^d) = \mathcal{O}_{\alpha}(n^{c_{\alpha,2}}).$$

Therefore by definition, changing one point in the sample changes the value of $\hat{P}_\alpha^{(s)}$ by at most

$$2 \left(\sum_{m=0}^d \frac{|a_m|(2\tau)^{\alpha-m}(4\tau)^m}{n^\alpha} \right) \leq \sum_{m=0}^d \frac{|a_m|(2\tau)^\alpha 2^{m+1}}{n^\alpha} = \mathcal{O}_\alpha(n^{2c_{\alpha,2}-\alpha}(\log n)^\alpha).$$

Let $\lambda \in (0, 1/4)$ be an arbitrary absolute constant. For sufficiently small $c_{\alpha,2}$, the right-hand side is at most $\mathcal{O}_\alpha(n^{\lambda-\alpha})$. The McDiarmid's inequality together with the concentration of Poisson random variables implies that for all $\varepsilon \geq 0$,

$$\Pr\left(|\hat{P}_\alpha^{(s)} - \mathbb{E}[\hat{P}_\alpha^{(s)}]| \geq \varepsilon P_\alpha(p)\right) \leq 2 \exp(-\Omega_\alpha(\varepsilon^2 P_\alpha^2(p) n^{2\alpha-1-2\lambda})).$$

Note that $n = \Omega_\alpha(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$ and $P_\alpha(p) \geq 1$, which follows from the fact that z^α is a concave function over $[0, 1]$ for $\alpha \in (0, 1)$. Hence we obtain

$$\Pr\left(|\hat{P}_\alpha^{(s)} - \mathbb{E}[\hat{P}_\alpha^{(s)}]| \geq \varepsilon P_\alpha(p)\right) \leq 3 \exp(-\Omega_\alpha(\varepsilon^2 n^{2\alpha-1-2\lambda})).$$

For $\alpha > 3/4$, we can set $\lambda = (4\alpha - 3)/8$. Direct calculation shows that for sufficiently large k , the right-hand side is no more than $\exp(-8\sqrt{n})$. Analogously, we can show that for $\alpha > 5/6$, the probability bound can be improved to $\exp(-\Theta(n^{2/3}))$.

Second, we consider the large probability estimator. To begin with, we set $n = \Theta_\alpha(k^{1/3})$. By the results in [6], for sufficiently large $c_{\alpha,1}$, the bias of $\hat{P}_\alpha^{(\ell)}$ in estimating $P_\alpha^{(\ell)}(p)$ satisfies

$$|\mathbb{E}[\hat{P}_\alpha^{(\ell)}] - P_\alpha^{(\ell)}(p)| \leq \mathcal{O}_\alpha\left(\frac{P_\alpha(p)}{\tau}\right) + \frac{1}{n^\alpha},$$

which, for sufficiently large k , is at most $\varepsilon P_\alpha(p)$. Under the same conditions, the variance of $\hat{P}_\alpha^{(\ell)}$ is at most

$$\text{Var}(\hat{P}_\alpha^{(\ell)}) \leq \mathcal{O}_\alpha\left(\sum_x \frac{p(x)^{2\alpha}}{\tau}\right) + \frac{1}{n^{2\alpha}} \leq \frac{(\varepsilon P_\alpha(p))^2}{3}.$$

Then, the Chebyshev's inequality yields

$$\Pr\left(|\mathbb{E}[\hat{P}_\alpha^{(\ell)}] - \hat{P}_\alpha^{(\ell)}| \geq \varepsilon P_\alpha(p)\right) \leq \frac{1}{3}.$$

The triangle inequality combines this tail bound with the above bias bound and implies

$$\Pr\left(|P_\alpha^{(\ell)}(p) - \hat{P}_\alpha^{(\ell)}| \geq 2\varepsilon P_\alpha(p)\right) \leq \frac{1}{3}.$$

Therefore, utilizing the median trick and $\alpha < 1$, we can construct another estimator $\hat{P}_\alpha^{(\ell,1)}$ that takes a sample of size $n = \Omega_\alpha(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$, and satisfies

$$\Pr\left(|P_\alpha^{(\ell)}(p) - \hat{P}_\alpha^{(\ell,1)}| \geq 2\varepsilon P_\alpha(p)\right) \leq 2 \exp(-\Omega_\alpha(n/k^{1/3})) \leq 2 \exp(-\Theta(n^{2/3})).$$

Recall that $P_\alpha(p) = P_\alpha^{(s)}(p) + P_\alpha^{(\ell)}(p)$. By the union bound and the triangle inequality, under Poisson sampling with parameter $n = \Theta_\alpha(k^{1/\alpha}/(\varepsilon^{1/\alpha} \log k))$,

$$\Pr\left(|P_\alpha(p) - (\hat{P}_\alpha^{(s)} + \hat{P}_\alpha^{(\ell,1)})| \geq 4\varepsilon P_\alpha(p)\right) \leq \exp(-8\sqrt{n}).$$

Since both N and N' are Poisson random variables with mean n , we must have $N + N' \sim \text{Poi}(2n)$, implying that $\Pr(N + N' = 2n) = e^{-2n}(2n)^{2n}/(2n)!$. A variant of the well-known Stirling's formula states that $m! \geq em^{m+1/2}e^{-m}$ for all positive integers m . We obtain $\Pr(N + N' = 2n) \geq e^{-2n}(2n)^{2n} \cdot (e(2n)^{2n+1/2}e^{-2n})^{-1} \geq 1/(e\sqrt{2n}) > 1/(4n)$. Hence, under fixed sampling with a sample size of $2n$, the estimator $\hat{P}_\alpha^{(1)} := (\hat{P}_\alpha^{(s)} + \hat{P}_\alpha^{(\ell,1)})$ satisfies

$$\Pr\left(|P_\alpha(p) - \hat{P}_\alpha^{(1)}| \geq 4\varepsilon P_\alpha(p)\right) \leq 4n \exp(-8\sqrt{n}).$$

Replacing n with $n/2$ and ε with $\varepsilon/4$, the *sufficiency of profiles* [6] implies the existence of a profile-based estimator \hat{P}_α^* such that for any $p \in \Delta_{\mathcal{X}}$,

$$\Pr_{X^n \sim p} \left(|P_\alpha(p) - \hat{P}_\alpha^*(X^n)| \geq \varepsilon P_\alpha(p) \right) \leq 2n \exp(-4\sqrt{2n}) < \exp(-4\sqrt{n}).$$

Let δ denote the quantity on the right-hand side. For any x^n with profile φ satisfying both $p(\varphi) > \delta$, we must have $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$. By definition, we also have $p_\varphi(\varphi) \geq p(\varphi) > \delta$ and hence $|\hat{P}_\alpha^*(x^n) - P_\alpha(p_\varphi)| \leq \varepsilon P_\alpha(p_\varphi)$. For any $\varepsilon \in (0, 1/2)$, simple algebra combines the two property inequalities and yields

$$|P_\alpha(p) - P_\alpha(p_\varphi)| \leq 2\varepsilon P_\alpha(p).$$

On the other hand, for a sample $X^n \sim p$ with profile φ' , the probability that we have $p(\varphi') \leq \delta$ is at most δ times the cardinality of the set $\Phi^n := \{\varphi(x^n) : x^n \in \mathcal{X}^n\}$. The latter quantity corresponds to the number of integer partitions of n , which, by the well-known result of Hardy and Ramanujan [49], is at most $\exp(3\sqrt{n})$. Hence, the probability that $p(\varphi') \leq \delta$ is upper bounded by $\exp(-\sqrt{n})$. To conclude, we have shown that

$$\Pr(|P_\alpha(p) - P_\alpha(p_\varphi)| \geq 2\varepsilon P_\alpha(p)) \leq \exp(-\sqrt{n}).$$

In terms of Rényi entropy values, applying the inequality $e^z - 1 \geq 1 - e^{-z} \geq z/2$ for all $z \geq 0$, we establish that for $\alpha > 3/4$ and $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$,

$$\Pr(|H_\alpha(p) - H_\alpha(p_\varphi)| \geq \varepsilon) = \Pr\left(P_\alpha(p_\varphi)e^{-(\alpha-1)\varepsilon} \leq P_\alpha(p) \leq P_\alpha(p_\varphi)e^{(\alpha-1)\varepsilon}\right) \leq \exp(-\sqrt{n}).$$

6.2 Proof of Theorem 3: Non-integer $\alpha > 1$

The proof of the following theorem is essentially the same as that shown in the previous section. However, for completeness, we still include a full-length proof.

For any distribution $p \in \Delta_{\mathcal{X}}$, error parameter $\varepsilon \in (0, 1)$, absolute constant $\lambda \in (0, 0.1)$, and sampling parameter n , draw a sample $X^n \sim p$ and denote its profile by φ . Then for sufficiently large k ,

Theorem 3. *For a non-integer $\alpha > 1$, if $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$,*

$$\Pr(|H_\alpha(p_\varphi) - H_\alpha(p)| \geq \varepsilon) \leq \exp(-n^{1-\lambda}).$$

We establish this theorem in the remaining section. Let n be a sampling parameter and $p \in \Delta_{\mathcal{X}}$ be an unknown distribution. For some α -dependent positive constants $c_{\alpha,1}$ and $c_{\alpha,2}$ to be determined later, let $\tau := c_{\alpha,1} \log n$ and $d := c_{\alpha,2} \log n$ be threshold and degree parameters, respectively. Let N, N' be independent Poisson random variables with mean n . Consider Poisson sampling with two samples drawn from p , first of size N and the second N' . Suppressing the sample representations, for each $x \in \mathcal{X}$, we denote by μ_x and μ'_x the multiplicities of symbol x in the first and second samples, respectively. Denote by $q(z) := \sum_{m=0}^d a_m z^m$ be the degree- d min-max polynomial approximation of z^α over $[0, 1]$. We consider the following variant of the estimator proposed in [6].

$$\hat{P}_\alpha := \sum_x \left(\sum_{m=0}^d \frac{a_m (2\tau)^{\alpha-m} \mu_x^m}{n^\alpha} \right) \mathbb{1}_{\mu_x \leq 4\tau} \cdot \mathbb{1}_{\mu'_x \leq \tau} + \sum_x \left(\frac{\mu_x}{n} \right)^\alpha \mathbb{1}_{\mu'_x > \tau}.$$

The smaller the value of μ'_x is, the smaller we expect the value of $p(x)$ to be. In view of this, we denote the first and second components of \hat{P}_α by $\hat{P}_\alpha^{(s)}$ and $\hat{P}_\alpha^{(\ell)}$, and refer to them as small- and large-probability estimators, respectively. Note that our estimator differs from that in [6] only by the additional $\mathbb{1}_{\mu_y \leq 4\tau}$ term, which for sufficiently large $c_{\alpha,1}$, only modifies $\mathbb{E}[\hat{P}_\alpha^{(s)}]$ by at most $n^{-2\alpha}$.

Note that μ' naturally induces a partition over \mathcal{X} . For symbols x with $\mu_x \leq 4\tau$, we denote by

$$P_{a,\mu'}^{(s)}(p) := \sum_{x:\mu_x \leq 4\tau} p(x)^\alpha$$

the small-probability power sum. Analogously, for symbols x with $\mu_x > 4\tau$, we denote by

$$P_{a,\mu'}^{(\ell)}(p) := \sum_{x:\mu_x > 4\tau} p(x)^\alpha$$

the large-probability power sum. These are random properties with non-trivial variances and are hard to be analyzed. To address this, we apply an ‘‘expectation trick’’ and denote by $P_\alpha^{(s)}(p) := \mathbb{E}[P_{\alpha,\mu'}^{(s)}(p)]$ and $P_\alpha^{(\ell)}(p) := \mathbb{E}[P_{\alpha,\mu'}^{(\ell)}(p)]$ their expected values, both of which are additive symmetric properties.

Let ε be a given error parameter and $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$ be a sampling parameter. First we consider the small probability estimator. By the results in [6], for sufficiently large $c_{\alpha,1}$, the bias of $\hat{P}_\alpha^{(s)}$ in estimating $P_\alpha^{(s)}(p)$ satisfies

$$|\mathbb{E}[\hat{P}_\alpha^{(s)}] - P_\alpha^{(s)}(p)| \leq \mathcal{O}_\alpha(1) \cdot P_\alpha(p) \left(\frac{k}{n \log n} \right)^\alpha + n^{-\alpha} \leq \varepsilon P_\alpha(p),$$

where we have used $n^{-\alpha} = \mathcal{O}_\alpha(\varepsilon k^{-\alpha} (\log k)^\alpha) \leq \varepsilon P_\alpha(p)$. To show concentration, we bound the sensitivity of estimator $\hat{P}_\alpha^{(s)}$. For $m \geq 0$, we can bound the coefficients of $q(x)$ as follows.

$$|a_m| \leq \mathcal{O}_\alpha((\sqrt{2} + 1)^d) = \mathcal{O}_\alpha(n^{c_{\alpha,2}}).$$

Therefore by definition, changing one point in the sample changes the value of $\hat{P}_\alpha^{(s)}$ by at most

$$2 \left(\sum_{m=0}^d \frac{|a_m| (2\tau)^{\alpha-m} (4\tau)^m}{n^\alpha} \right) \leq \sum_{m=0}^d \frac{|a_m| (2\tau)^\alpha 2^{m+1}}{n^\alpha} \leq \mathcal{O}_\alpha(n^{2c_{\alpha,2}-\alpha} (\log n)^\alpha).$$

Let $\lambda \in (0, 1/4)$ be an arbitrary absolute constant. For sufficiently small $c_{\alpha,2}$, the right-hand side is at most $\mathcal{O}_\alpha(n^{\lambda-\alpha})$. The McDiarmid’s inequality together with the concentration of Poisson random variables implies that for all $\varepsilon \geq 0$,

$$\Pr \left(|\hat{P}_\alpha^{(s)} - \mathbb{E}[\hat{P}_\alpha^{(s)}]| \geq \varepsilon P_\alpha(p) \right) \leq 2 \exp(-\Omega_\alpha(\varepsilon^2 P_\alpha^2(p) n^{2\alpha-1-2\lambda})).$$

Note that $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$ and $P_\alpha(p) \geq k^{1-\alpha}$. Hence we obtain

$$\Pr \left(|\hat{P}_\alpha^{(s)} - \mathbb{E}[\hat{P}_\alpha^{(s)}]| \geq \varepsilon P_\alpha(p) \right) \leq 3 \exp(-\Omega_\alpha(\varepsilon^2 k^{2-2\alpha} n^{2\alpha-1-2\lambda})).$$

By simple algebra, for sufficiently large k , the right-hand side is at most $\exp(-n^{1-3\lambda})$.

Second, we consider the large probability estimator. To begin with, we set $n = \Theta_\alpha(k^\lambda)$. By the results in [6], for sufficiently large $c_{\alpha,1}$, the bias of $\hat{P}_\alpha^{(\ell)}$ in estimating $P_\alpha^{(\ell)}(p)$ satisfies

$$|\mathbb{E}[\hat{P}_\alpha^{(\ell)}] - P_\alpha^{(\ell)}(p)| \leq \mathcal{O}_\alpha \left(\frac{P_\alpha(p)}{\tau} \right) + \frac{1}{n^{4\alpha}},$$

which, for sufficiently large k , is at most $\varepsilon P_\alpha(p)$. Under the same conditions, the variance of $\hat{P}_\alpha^{(\ell)}$ is at most

$$\text{Var}(\hat{P}_\alpha^{(\ell)}) \leq \mathcal{O}_\alpha \left(\sum_x \frac{p(x)^{2\alpha}}{\tau} \right) + \frac{1}{n^{8\alpha}} \leq \frac{(\varepsilon P_\alpha(p))^2}{3}.$$

Then, the Chebyshev’s inequality yields

$$\Pr \left(|\mathbb{E}[\hat{P}_\alpha^{(\ell)}] - \hat{P}_\alpha^{(\ell)}| \geq \varepsilon P_\alpha(p) \right) \leq \frac{1}{3}.$$

The triangle inequality combines this tail bound with the above bias bound and implies

$$\Pr \left(|P_\alpha^{(\ell)}(p) - \hat{P}_\alpha^{(\ell)}| \geq 2\varepsilon P_\alpha(p) \right) \leq \frac{1}{3}.$$

Therefore, utilizing the median trick, we can construct another estimator $\hat{P}_\alpha^{(\ell,1)}$ that takes a sample of size $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$, and for sufficiently large k , satisfies

$$\Pr \left(|P_\alpha^{(\ell)}(p) - \hat{P}_\alpha^{(\ell,1)}| \geq 2\varepsilon P_\alpha(p) \right) \leq 2 \exp(-\Omega_\alpha(n/k^\lambda)) \leq \exp(-n^{1-2\lambda}).$$

Recall that $P_\alpha(p) = P_\alpha^{(s)}(p) + P_\alpha^{(\ell)}(p)$. By the union bound and the triangle inequality, under Poisson sampling with parameter $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$,

$$\Pr\left(|P_\alpha(p) - (\hat{P}_\alpha^{(s)} + \hat{P}_\alpha^{(\ell,1)})| \geq 4\varepsilon P_\alpha(p)\right) \leq \exp(-n^{1-3\lambda}).$$

Since both N and N' are Poisson random variables with mean n , we must have $N + N' \sim \text{Poi}(2n)$, implying that $\Pr(N + N' = 2n) = e^{-2n}(2n)^{2n}/(2n)!$. A variant of the well-known Stirling's formula states that $m! \geq em^{m+1/2}e^{-m}$ for all positive integers m . We obtain $\Pr(N + N' = 2n) \geq e^{-2n}(2n)^{2n} \cdot (e(2n)^{2n+1/2}e^{-2n})^{-1} \geq 1/(e\sqrt{2n}) > 1/(4n)$. Hence, under fixed sampling with a sample size of $2n$, the estimator $\hat{P}_\alpha^{(1)} := (\hat{P}_\alpha^{(s)} + \hat{P}_\alpha^{(\ell,1)})$ satisfies

$$\Pr\left(|P_\alpha(p) - \hat{P}_\alpha^{(1)}| \geq 4\varepsilon P_\alpha(p)\right) \leq 4n \exp(-n^{1-3\lambda}).$$

Replacing ε with $\varepsilon/4$ and λ with $\lambda/5$, the sufficiency of profiles implies the existence of a profile-based estimator \hat{P}_α^* such that for sufficiently large k and any $p \in \Delta_{\mathcal{X}}$,

$$\Pr_{X^n \sim p}\left(|P_\alpha(p) - \hat{P}_\alpha^*(X^n)| \geq \varepsilon P_\alpha(p)\right) \leq 4n \exp(-n^{1-3\lambda/5}) < \exp(-n^{1-4\lambda/5}).$$

Let δ denote the quantity on the right-hand side. For any x^n with profile φ satisfying both $p(\varphi) > \delta$, we must have $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$. By definition, we also have $p_\varphi(\varphi) \geq p(\varphi) > \delta$ and hence $|\hat{P}_\alpha^*(x^n) - P_\alpha(p_\varphi)| \leq \varepsilon P_\alpha(p_\varphi)$. For any $\varepsilon \in (0, 1/2)$, simple algebra combines the two property inequalities and yields

$$|P_\alpha(p) - P_\alpha(p_\varphi)| \leq 2\varepsilon P_\alpha(p).$$

On the other hand, for a sample $X^n \sim p$ with profile φ' , the probability that we have $p(\varphi') \leq \delta$ is at most δ times the cardinality of the set $\Phi^n := \{\varphi(x^n) : x^n \in \mathcal{X}^n\}$. The latter quantity corresponds to the number of integer partitions of n , which, by the well-known result of Hardy and Ramanujan [49], is at most $\exp(3\sqrt{n})$. Hence, the probability that $p(\varphi') \leq \delta$ is upper bounded by $\exp(-n^{1-\lambda})$. To conclude, we have shown that

$$\Pr(|P_\alpha(p) - P_\alpha(p_\varphi)| \geq 2\varepsilon P_\alpha(p)) \leq \exp(-n^{1-\lambda}).$$

In terms of Rényi entropy values, applying the inequality $e^z - 1 \geq 1 - e^{-z} \geq z/2$ for all $z \geq 0$, we establish that for $n = \Omega_\alpha(k/(\varepsilon^{1/\alpha} \log k))$,

$$\Pr(|H_\alpha(p) - H_\alpha(p_\varphi)| \geq \varepsilon) = \Pr\left(P_\alpha(p_\varphi)e^{-(\alpha-1)\varepsilon} \leq P_\alpha(p) \leq P_\alpha(p_\varphi)e^{(\alpha-1)\varepsilon}\right) \leq \exp(-n^{1-\lambda}).$$

6.3 Proof of Theorem 4: Integer $\alpha > 1$

For an integer $\alpha > 1$, the following theorem characterizes the performance of the PML-plug-in estimator. For any $p \in \Delta_{\mathcal{X}}$, $\varepsilon \in (0, 1)$, and a sample $X^n \sim p$ with profile φ ,

Theorem 4. *If $n = \Omega_\alpha(k^{1-1/\alpha}(\varepsilon^2 |\log \varepsilon|)^{-(1+\alpha)})$ and $H_\alpha(p) \leq (\log n)/4$,*

$$\Pr(|H_\alpha(p_\varphi) - H_\alpha(p)| \geq \varepsilon) \leq 1/3.$$

Due to the lower bounds in [6], for all possible values of α , the sample complexity of the PML plug-in estimator has the optimal dependency in k . The remaining section is devoted to proving the above theorem. Note that estimating the Rényi entropy $H_\alpha(p)$ to an additive error is equivalent to estimating the power sum $P_\alpha(p)$ to a corresponding multiplicative error. Given this fact, we consider the estimator \hat{P}_α in [6] that maps each sequence $x^n \in \mathcal{X}^*$ to

$$\hat{P}_\alpha(x^n) := \sum_x \frac{\mu_x(x^n)^\alpha}{n^\alpha},$$

where for any real number z , the expression z^α denotes the falling factorial of z to the power α . For a sample $X^n \sim p$, we have $\mathbb{E}[\hat{P}_\alpha(X^n)] = P_\alpha(p)$. The following lemma [63, 6] states that $\hat{P}_\alpha(X^n)$ often estimates $P_\alpha(p)$ to a small multiplicative error when n is large.

Lemma 7. *Under the above conditions, for any $\varepsilon, n > 0$,*

$$\Pr \left(|\hat{P}_\alpha(X^n) - P_\alpha(p)| \geq \varepsilon P_\alpha(p) \right) = \mathcal{O}_\alpha(\varepsilon^{-2} n^{-1} (P_\alpha(p))^{-1/\alpha}).$$

For sufficiently large $n = \Omega_\alpha(k^{(\alpha-1)/\alpha})$, this inequality together with $P_\alpha(p) \leq k^{1-\alpha}$ implies that

$$\Pr \left(|\hat{P}_\alpha(X^n) - P_\alpha(p)| \geq \frac{1}{2} \cdot P_\alpha(p) \right) \leq \frac{1}{4}.$$

The following corollary is a consequence of the above lemma, the sufficiency of profiles, and the standard median trick.

Corollary 2. *Under the above conditions, there is an estimator \hat{P}_α^* such that for any $\varepsilon, n > 0$,*

$$\Pr \left(|\hat{P}_\alpha^*(X^n) - P_\alpha(p)| \geq \varepsilon P_\alpha(p) \right) \leq 2 \exp \left(-\Omega_\alpha(\varepsilon^2 n (P_\alpha(p))^{1/\alpha}) \right).$$

In addition, the estimator \hat{P}_α^ is profile-based.*

For simplicity, suppress X^n in $p_\mu(X^n)$. Since the profile probability $p(\varphi)$ is invariant to symbol permutation, for our purpose, we can assume that $p_\mu(y) \leq p_\mu(z)$ iff $p_\varphi(x) \leq p_\varphi(y)$, for all $x, y \in \mathcal{X}$. Under this assumption, the following lemma [67, 8] relates p_φ to p_μ .

Lemma 8. *For a distribution p and sample $X^n \sim p$ with profile φ ,*

$$\Pr \left(\max_x |p_\varphi(x) - p_\mu(x)| > \frac{2 \log n}{n^{1/4}} \right) = \mathcal{O} \left(\frac{1}{n} \right).$$

Consider $\varepsilon \in (0, 1/2)$ and x^n satisfying $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$. If we further have $P_\alpha(p) \geq 2(n^{1/4}(4 \log n)^{-1})^{1-\alpha}$ and $\max_y |p_\varphi(y) - p_\mu(y)| \leq 2(\log n)n^{-1/4}$, then,

$$\frac{P_\alpha(p)}{2} \stackrel{(a)}{\leq} \hat{P}_\alpha(x^n) \stackrel{(b)}{\leq} P_\alpha(p_\mu) \stackrel{(c)}{\leq} 2^{1+\alpha} P_\alpha(p_\varphi),$$

where (a) follows from the above assumptions; (b) follows from $A^B \leq A^B$ for any $A, B \geq 0$; and (c) follows from the reasoning below.

- Let S denote the the collection of symbols x such that $p_\mu(x) \leq 4(\log n)n^{-1/4}$. Then a convexity argument yields $\sum_{x \in S} (p_\mu(x))^\alpha \leq (n^{1/4}(4 \log n)^{-1})^{1-\alpha}$.
- Using (a), (b), and $P_\alpha(p) \geq 4(n^{1/4}(4 \log n)^{-1})^{1-\alpha}$, we immediately obtain $P_\alpha(p_\mu) \geq 2(n^{1/4}(4 \log n)^{-1})^{1-\alpha}$ and thus $2 \sum_{x \in S} (p_\mu(x))^\alpha \leq P_\alpha(p_\mu) \leq 2 \sum_{x \notin S} (p_\mu(x))^\alpha$.
- For any symbol $x \notin S$, we have $p_\mu(x) > 4(\log n)n^{-1/4}$. This together with the assumption that $\max_x |p_\varphi(x) - p_\mu(x)| \leq 2(\log n)n^{-1/4}$ implies $p_\mu(x) \leq 2p_\varphi(x)$.
- Therefore, the inequality $\sum_{x \notin S} (p_\mu(x))^\alpha \leq 2^\alpha \sum_{x \notin S} (p_\varphi(x))^\alpha \leq 2^\alpha P_\alpha(p_\varphi)$ holds.
- Consequently, we establish $P_\alpha(p_\mu(x)) \leq 2 \sum_{x \notin S} (p_\mu(x))^\alpha \leq 2^{1+\alpha} P_\alpha(p_\varphi)$.

By the inequality $P_\alpha(p)/2 \leq 2^{1+\alpha} P_\alpha(p_\varphi)$ and Corollary 2, if $|\hat{P}_\alpha^*(x^n) - P_\alpha(p_\varphi)| \geq \varepsilon P_\alpha(p_\varphi)$,

$$p_\varphi(\varphi) \leq 2 \exp \left(-\Omega_\alpha(\varepsilon^2 n (P_\alpha(p_\varphi))^{1/\alpha}) \right) \leq 2 \exp \left(-\Omega_\alpha(\varepsilon^2 n (P_\alpha(p))^{1/\alpha}) \right).$$

Let δ_p denote the quantity on the right-hand side. If we further have $p(\varphi) > \delta_p$, then by definition, $p_\varphi(\varphi) \geq p(\varphi) > \delta_p$. Hence for any x^n with profile φ satisfying both $p(\varphi) > \delta_p$ and $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$, we must have $|\hat{P}_\alpha^*(x^n) - P_\alpha(p_\varphi)| \leq \varepsilon P_\alpha(p_\varphi)$. Simple algebra combines the last two inequalities and yields

$$|P_\alpha(p) - P_\alpha(p_\varphi)| \leq 4\varepsilon P_\alpha(p).$$

On the other hand, for a sample $X^n \sim p$ with profile φ' , the probability that we have both $p(\varphi') \leq \delta_p$ and $|\hat{P}_\alpha^*(X^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$ is at most δ_p times the cardinality of the set $\Phi_{\alpha, \varepsilon}^n(p) := \{\varphi(x^n) :$

$x^n \in \mathcal{X}^n$ and $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$. Below we complete this argument by finding a tight upper bound on $|\Phi_{\alpha,\varepsilon}^n(p)|$ in terms of its parameters.

For any sequence x^n such that $\varphi(x^n) \in \Phi_{\alpha,\varepsilon}^n(p)$, let $N_\varphi(x^n)$ denote the number of prevalences $\varphi_j(x^n)$ that are non-zero. Then by definition, we obtain

$$\sum_{j=0}^{N_\varphi(x^n)} \frac{j^\alpha}{n^\alpha} \leq \sum_j \frac{j^\alpha}{n^\alpha} \cdot \varphi_j(x^n) = \hat{P}_\alpha^*(x^n) \leq \frac{3}{2} P_\alpha(p).$$

Using the standard falling-factorial identity $((j+1)^{1+\alpha} - j^{1+\alpha})/(1+\alpha) = j^\alpha$, we can further simplify the expression on the left-hand side:

$$\sum_{j=0}^{N_\varphi(x^n)} \frac{j^\alpha}{n^\alpha} = \frac{(N_\varphi(x^n) + 1)^{1+\alpha}}{(1+\alpha)n^\alpha}.$$

This together with the inequality above yields $N_\varphi(x^n) \leq T_\alpha^n(p) := (3(1+\alpha)n^\alpha \cdot P_\alpha(p)/2)^{1/(1+\alpha)}$. Further note that each prevalence in $\varphi(x^n) = (\varphi_1(x^n), \dots, \varphi_n(x^n))$ can only take values in $[n] := \{0, 1, \dots, n\}$. Therefore, $|\Phi_{\alpha,\varepsilon}^n(p)|$ is at most the number of $T_\alpha^n(p)$ -sparse vectors over $[n]^n$, which admits the following upper bound

$$\binom{n}{T_\alpha^n(p)} |[n]^{T_\alpha^n(p)} \leq (n+1)^{2T_\alpha^n(p)}.$$

Therefore, for $\delta_p \cdot |\Phi_{\alpha,\varepsilon}^n(p)|$ to be small, it suffices to have

$$\Omega_\alpha(\varepsilon^2 n (P_\alpha(p))^{1/\alpha}) \gg 2T_\alpha^n(p) \log(n+1) = 2(3(1+\alpha)n^\alpha \cdot P_\alpha(p)/2)^{1/(1+\alpha)} \log(n+1),$$

which in turn simplifies to

$$\varepsilon^2 n^{1/(1+\alpha)} (P_\alpha(p))^{1/(\alpha(1+\alpha))} \gg \Theta_\alpha(\log n).$$

Following this and $P_\alpha(p) \geq 4(n^{1/4}(4 \log n)^{-1})^{1-\alpha}$, we obtain the following lower bound on n .

$$n \gg \Theta_\alpha((\varepsilon^2 |\log \varepsilon|)^{-(1+\alpha)} (P_\alpha(p))^{-1/\alpha}).$$

In this case, the probability bound $\delta_p \cdot |\Phi_{\alpha,\varepsilon}^n(p)|$ is no larger than $1/6$.

Finally, let C denote the collection of sequences x^n with profile φ that do not satisfy $|\hat{P}_\alpha^*(x^n) - P_\alpha(p)| \leq \varepsilon P_\alpha(p)$ or $\max_x |p_\varphi(x) - \mu_x(x^n)/n| \leq 2(\log n)n^{-1/4}$. By Corollary 2, Lemma 9, and the union bound,

$$\Pr_{X^n \sim p}(X^n \in C) \leq 2 \exp\left(-\Omega_\alpha(\varepsilon^2 n (P_\alpha(p))^{1/\alpha})\right) + \mathcal{O}\left(\frac{1}{n}\right).$$

For n satisfying the lower-bound inequality above, the right-hand side is again no larger than $1/6$. This completes the proof of the theorem.

7 Sorted distribution estimation

7.1 Sorted ℓ_1 distance and Wasserstein duality

For convenience, we first restate the theorem.

Theorem 5. *If $n = \Omega(n(\varepsilon)) = \Omega(k/(\varepsilon^2 \log k))$ and $\varepsilon \geq n^{-c}$,*

$$\Pr(\ell_1^\zeta(p_\varphi, p) \geq \varepsilon) \leq \exp(-\Omega(n^{1/11})).$$

In this section, we relate the estimation of sorted distributions to that of distribution properties through a dual definition of the 1-Wasserstein distance.

Recall that we let $\{p\}$ denote the multiset of probability values of a distribution $p \in \Delta_{\mathcal{X}}$. The sorted ℓ_1 distance between two distributions $p, q \in \Delta_{\mathcal{X}}$ is

$$\ell_1^\zeta(p, q) := \min_{q' \in \Delta_{\mathcal{X}}: \{q'\} = \{q\}} \|p - q'\|_1,$$

which is invariant under domain-symbol permutations on either p or q .

For two distributions ω, ν over the unit interval $[0, 1]$, let $\Gamma'_{\omega, \nu}$ be the collection of distributions over $[0, 1] \times [0, 1]$ with marginals ω and ν on the first and second factors respectively. The 1-Wasserstein distance, also known as the *earth-mover distance*, between ω and ν is

$$\mathcal{W}_1(\omega, \nu) := \inf_{\gamma \in \Gamma'_{\omega, \nu}} \mathbb{E}_{(X, Y) \sim \gamma} |X - Y|.$$

Equivalently, let \mathcal{L}_1 denote the collection of real functions that are 1-Lipschitz on $[0, 1]$. Through duality, one can also define the 1-Wasserstein distance [54] as

$$\mathcal{W}_1(\omega, \nu) = \sup_{f \in \mathcal{L}_1} \left(\mathbb{E}_{X \sim \omega} f(X) - \mathbb{E}_{Y \sim \nu} f(Y) \right).$$

For any $p \in \Delta_{\mathcal{X}}$, let $u_{\{p\}}$ denote the distribution induced by the uniform measure on $\{p\}$. For any distributions $p, q \in \Delta_{\mathcal{X}}$, one can verify [81, 38, 43] that

$$\ell_1^{\leq}(p, q) = k \cdot \mathcal{W}_1(u_{\{p\}}, u_{\{q\}}) \leq R(p, q).$$

Combining this with the dual definition of \mathcal{W}_1 , we obtain

$$\ell_1^{\leq}(p, q) = k \cdot \sup_{f \in \mathcal{L}_1} \left(\mathbb{E}_{X \sim u_{\{p\}}} f(X) - \mathbb{E}_{Y \sim u_{\{q\}}} f(Y) \right) = \sup_{f \in \mathcal{L}_1} \left(\sum_x f(p(x)) - \sum_x f(q(x)) \right).$$

7.2 Proof of Theorem 5

For a real function $f \in \mathcal{L}_1$, we denote by $f(p) := \sum_x f(p(x))$ the corresponding additive symmetric property. The previous reasoning also shows that for any $p, q \in \Delta_{\mathcal{X}}$,

$$R(p, q) \geq \ell_1^{\leq}(p, q) \geq |f(p) - f(q)|.$$

Therefore, property f is 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$.

Set $n := \sup_{f \in \mathcal{L}_1} n_f(\varepsilon)$. The results in [43] imply that if $\varepsilon > n^{-0.3}$,

$$n = \Theta \left(\frac{k}{\varepsilon^2 \log k} \right).$$

Clearly, we only need to consider $\varepsilon \leq 2$, implying $k = \mathcal{O}(n \log n)$. Let α, γ be absolute constants in $[1/100, 1/6)$ and $\varepsilon > 0$ be an error parameter.

By the proof of Theorem 1 in Section 5.2, for any distribution $p \in \Delta_{\mathcal{X}}$ and $X^{n/\alpha} \sim p$, with probability at least $1 - 2 \exp(-4n^{1-2\gamma})$, the PML (or APML) plug-in estimator will satisfy

$$|f(p) - f(p_{\varphi(X^{n/\alpha})})| < \varepsilon(2 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n) + 4n^{1-\gamma} \tau(n),$$

where $c_1 \in (1/100, 1/32]$, $c_2 = 1/2 + 6c_1$, and $\tau(n) = \mathcal{O}(n^{\alpha c_2 + (2-\alpha)c_1 - 1} \log^3 n)$. Additionally, in the previous section, we have proved that

$$\ell_1^{\leq}(p, q) = \sup_{f \in \mathcal{L}_1} (f(p) - f(q)) = \sup_{f \in \mathcal{L}_1} |f(p) - f(q)|.$$

Though it seems that the above inequality and equation imply the optimality of PML (since f is chosen arbitrarily), such direct implication actually does not hold. The reason is a little bit subtle: The inequality on $|f(p) - f(p_{\varphi(X^{n/\alpha})})|$ holds for any fixed function f and $p \in \Delta_{\mathcal{X}}$, while the function that achieves the corresponding supremum in

$$\sup_{f \in \mathcal{L}_1} |f(p) - f(p_{\varphi(X^{n/\alpha})})| = \ell_1^{\leq}(p, p_{\varphi(X^{n/\alpha})})$$

depends on both p and $X^{n/\alpha}$, and hence is a random function. To address this discrepancy, we provide a more involved argument below.

Let f be a function in \mathcal{L}_1 . Without loss of generality, we also assume that $f(0) = 0$. Let $\eta \in (0, 1)$ be a threshold parameter to be determined later. An η -truncation of f is a function

$$f_{\eta}(z) := f(z) \mathbb{1}_{z \leq \eta} + f(\eta) \mathbb{1}_{z > \eta}.$$

One can easily verify that $f_\eta \in \mathcal{L}_1$. Next, we find a finite subset of \mathcal{L}_1 so that the η -truncation of any $f \in \mathcal{L}_1$ is close to at least one of the functions in this subset.

For a parameter $s > 3$ to be chosen later. Partition the interval $[0, \eta]$ into s disjoint sub-intervals of equal length, and define the sequence of end points as $z_j := \eta \cdot j/s, j \in [s]$ where $[s] := \{0, 1, \dots, s\}$. Then, for each $j \in [s]$, we find the integer j' such that $|f_\eta(z_j) - z_{j'}|$ is minimized and denote it by j^* . Since f_η is 1-Lipschitz, we must have $|j^*| \in [j]$. Finally, we connect the points $Z_j := (z_j, z_{j^*})$ sequentially. This curve is continuous and corresponds to a particular η -truncation $\tilde{f}_\eta \in \mathcal{L}_1$, which we refer to as the *discretized η -truncation* of f . Intuitively, we have constructed an $(s+1) \times (s+1)$ grid and “discretized” function f by finding its closest approximation in \mathcal{L}_1 whose curve only consists of edges and diagonals of the grid cells. By construction,

$$\max_{z \in [0,1]} |f_\eta(z) - \tilde{f}_\eta(z)| \leq \eta/s.$$

Therefore, for any $p \in \Delta_{\mathcal{X}}$, the corresponding properties of f_η and \tilde{f}_η satisfy

$$|f_\eta(p) - \tilde{f}_\eta(p)| \leq k \cdot \eta/s.$$

Note that $|j^*| \in [j]$ for all $j \in [s]$, and $\tilde{f}_\eta(z) = z_{s^*}$ for $z \geq \eta$. While there are infinitely many η -truncations, the cardinality of the discretized η -truncations of functions in \mathcal{L}_1 is at most

$$\prod_{j=0}^s (2j+1) = (s+1) \prod_{j=0}^{s-1} (2j+1)(2s-2j+1) \leq (s+1)^{2s+1} = e^{(2s+1)\log(s+1)} \leq e^{3s \log s}.$$

Consider any $p \in \Delta_{\mathcal{X}}$ and $X^{n/\alpha} \sim p$ with a profile φ . Consolidate the previous results, and apply the union bound and triangle inequality. With probability at least $1 - 2 \exp(3s \log s - 4n^{1-2\gamma})$, the PML plug-in estimator will satisfy

$$\begin{aligned} |f_\eta(p) - f_\eta(p_\varphi)| &\leq |f_\eta(p) - \tilde{f}_\eta(p)| + |\tilde{f}_\eta(p) - \tilde{f}_\eta(p_\varphi)| + |\tilde{f}_\eta(p_\varphi) - f_\eta(p_\varphi)| \\ &\leq 2k \cdot \eta/s + \varepsilon(2 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n) + 4n^{1-\gamma} \tau(n), \end{aligned}$$

for all functions f in \mathcal{L}_1 .

Next we consider the “second part” of a function $f \in \mathcal{L}_1$, namely,

$$\bar{f}_\eta(z) := f(z) - f_\eta(z) = (f(z) - f(\eta)) \mathbb{1}_{z > \eta}.$$

Again, we can verify that $\bar{f}_\eta \in \mathcal{L}_1$. To establish the corresponding guarantees, we make use of the following result. Since the profile probability $p(\varphi)$ is invariant to symbol permutation, for our purpose, we can assume that $p(y) \leq p(z)$ iff $p_\varphi(x) \leq p_\varphi(y)$, for all $x, y \in \mathcal{X}$. Under this assumption, the next lemma, which follows from the consistency results in [67, 8], relates p_φ to p . Let $\gamma' \in (0, 1/4)$ be an absolute constant to be determined later. Then,

Lemma 9. *For any distribution p and sample $X^m \sim p$ with profile φ ,*

$$\Pr \left(\max_x |p_\varphi(x) - p(x)| > m^{\gamma'-1/4} \right) = \mathcal{O} \left(m^{1/4} \exp(-\Omega(m^{1/2+2\gamma'})) \right).$$

Simply following the proofs in [67, 8], we obtain: Changing $1/4$ to any (fixed) number greater than $1/6$, the above lemma also holds for APML with $m^{1/2+2\gamma'}$ replaced by $m^{2/3+2\gamma'}$.

Set $m = n/\alpha$ in this lemma. With probability at least $1 - \mathcal{O} \left((n/\alpha)^{1/4} \exp(-\Omega((n/\alpha)^{1/2+2\gamma'})) \right)$,

$$\begin{aligned} |\bar{f}_\eta(p) - \bar{f}_\eta(p_\varphi)| &= \left| \sum_x \bar{f}_\eta(p(x)) - \bar{f}_\eta(p_\varphi(x)) \right| \\ &\leq \sum_{x:p(x) > \eta \text{ or } p_\varphi(x) > \eta} |\bar{f}_\eta(p(x)) - \bar{f}_\eta(p_\varphi(x))| \\ &\leq \sum_{x:p(x) > \eta \text{ or } p_\varphi(x) > \eta} |p(x) - p_\varphi(x)| \\ &\leq (2/\eta)(n/\alpha)^{\gamma'-1/4}, \end{aligned}$$

for all functions f in \mathcal{L}_1 .

Consolidate the previous results. By the triangle inequality and the union bound, with probability at least $1 - 2 \exp(3s \log s - 4n^{1-2\gamma}) - \mathcal{O}\left((n/\alpha)^{1/4} \exp(-\Omega((n/\alpha)^{1/2+2\gamma'})\right)$,

$$\begin{aligned} |f(p) - f(p_\varphi)| &\leq |f_\eta(p) - f_\eta(p_\varphi)| + |\bar{f}_\eta(p) - \bar{f}_\eta(p_\varphi)| \\ &\leq 2k\eta/s + \varepsilon(2 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^2 n) + 4n^{1-\gamma} \tau(n) + (2/\eta)(n/\alpha)^{\gamma'-1/4}, \end{aligned}$$

for all functions f in \mathcal{L}_1 . Now we can conclude that $\ell_1^\zeta(p, p_\varphi)$ is also at most the error bound on the right-hand side. The reason is straightforward: Since with high probability, the above guarantee holds for all functions in \mathcal{L}_1 , it must also hold for the function that achieves the supremum in

$$\sup_{f \in \mathcal{L}_1} |f(p) - f(p_\varphi)| = \ell_1^\zeta(p, p_\varphi).$$

It remains to make sure that all the quantities in the error bound except $\varepsilon(2 + o(1))$ vanish with n , and the probability bound converges to 1 as n increases. Recall that $k = \mathcal{O}(n \log n)$, $c_1 \in (1/100, 1/25]$, $c_2 = 1/2 + 6c_1$, and $\tau(n) = \mathcal{O}(n^{\alpha c_2 + (2-\alpha)c_1 - 1} \log^3 n)$.

By direct computation, we can choose $\alpha = 1/100$, $c_1 = 1/26$, $\gamma' = 1/200$, $\gamma = (5/2 + 5\alpha)c_1 + \alpha/2$, $s = n^{\gamma' + 3/4 + c_1}$, and $\eta = n^{\gamma' - 1/4 + c_1/2}$. Note that this is just one possible set of parameters. Given this choice, we have

$$\ell_1^\zeta(p, p_\varphi) \leq \varepsilon(2 + o(1)) + \mathcal{O}(n^{-c_1/2} \log^3 n),$$

with probability at least $1 - \exp(-\Omega(n^{1/2}))$. Additionally, the equation

$$\sup_{f \in \mathcal{L}_1} |f(p) - f(p_\varphi)| = \ell_1^\zeta(p, p_\varphi)$$

clearly yields that $n(\varepsilon) \geq \sup_{f \in \mathcal{L}_1} n_f(\varepsilon)$. Hence for $\varepsilon \geq \mathcal{O}(n^{-c_1/2} \log^4 n)$,

$$n(p_\varphi, (2 + o(1))\varepsilon) \leq 100n(\varepsilon).$$

8 Uniformity testing

8.1 PML-based tester

Let ε be an arbitrary accuracy parameter and \mathcal{X} be a finite set. Let p_u denote the uniform distribution over \mathcal{X} . Given sample access to an unknown distribution $p \in \Delta_{\mathcal{X}}$, the uniformity testing distinguishes between the null hypothesis

$$H_0 : p = p_u$$

and the alternative hypothesis

$$H_1 : \|p - p_u\|_1 \geq \varepsilon.$$

After a sequence of research works [37, 13, 71, 3, 19, 78, 31, 4, 30, 32], it is shown that to achieve a $k^{-\Theta(1)}$ bound on the error probability, this task requires a worst-case sample size of order $\sqrt{k \log k}/\varepsilon^2$. The uniformity tester $T_{\text{PML}}(X^n)$ in Figure 7 is purely based on PML, and takes as input parameters k and ε , and a sample $X^n \sim p$.

Input: parameters k, ε , and a sample $X^n \sim p$ with profile φ .

1. If $\max_x \mu_x(X^n) \geq 3 \max\{1, n/k\} \log k$, return 1;
2. Elif $\|p_\varphi - p_u\|_2 \geq 3\varepsilon/(4\sqrt{k})$, return 1;
3. Else return 0.

Figure 7: Uniformity tester T_{PML}

In the rest of this section, we establish the following theorem.

Theorem 6. *If $\varepsilon = \tilde{\Omega}(k^{-1/4})$ and $n = \tilde{\Omega}(\sqrt{k}/\varepsilon^2)$, then the tester $T_{\text{PML}}(X^n)$ will be correct with probability at least $1 - k^{-2}$. The tester also distinguishes between $p = p_u$ and $\|p - p_u\|_2 \geq \varepsilon/\sqrt{k}$.*

8.2 Proof of Theorem 6

Assume that $\varepsilon \geq (\log k)/k^{1/4}$. For a sample $X^n \sim p_u$, the multiplicity of each symbol x follows a binomial distribution $\text{bin}(n, k^{-1})$ with mean n/k . The following lemma [25] bounds the tail probability of a binomial random variable.

Lemma 10. *For a binomial random variable Y with mean M and any $t \geq 1$,*

$$\Pr(Y \geq (1+t)M) \leq \exp(-t(2/t + 2/3)^{-1}M).$$

Applying the above lemma to $Y = \mu_x(X^n)$ and $t = 3 \max\{k/n, 1\} \log k$ immediately yields that $\Pr(\mu_x(X^n) \geq (1+t)n/k) \leq k^{-3}$. By symmetry and the union bound, we then have $\Pr(\max_x \mu_x(X^n) \geq (1+t)n/k) \leq k^{-2}$. In the subsequent discussion, we denote by $\Phi_{\mathcal{X}}^n$ the profile set $\{\varphi(x^n) : x^n \in \mathcal{X}^n \text{ and } \max_x \mu_x(x^n) < (1+t)n/k\}$.

Consider the problem of estimating the ℓ_2 -distance between an unknown distribution and the uniform distribution p_u , for which we have the following result [36].

Lemma 11. *There is a profile-based estimator $\hat{\ell}_2$ such that for any $\varepsilon_0 \leq k^{-1/2}$, $n = \Omega(k^{-1/2}/\varepsilon_0^2)$, $p \in \Delta_{\mathcal{X}}$ satisfying $P_2(p) = \mathcal{O}(k^{-1})$, and $X^n \sim p$,*

- if $\|p - p_u\|_2 > \varepsilon_0$, then $\hat{\ell}_2(X^n) \geq 0.9\varepsilon_0$,
- if $\|p - p_u\|_2 < \varepsilon_0/2$, then $\hat{\ell}_2(X^n) \leq 0.6\varepsilon_0$,

with probability at least $2/3$.

Set $\varepsilon_0 = \varepsilon/\sqrt{k}$ in the above lemma. Then, by the sufficiency of profiles and the standard median trick, there exists another profile-based estimator $\hat{\ell}_2^*$ that under the same conditions, provides the estimation guarantees stated above, with probability at least $1 - \delta$ for $\delta := 2 \exp(-\Omega(n\varepsilon^2/\sqrt{k}))$. Scaling ε_0 by positive absolute constant factors yields: If $\|p - p_u\|_2 > 0.67\varepsilon_0$, then $\hat{\ell}_2(X^n) \leq 0.6\varepsilon_0$ with probability at most δ ; if $\|p - p_u\|_2 < 0.75\varepsilon_0$, then $\hat{\ell}_2(X^n) \geq 0.9\varepsilon_0$ with probability at most δ .

Let φ' be a profile. If we further have $p(\varphi') > \delta$, then by definition, $p_{\varphi'}(\varphi') \geq p(\varphi') > \delta$. Hence for any x^n with profile φ' , if $\|p - p_u\|_2 > \varepsilon_0$, we must have both $\hat{\ell}_2(x^n) \geq 0.9\varepsilon_0$ and $\|p_{\varphi'} - p_u\|_2 \geq 0.75\varepsilon_0$; if $\|p - p_u\|_2 < \varepsilon_0/2$, we must have both $\hat{\ell}_2(x^n) \leq 0.6\varepsilon_0$ and $\|p_{\varphi'} - p_u\|_2 \leq 0.67\varepsilon_0$.

On the other hand, for a sample $X^n \sim p$ with profile φ , the probability that we have both $p(\varphi) \leq \delta$ and $\varphi \in \Phi_{\mathcal{X}}^n$ is at most δ times the cardinality of the set $\Phi_{\mathcal{X}}^n$. By definition, if $\varphi \in \Phi_{\mathcal{X}}^n$, then $\varphi_i = 0$ for $i \geq (1+t)n/k$. In addition, each φ_i can only take values in $[k] = \{0, 1, \dots, k\}$, implying that $|\Phi_{\mathcal{X}}^n| \leq |k|^{(1+t)n/k} \leq \exp(6 \max\{n/k, 1\} \log^2 k)$. Therefore, we obtain the following upper bound on the probability of interest: $\delta \cdot |\Phi_{\mathcal{X}}^n| \leq 2 \exp(-\Omega(n\varepsilon^2/\sqrt{k})) + 6 \max\{n/k, 1\} \log^2 k$. In order to make the probability bound vanish, we need to consider two cases: $n \leq k$ and $n > k$. If $n \leq k$, it suffices to have $n \gg (\log^2 k)\sqrt{k}/\varepsilon^2$; If $n > k$, it suffices to have $\varepsilon \gg (\log k)/k^{1/4}$. In both cases, the probability bound is at most $\exp(-\log^2 k)$.

Next, consider estimating the power sum $P_2(p)$, which is at least $k^{-1/2}$ for $p \in \Delta_{\mathcal{X}}$. By Corollary 2, there is a profile-based estimator \hat{P}_2^* such that $\Pr_{X^n \sim p}(|\hat{P}_2^*(X^n) - P_2(p)| \geq (\varepsilon/8) \cdot P_2(p)) \leq 2 \exp(-\Omega(n\varepsilon^2/\sqrt{k})) = \delta$. Following the same derivations as above and in Section 6.3 with $\Phi_{\alpha, \varepsilon}^n(p)$ replaced by $\Phi_{\mathcal{X}}^n$, we establish that

$$\Pr(|P_2(p_{\varphi}) - P_2(p)| > P_2(p)/2 \text{ and } \varphi \in \Phi_{\mathcal{X}}^n) \leq \delta \cdot |\Phi_{\mathcal{X}}^n| \leq \exp(-\log^2 k).$$

Now we are ready to characterize the performance of the tester $T_{\text{PML}}(X^n)$. For clarity, we divide our analysis into two parts based on which hypothesis is true.

- **Case 1:** The null hypothesis H_0 is true, i.e., $p = p_u$.
 - **Step 1:** By Lemma 10 and its implications, given $p = p_u$, the probability of failure at this step is at most $\Pr_{X^n \sim p_u}(\exists x \in \mathcal{X} \text{ s.t. } \mu_x(X^n) \geq (1+t)n/k) \leq k^{-2}$.
 - **Step 2:** Note that $P_2(p) = k^{-1}$ and $\|p - p_u\|_2 = 0$, and recall that $\varphi = \varphi(X^n)$. The tester accepts H_1 in this step iff $\varphi \in \Phi_{\mathcal{X}}^n$ and $\|p_{\varphi} - p_u\|_2 \geq 0.75\varepsilon_0$. By Lemma 11 and the subsequent arguments, this happens with probability at most $\exp(-\log^2 k)$.

- **Step 3:** The tester always accepts H_0 in this step. Hence by the union bound, if the null hypothesis H_0 is true, then the tester succeeds with probability at least $1 - k^{-2}$.
- **Case 2:** The alternative hypothesis H_1 is true, i.e., $\|p - p_u\|_1 \geq \varepsilon$.
 - **Step 1 to 2:** The tester accepts H_1 if the conditions in either Step 1 or Step 2 are satisfied, and hence incurs no error.
 - **Step 3:** According to the value of $P_2(p)$, we further divide our analysis into two parts:
 - * If $P_2(p) \geq 10k^{-1}$, then $\|p_\varphi - p_u\|_2 < 0.75\varepsilon/\sqrt{k}$ implies that $P_2(p_\varphi) < 1.6k^{-1}$ and $|P_2(p_\varphi) - P_2(p)| > P_2(p)/2$. Hence, the tester accepts H_0 only if both $|P_2(p_\varphi) - P_2(p)| > P_2(p)/2$ and $\varphi \in \Phi_{\mathcal{X}}^n$ happen, whose probability, by the above discussion, is at most $\exp(-\log^2 k)$.
 - * If $P_2(p) < 10k^{-1}$, then all the conditions in Lemma 11 are satisfied. In addition, by the Cauchy-Schwarz inequality, we have $\|p - p_u\|_2 \geq \|p - p_u\|_1 \cdot k^{-1/2} \geq \varepsilon \cdot k^{-1/2}$. The tester accepts H_0 iff both $\|p_\varphi - p_u\|_2 < 0.75\varepsilon \cdot k^{-1/2}$ and $\varphi \in \Phi_{\mathcal{X}}^n$ hold, which happen, by Lemma 11 and the subsequent arguments, with probability at most $\exp(-\log^2 k)$.

This completes the proof of the theorem.

Conclusion

We studied three fundamental problems in statistical learning: distribution estimation, property estimation, and property testing. We established the profile maximum likelihood (PML) as the first universally sample-optimal approach for several important learning tasks: distribution estimation under the sorted ℓ_1 distance, additive property estimation, Rényi entropy estimation, and identity testing. Several future directions are promising. We believe that neither the factor of 4 in the sample size in Theorem 1, nor the lower bounds on ε in Theorem 1, 5, and 6 are necessary. In other words, the PML approach is universally sample-optimal for these tasks in all ranges of parameters. It is also of interest to extend the PML's optimality to estimating symmetric properties not covered by Theorem 1 to 4, such as *generalized distance to uniformity* [10, 44], the ℓ_1 distance between the unknown distribution and the closest uniform distribution over an arbitrary subset of \mathcal{X} .

Another important direction is *competitive (or instance-optimal) property estimation*. It should be noted that all the referenced works including this paper are of the worst-case nature, namely, designing estimators with near-optimal worst-case performances. On the contrary, practical and natural distributions often possess simple structures, and are rarely the worst possible. To address this discrepancy, the recent work [45, 48] took a competitive approach and constructed estimators whose performances are adaptive to the simplicity of the underlying distributions. Specifically, for any property in a broad class and *every* distribution in $\Delta_{\mathcal{X}}$, the expected error of the proposed estimator with a sample of size $n/\log n$ is at most that of the empirical estimator with a sample of size n , plus a distribution-free vanishing function of n . These results not only cover \tilde{S} , \tilde{C}_m , H , and D , for which the $\log n$ -factor is optimal up to constants, but also apply to any *non-symmetric additive* property $\sum_x f_x(p_x)$ where f_x is 1-Lipschitz for all $x \in \mathcal{X}$, such as the ℓ_1 -distance to a given distribution. It would be of interest to study the optimality of the PML approach under this formulation as well. Readers interested in estimating non-symmetric properties may also find the paper [47] helpful.

A Proof of Lemma 3

The proof closely follows that of Proposition 6.19 in [83] (page 131–136), which we refer to as *the proposition's proof*. Note that in the work [83], the definitions of k and n are swapped, i.e., k stands for the sample size, and n denotes the alphabet size. For consistency, we still keep our notation.

Recall that we set $t_n := 2n^{-c_1} \log n$ and $\alpha \in (0, 1)$, and define

$$\beta_i := (1 - e^{-t_n \alpha^i}) f\left(\frac{(i+1)\alpha}{n}\right) \frac{n}{(i+1)\alpha} + \sum_{\ell=0}^i z_\ell (1 - t_n)^\ell \alpha^\ell (1 - \alpha)^{i-\ell} \binom{i}{\ell}.$$

for any $i \leq n$, and $\beta_i := \beta_n$ for $i > n$. Let $w(i)$ denote the first quantity on the right-hand side, and $w := (w(0), w(1), \dots)$ be the corresponding vector. Similarly, let $\tilde{z}_\alpha(i)$ denote the second quantity on the right-hand side, and \tilde{z}_α be the corresponding vector. Assume that $v \leq \log^2 n$.

First part of the proposition's proof remains unchanged, which corresponds to the content from page 131 to the second last paragraph on page 132, showing that

$$\sqrt{\alpha} \|\tilde{z}_\alpha\|_2 = \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \cdot \log^3 n).$$

The assumption that $\alpha \in [1/100, 1)$ implies $\sqrt{\alpha} \geq 1/10$, and hence we have $|\tilde{z}_\alpha(i)| \leq \|\tilde{z}_\alpha\|_2 = \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \cdot \log^3 n)$. Recall that for lemma 2 to hold, the coefficients β_i must satisfy the following two conditions,

1. $|\varepsilon(y)| \leq a' + b'/y$,
2. $|\beta_j^* - \beta_\ell^*| \leq c' \sqrt{j/n}$ for any j and ℓ such that $|j - \ell| \leq \sqrt{j} \log n$,

where $\varepsilon(y) := f(y)/y - e^{-ny} \sum_{i \geq 0} \beta_i \cdot (ny)^i / i!$, and $\beta_i^* := \beta_{i-1} \cdot i/n, \forall i \geq 1$, and $\beta_0^* := 0$.

We first consider the second condition and find a proper parameter c' .

Our objective is to find $c' > 0$ such that $c' > \sqrt{n/j} |\beta_j^* - \beta_\ell^*|$. By the triangle inequality,

$$\sqrt{\frac{n}{j}} |\beta_j^* - \beta_\ell^*| \leq \sqrt{\frac{n}{j}} \left| \frac{j}{n} \tilde{z}_\alpha(j-1) - \frac{\ell}{n} \tilde{z}_\alpha(\ell-1) \right| + \sqrt{\frac{n}{j}} \left| \frac{j}{n} w(j-1) - \frac{\ell}{n} w(\ell-1) \right|$$

We bound the two quantities on the right-hand side separately and consider two cases for each. If both j and ℓ are at most $400n^{c_1}$, then

$$\sqrt{\frac{n}{j}} \left| \frac{j}{n} \tilde{z}_\alpha(j-1) - \frac{\ell}{n} \tilde{z}_\alpha(\ell-1) \right| \leq \mathcal{O}(n^{c_1/2-1/2}) \cdot \max_i |z_\alpha(i)| \leq \mathcal{O}(n^{\alpha c_2 + (3/2-\alpha)c_1-1/2} \log^3 n).$$

Recall that $|z_\ell| \leq v \cdot n^{c_2}, \forall \ell \geq 0$. If one of j and ℓ is larger than $400n^{c_1}$, say $j > 400n^{c_1}$, then

$$\begin{aligned} \sqrt{\frac{n}{j}} \left| \frac{j}{n} \tilde{z}_\alpha(j-1) \right| &\leq \sqrt{\frac{j}{n}} \sum_{\ell=0}^{j-1} |z_\ell| (1-t_n)^\ell \alpha^\ell (1-\alpha)^{j-1-\ell} \binom{j-1}{\ell} \\ &\leq \sqrt{j} n^{c_2-1/2} (\log^2 n) \sum_{\ell=0}^{j-1} (1-t_n)^\ell \alpha^\ell (1-\alpha)^{j-1-\ell} \binom{j-1}{\ell} \\ &= \sqrt{j} n^{c_2-1/2} (\log^2 n) (1-t_n \alpha)^{j-1} \\ &\leq \sqrt{j} n^{c_2-1/2} (\log^2 n) (1 - \log n / (50n^{c_1}))^{400n^{c_1}} \\ &\leq \sqrt{j} n^{c_2-1/2} (\log^2 n) n^{-8}. \end{aligned}$$

For $j < 2n^2$, the last quantity is at most n^{-1} . For $j > 2n^2$, we have $\ell > n^2$ and hence

$$\sqrt{\frac{n}{j}} \left| \frac{j}{n} \tilde{z}_\alpha(j-1) - \frac{\ell}{n} \tilde{z}_\alpha(\ell-1) \right| = \sqrt{\frac{n}{j}} |j - \ell| \tilde{z}_\alpha(n-1) \leq \sqrt{n} (\log n) n^{-1} = (\log n) n^{-1/2}.$$

Similarly, we can bound the other quantity, i.e.,

$$\sqrt{\frac{n}{j}} \left| \frac{j}{n} w(j-1) - \frac{\ell}{n} w(\ell-1) \right| = \sqrt{\frac{n}{\alpha^2 j}} \left| (1 - e^{-t_n \alpha(j-1)}) f\left(\frac{j\alpha}{n}\right) - (1 - e^{-t_n \alpha(\ell-1)}) f\left(\frac{\ell\alpha}{n}\right) \right|.$$

Since f (the property) is 1-Lipschitz on $(\Delta_{\mathcal{X}}, R)$ and $f(p) = 0$ if $p(x) = 1$ for some $x \in \mathcal{X}$, one can verify that $|f(x)| \leq x |\log x| \leq e^{-1}$ and $|f(x)/x - f(y)/y| \leq |\log(x/y)|$ for $x, y \in [0, 1]$ (the corresponding real function). We consider two cases and bound the quantity of interest. If $j \geq \sqrt{n}$,

$$\sqrt{\frac{n}{\alpha^2 j}} \left| (1 - e^{-t_n \alpha(j-1)}) f\left(\frac{j\alpha}{n}\right) \right| \leq \sqrt{\frac{n}{\alpha^2 j}} \left| f\left(\frac{j\alpha}{n}\right) \right| \leq \sqrt{\frac{n}{\alpha^2 j}} \frac{j\alpha}{n} \log\left(\frac{j\alpha}{n}\right) \leq \mathcal{O}(n^{-1/4} \log n).$$

The same bound also applies to the other term where j is replaced by ℓ . If $j > \sqrt{n}$, then $e^{-t_n \alpha(j-1)} \leq \exp(-2\alpha(\log n)n^{1/2-c_1}) = \mathcal{O}(n^{-2})$. Analogously, the same upper bound holds for the other term $e^{-t_n \alpha(\ell-1)}$. Hence, we ignore these two terms and consider only

$$\begin{aligned}
\sqrt{\frac{n}{\alpha^2 j}} \left| f\left(\frac{j\alpha}{n}\right) - f\left(\frac{\ell\alpha}{n}\right) \right| &\leq \sqrt{\frac{j}{n}} \left| \frac{n}{j\alpha} f\left(\frac{j\alpha}{n}\right) - \frac{n}{\ell\alpha} f\left(\frac{\ell\alpha}{n}\right) \right| + \sqrt{\frac{j}{n}} \left| \frac{n}{j\alpha} - \frac{n}{\ell\alpha} \right| f\left(\frac{\ell\alpha}{n}\right) \\
&\leq \sqrt{\frac{j}{n}} \left| \log \frac{j}{\ell} \right| + \sqrt{\frac{j}{n}} \left| \frac{n}{j\alpha} - \frac{n}{\ell\alpha} \right| f\left(\frac{\ell\alpha}{n}\right) \\
&\leq \sqrt{\frac{j}{n}} \frac{|j-\ell|}{j} + \frac{\sqrt{jn}}{\alpha} \frac{|j-\ell|}{j\ell} f\left(\frac{\ell\alpha}{n}\right) \\
&\leq \sqrt{\frac{j}{n}} \frac{|j-\ell|}{j} + \sqrt{\frac{j}{n}} \frac{|j-\ell|}{j} \left| \log\left(\frac{\ell\alpha}{n}\right) \right| \\
&\leq \frac{\log n}{\sqrt{n}} + \frac{\log n}{n} \left| \log\left(\frac{\ell\alpha}{n}\right) \right| \\
&= \mathcal{O}(n^{-1/2} \log n).
\end{aligned}$$

By the assumption that $\alpha c_2 + (3/2 - \alpha)c_1 \leq 1/4$, we have $\mathcal{O}(n^{\alpha c_2 + (3/2 - \alpha)c_1 - 1/2} \log^3 n) = \mathcal{O}(n^{-1/4} \log^3 n)$. Hence, we can set the latter quantity to be c' . The above derivations also show that

$$|w(i)| = \left| (1 - e^{-t_n \alpha i}) f\left(\frac{(i+1)\alpha}{n}\right) \frac{n}{(i+1)\alpha} \right| \leq \left| \log\left(\frac{(i+1)\alpha}{n}\right) \right| = \mathcal{O}(\log n).$$

Together with $\beta_i = w(i) + \tilde{z}_\alpha(i)$ and $|\tilde{z}_\alpha(i)| = \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \cdot \log^3 n)$, this inequality implies

$$|\beta_i| \leq \mathcal{O}(n^{\alpha c_2 + (1-\alpha)c_1} \log^3 n).$$

It remains to analyze the first condition of Lemma 2 and find proper values for a' and b' . For this part, the corresponding proof in [83] also holds for $\alpha \in [1/100, 1/2]$ (page 134 to the second last paragraph on page 135), hence no change is needed. One thing to note is that $1/\alpha$ and $1/\sqrt{\alpha}$ are both $\mathcal{O}(1)$. For some $a'', b'' \geq 0$ such that $a'' + b''k \leq v$, we can set $a' = a'' + \mathcal{O}(n^{-c_1/2} \log^2 n)$ and $b' = b''(1 + \mathcal{O}(n^{-c_1} \log n))$. The proof of Lemma 3 is complete.

References

- [1] J. Acharya, H. Das, H. Mohimani, A. Orlitsky, and S. Pan. Exact calculation of pattern probabilities. In *Proceedings of 2010 IEEE International Symposium on Information Theory (ISIT)*, pages 1498–1502. IEEE, 2010.
- [2] J. Acharya, H. Das, A. Jafarpour, A. Orlitsky, and S. Pan. Estimating multiple concurrent processes. In *Proceedings 2012 IEEE International Symposium on Information Theory*, pages 1628–1632. IEEE, 2012.
- [3] J. Acharya, A. Jafarpour, A. Orlitsky, and A. Suresh. A competitive test for uniformity of monotone distributions. In *Artificial Intelligence and Statistics (AISTats)*, pages 57–65, 2013.
- [4] J. Acharya, C. Daskalakis, and G. Kamath. Optimal testing for properties of distributions. In *Advances in Neural Information Processing Systems*, pages 3591–3599. Society for Industrial and Applied Mathematics, 2015.
- [5] J. Acharya, H. Das, A. Orlitsky, and A. T. Suresh. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. In *International Conference on Machine Learning*, pages 11–21, 2017.
- [6] J. Acharya, A. Orlitsky, A. T. Suresh, and H. Tyagi. Estimating Rényi entropy of discrete distributions. *IEEE Transactions on Information Theory*, 63(1):38–56, 2017.
- [7] J. Acharya, Y. Bao, Y. Kang, and Z. Sun. Improved bounds for minimax risk of estimating missing mass. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 326–330. IEEE, 2018.

- [8] D. Anevski, R. D. Gill, and S. Zohren. Estimating a probability mass function with unknown labels. *The Annals of Statistics*, 45(6):2708–2735, 2017.
- [9] R. Armaanzas, I. Inza, R. Santana, Y. Saeys, J. L. Flores, J. A. Lozano, . . . , and P. Larrañaga. A review of estimation of distribution algorithms in bioinformatics. *BioData Mining*, 1(6), 2008.
- [10] T. Batu and C. L. Canonne. Generalized uniformity testing. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, pages 880–889. IEEE, 2017.
- [11] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 259–269. IEEE, 2000.
- [12] T. Batu, E. Fischer, Kumar Fortnow, L., R. R., Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proceedings 2001 IEEE International Conference on Cluster Computing*, pages 442–451. IEEE, 2001.
- [13] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proceedings of 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451. IEEE, 2001.
- [14] D. Braess and T. Sauer. Bernstein polynomials and learning theory. *Journal of Approximation Theory*, 128(2):187–206, 2004.
- [15] G. Bresler. Efficiently learning ising models on arbitrary graphs. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, pages 771–782. ACM, 2015.
- [16] C. L. Canonne. A survey on distribution testing. *Your Data is Big. But is it Blue.*, 2017.
- [17] A. Carlton. On the bias of information estimates. *Psychological Bulletin*, 71(2):108, 1969.
- [18] A. G. Carlton. On the bias of information estimates. *Psychological Bulletin*, 71(2):108, 1969.
- [19] S. O. Chan, I. Diakonikolas, P. Valiant, and G. Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1203. Society for Industrial and Applied Mathematics, 2014.
- [20] A. Chao. Nonparametric estimation of the number of classes in a population. *Scandinavian Journal of Statistics*, pages 265–270, 1984.
- [21] A. Chao and C. H. Chiu. Species richness: estimation and comparison. *Wiley StatsRef: Statistics Reference Online*, pages 1–26, 2014.
- [22] A. Chao and S. M. Lee. Estimating the number of classes via sample coverage. *Journal of the American Statistical Association*, 87(417):210–217, 1992.
- [23] M. S. Charikar, K. Shiragur, and A. Sidford. Efficient profile maximum likelihood for universal symmetric property estimation. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 780–791. ACM, 2019.
- [24] S. F. Chen and J. Goodman. An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13(4):359–394, 1999.
- [25] H. Chernoff. A note on an inequality involving the normal distribution. *The Annals of Probability*, pages 533–535, 1981.
- [26] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.
- [27] R. K. Colwell, A. Chao, N. J. Gotelli, S. Y. Lin, C. X. Mao, R. L. Chazdon, and J. T. Longino. Models and estimators linking individual-based and sample-based rarefaction, extrapolation and comparison of assemblages. *Journal of Plant Ecology*, 5(1):3–21, 2012.
- [28] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

- [29] H. Das. Competitive tests and estimators for properties of distributions. (*Doctoral dissertation, UC San Diego*), 2012.
- [30] I. Diakonikolas and D. M. Kane. A new approach for testing properties of discrete distributions. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 685–694. IEEE, 2016.
- [31] 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 (SODA)*, pages 1841–1854. Society for Industrial and Applied Mathematics, 2015.
- [32] 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.
- [33] B. Efron and R. Thisted. Estimating the number of unseen species: How many words did Shakespeare know? *Biometrika*, 63(3):435–447, 1976.
- [34] W. Gerstner and W. M. Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge University Press, 2002.
- [35] O. Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, page 1. IEEE, 2016.
- [36] O. Goldreich. *Introduction to property testing (chapter 11)*. Cambridge University Press, 2017.
- [37] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. In *Technical Report TR00-020, Electronic Colloquium on Computational Complexity (ECCC)*, 2000.
- [38] O. Goldreich and D. Ron. On the relation between the relative earth mover distance and the variation distance (an exposition). *Available Online*, 2016.
- [39] I. J. Good. The population frequencies of species and the estimation of population parameters. *Biometrika*, 40(3-4):237–264, 1953.
- [40] I. J. Good and G. H. Toulmin. The number of new species, and the increase in population coverage, when a sample is increased. *Biometrika*, 43(1-2):45–63, 1956.
- [41] P. Grassberger. Finite sample corrections to entropy and dimension estimates. *Physics Letters A*, 128(6-7):369–373, 1988.
- [42] P. J. Haas, J. F. Naughton, S. Seshadri, and L. Stokes. Sampling-based estimation of the number of distinct values of an attribute. *VLDB*, 95:311–322, 1995.
- [43] Y. Han, J. Jiao, and T. Weissman. Local moment matching: A unified methodology for symmetric functional estimation and distribution estimation under Wasserstein distance. In *Conference on Learning Theory*, pages 3189–3221, 2018.
- [44] Y. Hao and A. Orlitsky. Adaptive estimation of generalized distance to uniformity. In *2018 IEEE International Symposium on Information Theory*, pages 1076–1080. IEEE, 2018.
- [45] Y. Hao and A. Orlitsky. Data amplification: Instance-optimal property estimation. In *arXiv preprint arXiv:1903.01432*, 2019.
- [46] Y. Hao and A. Orlitsky. Doubly-competitive distribution estimation. In *International Conference on Machine Learning*, pages 2614–2623, 2019.
- [47] Y. Hao and A. Orlitsky. Unified sample-optimal property estimation in near-linear time. In *Advances in Neural Information Processing Systems*, 2019.
- [48] Y. Hao, A. Orlitsky, A. T. Suresh, and Y. Wu. Data amplification: A unified and competitive approach to property estimation. In *Advances in Neural Information Processing Systems*, pages 8848–8857, 2018.

- [49] G. H. Hardy and S. Ramanujan. Asymptotic formulae in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.
- [50] R. Jenssen, K. E. Hild, D. Erdogmus, J. C. Principe, and T. Eltoft. Clustering using Rényi’s entropy. In *Proceedings of the International Joint Conference on Neural Networks*, volume 1, pages 523–528. IEEE, 2003.
- [51] J. Jiao, K. Venkat, Y. Han, and T. Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015.
- [52] D. Källberg, N. Leonenko, and O. Seleznev. Statistical inference for Rényi entropy functionals. In *Conceptual Modelling and Its Theoretical Foundations*. Springer, Berlin, Heidelberg., pages 36–51, 2012.
- [53] S. Kamath, A. Orlitsky, D. Pichapati, and A. T. Suresh. On learning distributions from their samples. In *Conference on Learning Theory*, pages 1066–1100, 2015.
- [54] L. V. Kantorovich and G. S. Rubinstein. On a space of completely additive functions. *Vestnik Leningrad. Univ*, 13(7), 1958.
- [55] R. Krichevsky and V. Trofimov. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207, 1981.
- [56] I. Kroes, P. W. Lepp, and D. A. Relman. Bacterial diversity within the human subgingival crevice. *Proceedings of the National Academy of Sciences*, 96(25):14547–14552, 1999.
- [57] B. Ma, A. Hero, J. Gorman, and O. Michel. Image registration with minimum spanning tree algorithm. In *Proceedings 2000 International Conference on Image Processing (Cat. No. 00CH37101)*, volume 1, pages 481–484. IEEE, 2000.
- [58] Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- [59] C. X. Mao and B. G. Lindsay. Estimating the number of classes. *The Annals of Statistics*, pages 917–930, 2007.
- [60] C. McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1): 148–188, 1989.
- [61] D. R. McNeil. Estimating an author’s vocabulary. *Journal of the American Statistical Association*, 68(341):92–96, 1973.
- [62] H. Neemuchwala, A. Hero, S. Zabuawala, and P. Carson. Image registration methods in high-dimensional space. *International Journal of Imaging Systems and Technology*, 16(5):130–145, 2006.
- [63] M. Obremski and M. Skorski. Rényi entropy estimation revisited. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [64] A. Orlitsky and A. T. Suresh. Competitive distribution estimation: Why is Good-Turing good. In *Advances in Neural Information Processing Systems*, pages 2143–2151, 2015.
- [65] A. Orlitsky, S. Sajama, N. P. Santhanam, K. Viswanathan, and J. Zhang. Algorithms for modeling distributions over large alphabets. In *Proceedings of 2004 IEEE International Symposium on Information Theory (ISIT)*, pages 304–304. IEEE, 2004.
- [66] A. Orlitsky, N. P. Santhanam, K. Viswanathan, and J. Zhang. On modeling profiles instead of values. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 426–435. AUAI Press, 2004.
- [67] A. Orlitsky, N. P. Santhanam, K. Viswanathan, and J. Zhang. On estimating the probability multiset. *Online Draft*, 2011. URL <http://alou.ucsd.edu/papers/pm11.pdf>.

- [68] A. Orlitsky, A. T. Suresh, and Y. Wu. Optimal prediction of the number of unseen species. *Proceedings of the National Academy of Sciences*, 113(47):13283–13288, 2016.
- [69] S. Pan. On the theory and application of pattern maximum likelihood. (*Doctoral dissertation, UC San Diego*), 2012.
- [70] L. Paninski. Estimation of entropy and mutual information. *Neural Computation*, 15(6):1191–1253, 2003.
- [71] L. Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- [72] D. S. Pavlichin, J. Jiao, and T. Weissman. Approximate profile maximum likelihood. *arXiv preprint, arXiv:1712.07177*, 2017.
- [73] C. J. Quinn, N. Kiyavash, and T. P. Coleman. Efficient methods to compute optimal tree approximations of directed information graphs. *IEEE Transactions on Signal Processing*, 61(12):3173–3182, 2013.
- [74] A. Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1961.
- [75] D. Ron. *Algorithmic and analysis techniques in property testing*. Number 5 in 2. Foundations and Trends in Theoretical Computer Science, 2010.
- [76] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [77] R. Thisted and B. Efron. Did Shakespeare write a newly-discovered poem? *Biometrika*, 74(3):445–455, 1987.
- [78] G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.
- [79] 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 43rd Annual ACM Symposium on Theory of Computing*, pages 685–694. ACM, 2011.
- [80] G. Valiant and P. Valiant. The power of linear estimators. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 403–412. IEEE, 2011.
- [81] G. Valiant and P. Valiant. Instance optimal learning of discrete distributions. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 142–155. ACM, 2016.
- [82] G. Valiant and P. Valiant. Estimating the unseen: Improved estimators for entropy and other properties. *Journal of the ACM (JACM)*, 64(6):37, 2017.
- [83] G. J. Valiant. Algorithmic approaches to statistical questions. (*Doctoral dissertation, UC Berkeley*), 2012.
- [84] R. K. Vinayak, W. Kong, G. Valiant, and S. M. Kakade. Maximum likelihood estimation for learning populations of parameters. *arXiv preprint, arXiv:1902.04553*, 2019.
- [85] P. O. Vontobel. The Bethe approximation of the pattern maximum likelihood distribution. In *Proceedings of 2012 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2012.
- [86] P. O. Vontobel. The Bethe and Sinkhorn approximations of the pattern maximum likelihood estimate and their connections to the Valiant-Valiant estimate. In *2014 Information Theory and Applications Workshop (ITA)*, pages 1–10. IEEE, 2014.
- [87] Y. Wu and P. Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, 2016.

- [88] Y. Wu and P. Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *The Annals of Statistics*, 47(2):857–883, 2019.
- [89] D. Xu. Energy, entropy and information potential for neural computation. *Doctoral dissertation, University of Florida*, 1999.
- [90] D. Xu and D. Erdogmus. Rényi’s entropy, divergence and their nonparametric estimators. In *Information Theoretic Learning*. Springer, New York, NY., pages 47–102, 2010.
- [91] P. Yang. Optimal entropy estimation on large alphabet: Fundamental limits and fast algorithms. (*Doctoral dissertation*), 2016.