arXiv:2002.11665v1 [stat.ML] 26 Feb 2020

Profile Entropy: A Fundamental Measure for the Learnability and Compressibility of Discrete Distributions

Yi Hao
Dept. of Electrical and Computer Engineering
University of California, San Diego
yih179@ucsd.edu

Alon Orlitsky
Dept. of Electrical and Computer Engineering
University of California, San Diego
alon@ucsd.edu

Abstract

The profile of a sample is the multiset of its symbol frequencies. We show that for samples of discrete distributions, profile entropy is a fundamental measure unifying the concepts of estimation, inference, and compression. Specifically, profile entropy a) determines the speed of estimating the distribution relative to the best natural estimator; b) characterizes the rate of inferring all symmetric properties compared with the best estimator over any label-invariant distribution collection; c) serves as the limit of profile compression, for which we derive optimal near-linear-time block and sequential algorithms. To further our understanding of profile entropy, we investigate its attributes, provide algorithms for approximating its value, and determine its magnitude for numerous structural distribution families.

Contents

1 Introduction **2**

1.1 Sample Profiles and Their Entropy 3

1.2 Applications and Prior Works 3

2 New Results **5**

2.1 Profile Dimension and Entropy 6

2.2 Adaptive Property Estimation 6

2.3 Competitive Distribution Estimation 7

2.4 Optimal Profile Compression 8

2.5 Attributes of Profile Entropy and Dimension 9

2.6 Profile Entropy for Structured Families 11

3 Extension and Conclusion **14**

3.1 Multi-Dimensional Profiles 14

3.2 Concluding Remarks 14

Preprint. Under review.

A Dimension and Entropy of Sample Profiles **15**
A.1 Profile Dimension and Its Concentration 15
A.2 Profile Entropy and Its Connection to Dimension 18
A.3 Symmetric Property Estimation and Sufficiency of Profiles 28
A.4 Median Trick 29
A.5 Profile Maximum Likelihood and Its Adaptiveness 29
A.6 PML and Sorted Distribution Estimation 32
A.7 PML and Uniformity Testing 34

B Attributes of Profile Entropy and Dimension **35**
B.1 Monotonicity 35
B.2 Lipschitzness 36
B.3 Log-Concave Distributions 38
B.4 Discretization of Continuous Distributions 39
B.5 Optimality of Theorem 13 40
B.6 Power-Law Distributions 42
B.7 Histogram Distributions 44

C Competitive Estimation of Distributions and Their Entropy **45**
C.1 Competitive Distribution Estimation 45
C.2 Competitive Entropy Estimation 47

D Optimal Profile Compression **48**
D.1 Block Compression 49
D.2 Sequential Compression 49

E Extensions and Additional Results **49**
E.1 Multi-Dimensional Profiles 49
E.2 Discrete Multi-Variate Gaussian 50

1 Introduction

Recent research in statistical machine learning, ranging from neural-network training and online learning, to density estimation and property testing, has advanced evaluation criteria beyond worst-case analysis. New performance measures apply more refined metrics relating the algorithm's accuracy and efficiency to the problem's inherent structure.

Consider for example learning an unknown discrete distribution from its i.i.d. samples. Classical worst-case analysis states that in the worst case, the number of samples required to estimate a distribution to a given KL-divergence grows linearly in the alphabet size.

However, this formulation is pessimistic. Distributions are rarely the worst possible, and many practical distributions can be estimated with significantly smaller samples. Furthermore, once the sample is drawn, it reveals the distribution's complexity and hence the hardness of the learning task.

Going beyond worst-case analysis, we design an *adaptive* learning algorithm whose theoretical guarantees vary according to the problem's simplicity. For example, Orlitsky and Suresh [2015] recently

proposed an estimator that instance-by-instance achieves nearly the same performance as a genie algorithm designed with prior knowledge of the underlying distribution.

We introduce *profile entropy*, a fundamental measure for the complexity of discrete distributions, and show that it connects three vital scientific tasks: estimation, inference, and compression. The resulting algorithms have guarantees directly relating to the data profile entropy, hence also adapt to the intrinsic simplicity of the tasks at hand.

The next subsections, formalize the relevant concepts and present relevant prior works.

1.1 Sample Profiles and Their Entropy

Consider an arbitrary sequence x^n over a finite or countably infinite alphabet $\mathcal{X}$. The *multiplicity* $\mu_y(x^n)$ of a symbol $y \in \mathcal{X}$ is the number of times y appears in x^n. The *prevalence* of an integer μ is the number $\varphi_\mu(x^n)$ of symbols in x^n with multiplicity μ. The *profile* of x^n is the multiset $\varphi(x^n)$ of multiplicities of the symbols in x^n. We refer to it as a profile of *length* n.

The number $\mathcal{D}(S)$ of distinct elements in a multiset S is its *dimension*. For convenience, we also write $\mathcal{D}(x^n)$ for profile dimension. Note that the dimension of a length-n profile is at most $\min\{\sqrt{2n}, |\mathcal{X}|\}$.

Let $\Delta_\mathcal{X}$ be the collection of distributions over $\mathcal{X}$, and p be an arbitrary distribution in $\Delta_\mathcal{X}$. The profile Φ^n of an i.i.d. sample $X^n \sim p$ is a random variable whose distribution depends on only p and n. We therefore write $\Phi^n \sim p$, and call $H(\Phi^n)$ the *profile entropy* with respect to (p, n). Analogously, we call $\mathcal{D}_n := \mathcal{D}(\Phi^n)$, the *profile dimension* associated with (p, n) and write $\mathcal{D}_n \sim p$.

Due to the dependence among multiplicities, the distributions of Φ^n and $\mathcal{D}_n$ are rather complex in general. To obtain clean expressions, we can adopt the standard *Poisson sampling* technique and make the sample size a Poisson variable $N \sim \mathrm{Poi}(n)$, independent of the sample. As an example,

$$\mathbb{E}\left[\mathcal{D}_N \sim p\right] = \sum_{i=1}^{\infty} \left(1 - \prod_{x \in \mathcal{X}} \left(1 - e^{-np_x} \frac{(np_x)^i}{i!}\right)\right),$$

where p_x denotes the probability of symbol x assigned by p. Note that sometimes we also write $p(x)$ instead of p_x for notational convenience. Despite the complex landscape of statistical dependency, in Theorem 1, we show that $\mathcal{D}_n \sim p$ and $H(\Phi^n \sim p)$ are of the same order, with high probability and for every $p \in \Delta_\mathcal{X}$. In Theorem 6 and 7, we show that $\mathcal{D}_n \sim p$ highly concentrates around a variant of its expectation. In Section 2.5, we provide a much simpler quantity $H_n^\mathcal{S}(p)$ that well approximates the expectation variant of $\mathcal{D}_n \sim p$. Combined, these three results provide a precise characterization of both profile entropy and dimension. Leveraging this in Section B.2, we derive nearly-tight bounds on the magnitude of profile entropy for several important structural distribution families, including log-concave and power-law.

1.2 Applications and Prior Works

In this section, we present several learning and compression applications in which the profile entropy would play an important role. We also review related prior works with an emphasize on adaptive algorithms, and suppress the discussions on the worst-case analysis for brevity.

Basics and Significance

The profile of a sample corresponds to the empirical distribution of symbols, and reflects the magnitudes of the actual symbol probabilities. Hence, the profile dimension, the number of distinct symbol frequencies, characterizes the variability of ranges the probabilities spread over. The sample profile's entropy, which by Theorem 1 is of the same order as its dimension, admits the same interpretation.

Intuitively, samples from simple distributions tend to have low profile entropy, such as those from a m-piecewise distribution with small m, or one whose probability masses concentrate over some sparse set. The profile entropy is also likely to decrease as one reduces the sample size, since the sample contains less information regarding the variability of distribution probabilities. See Theorem 9, 14, and 19 for a formal justification of these arguments.

From a statistical perspective, the profile of a sample is a sufficient statistic for estimating the probability multiset and any symmetric functional of the underlying distribution, such as entropy and support size. When we express a profile as a collection of multiplicity-prevalence pairs, the profile dimension is the size of this collection. Being of the same magnitude, the profile entropy is thus the effective size of a natural sufficient statistic for label-invariant inference.

Profile entropy also directly connects to adaptive testing and classification. Such connection arises from computing the *profile probability* Acharya et al. [2011, 2012], the probability of observing a sample with the given profile. If the profile has entropy of $\mathcal{H}$, we can show that this computation problem has a time complexity of $\mathcal{O}(\exp(\tilde{\Theta}(\mathcal{H})\log|\mathcal{X}|))$. The result follows by the equivalence of the problem and computing the permanent of a rank-$\tilde{\Theta}(\mathcal{H})$ matrix Barvinok [1996], Vontobel [2012, 2014], Barvinok [2016].

Below, we introduce two important applications that are more involved, in which the profile entropy has essential connection to the statistical efficacy of adaptive learning.

Distribution Estimation

Estimating unknown distributions from their samples is a statistical-inference cornerstone, and has numerous applications, ranging from biological studies Armañanzas et al. [2008] to language modeling Chen and Goodman [1999].

A learning algorithm in this setting is often referred to as a *distribution estimator*, which is a functional $\hat{p}$ associating with every sequence x^n over $\mathcal{X}$ a distribution $\hat{p}_{x^n} \in \Delta_{\mathcal{X}}$. Given a sample $X^n \sim p$, we measure the performance of $\hat{p}$ in estimating the (unknown) distribution p with a loss function $\ell(p, \hat{p}_{X^n})$, e.g., the ℓ_1 distance and KL divergence.

A classical worst-case type result shows that any estimator that achieves a small ℓ_1 loss of $\varepsilon > 0$ over $\Delta_{\mathcal{X}}$ in expectation requires a sample size of $\Omega(|\mathcal{X}|/\varepsilon^2)$. Recent research further shows that the naive empirical-distribution estimator attains the optimal sample efficiency, to the right constants.

The desire to design more efficient estimators for practical distributions such as Poisson mixtures leads to two adaptive estimation frameworks: structural and competitive.

Structural estimation focuses on distributions possessing a natural structure, such as monotonicity, m-modality, and log-concavity. In many cases including the mentioned, structural assumptions lead to effective estimators that provably perform better on the corresponding distribution classes. See Bühlmann et al. [2016] for a review of recent literature.

Competitive estimation aims to design estimators that are universally near-optimal. Without strong structural knowledge, a reasonable estimator should *naturally* assign the same probability to symbols appearing equal number of times. The objective here is to find an estimator that learns *every* distribution as well as the best natural estimator designed with knowledge of the true distribution. Discussion continues in Section 2.3 with a review of relevant works.

Property (Functional) Inference

Instead of recovering the underlying distribution, numerous practical applications require only inferring a particular *property value*, such as entropy for graphical modeling Koller and Friedman [2009], and support size for species richness estimation Magurran [2013].

Formally, a *distribution property* over a distribution collection $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$ is a functional $f : \mathcal{P} \to \mathbb{R}$ that associates with each distribution in $\mathcal{P}$ a real value. Given a sample X^n from an unknown distribution $p \in \mathcal{P}$, the problem of interest is to infer the value of $f(p)$. To do this, we employee another functional $\hat{f} : \mathcal{X}^* \to \mathbb{R}$, a *property estimator* that maps every sample to a real value.

The statistical efficiency of $\hat{f}$ in estimating f with respect to the distribution collection $\mathcal{P}$ is measured by its *sample complexity*. Specifically, for an accuracy $\varepsilon > 0$ and error tolerance $\delta \in (0, 1)$, the (ε, δ)-sample complexity of $\hat{f}$ with respect to $(f, \mathcal{P})$ is the minimal sample size n for which $\Pr_{X^n \sim p}(|\hat{f}(X^n) - f(p)| > \varepsilon) \leq \delta$ for all $p \in \mathcal{P}$. Note that for the special case of $\mathcal{P} = \{p\}$, the sample complexity directly characterizes the ability of $\hat{f}$ in estimating $f(p)$.

Recent years have shown interests in determining the sample complexities of inferring distribution properties. Built upon worst-case analysis, the major contribution of these works is establishing the sufficiency of sample sizes sub-linear in $|\mathcal{X}|$. As an example, in the vital sample-sparse regime and over $\Delta_{\mathcal{X}}$, the $(\varepsilon, 1/10)$-sample complexity of learning entropy is $\Theta(|\mathcal{X}|/(\varepsilon \log |\mathcal{X}|))$. We refer the readers to Verdú [2019] for a thorough survey of related works.

As the problem involves two components, the property and distribution, adaptive analysis also advances in two veins.

The first vein concerns constructing a universal plug-in estimator for all *symmetric properties*. A symmetric property is invariant under symbol permutations, hence it suffices to obtain an accurate estimate of the probability multiset. Recently, following the works of Das [2012], Acharya et al. [2017], Hao and Orlitsky [2019a] show that for any symmetric property that is additively separable and appropriately Lipschitz, the profile maximum likelihood (Section 2.2) achieves the optimal sample complexity up to small constant factors. Other major works include Valiant and Valiant [2011, 2013, 2016], Han et al. [2018], Charikar et al. [2019b]

The second vein is an analogy to the competitive distribution estimation framework, and aims to compete with the *instance-by-instance* performance of a genie having access to more information, but reasonably restricted. A natural choice for the genie is the best-known and most-used – the *empirical estimator* that evaluates the property at the sample empirical distribution. To empower the genie, we grant it access to a sample whose size is logarithmically larger than that available to the learner. One can show that this enables the genie to universally achieve the optimal sample complexities for numerous properties and hypothesis classes $\mathcal{P}$. Under this formulation, Hao et al. [2018], Hao and Orlitsky [2019b] provide a unified learning algorithm that achieves the optimal competitiveness guarantees in near-linear time.

In this work, we further both veins of works and show that: 1) the PML plug-in estimator possesses the amazing ability of adapting to the simplicity of data distributions in inferring all symmetric properties, over any label-invariant classes; 2) when plugged into entropy, the estimator in Hao and Orlitsky [2019c] approximates the property as well as the plug-in estimator whose distribution component is the best natural, for every distribution. See Theorem 3 and 4 for the formal statements.

2 New Results

We establish essential connections between profile entropy and the estimation of distributions, inference of their properties, and compression of profiles. To further our understanding of profile entropy, we then investigate its attributes, provide algorithms for approximating its value, and determine its magnitude for numerous structural distribution families.

For space considerations, we relegate most technical proofs to the appendices.

Permutation invariance By definition, both the profile of a sequence and its dimension are invariant to domain-symbol permutations. Since entropy is a symmetric property, the profile entropy of an i.i.d. sample is also permutation invariant. Consequently, a result in this section that holds for a distribution will also hold for *any distributions sharing the same probability multiset*.

This is desirable for practical applications, since samples often come as categorical data, while the symbol ordering under which the underlying distribution would exhibit certain structure is unknown to the learner. For example, in natural language processing, we observe words and punctuation marks. Given that the data comes from a power-law distribution Mitzenmacher [2004], we often don't know how to order the alphabet to realize such a condition.

Surprisingly, with a few exceptions such as Hao and Orlitsky [2019c], most previous works on learning structured discrete distributions do not address this crucial matter in their learning algorithms. The existing results are rather artificial and more like learning discretized continuous distributions. See Section B.2 for our discussion on distribution discretization.

2.1 Profile Dimension and Entropy

Denote by $\lceil x \rceil$ the smallest integer larger than x. Then,

Theorem 1. *For any distribution $p \in \Delta_{\mathcal{X}}$ and $\Phi^n \sim p$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,*

$$\lceil H(\Phi^n) \rceil = \tilde{\Theta}(\mathcal{D}(\Phi^n)),$$

where the notation $\tilde{\Theta}(\cdot)$ hides logarithmic factors of n.

The theorem shows that for every distribution and sampling parameter n, the induced profile entropy and profile dimension are of the same order, with high probability.

Taking expectation and noting that $\mathcal{D}(\Phi^n) \in [1, \sqrt{2n}]$ yield

$$\lceil H(\Phi^n \sim p) \rceil = \tilde{\Theta}(\mathop{\mathbb{E}}_{\Phi^n \sim p}[\mathcal{D}(\Phi^n)]), \ \forall p \in \Delta_{\mathcal{X}}.$$

2.2 Adaptive Property Estimation

Definitions A profile $\boldsymbol{\phi}$ is said to have length n if there exists $x^n \in \mathcal{X}^n$ satisfying $\boldsymbol{\phi} = \varphi(x^n)$. For every profile $\boldsymbol{\phi}$ of length n and distribution collection $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$, the *profile maximum likelihood* (PML) estimator Orlitsky et al. [2004] over $\mathcal{P}$ maps $\boldsymbol{\phi}$ to a distribution

$$\mathcal{P}_{\boldsymbol{\phi}} := \mathop{\arg\min}_{p \in \mathcal{P}} \Pr_{X^n \sim p}(\varphi(X^n) = \boldsymbol{\phi}),$$

that maximizes the probability of observing the profile $\boldsymbol{\phi}$. For any property f, let $\varepsilon_f(n, \delta, \mathcal{P})$ denote the smallest error that can be achieved by any estimator with a sample size n and tolerance δ on the error probability. This definition is equivalent to that of the sample complexity. Below, we assume that $\mathcal{P}$ is *label invariant*, i.e., for any $p \in \mathcal{P}$, collection $\mathcal{P}$ contains all its symbol-permuted versions.

We first show that profile-based estimators are sufficient for estimating symmetric properties.

Theorem 2 (Sufficiency of profiles)**.** *Let f be a symmetric property over $\mathcal{P}$. For any accuracy $\varepsilon > 0$ and tolerance $\delta \in (0, 1)$, if there exists an estimator $\hat{f}$ such that*

$$\Pr_{X^n \sim p}\left(\left|\hat{f}(X^n) - f(p)\right| > \varepsilon\right) < \delta, \ \forall p \in \mathcal{P},$$

there is an estimator $\hat{f}_\varphi$ over length-n profiles satisfying

$$\Pr_{X^n \sim p}\left(\left|\hat{f}_\varphi(\varphi(X^n)) - f(p)\right| > \varepsilon\right) < \delta, \ \forall p \in \mathcal{P}.$$

Note that both estimators can have independent randomness.

The second result shows that the PML estimator is adaptive to the simplicity of underlying distributions in inferring all symmetric properties, over any label-invariant $\mathcal{P}$. For clarity, we set $\delta = 1/10$ and suppress both δ and $\mathcal{P}$ in $\varepsilon_f(n, \delta, \mathcal{P})$.

Theorem 3 (Adaptiveness of PML)**.** *Let f be a symmetric property. For any $p \in \mathcal{P}$ and $\Phi^n \sim p$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,*

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\varepsilon_f\left(\frac{\tilde{\Omega}(n)}{\lceil H(\Phi^n) \rceil}\right).$$

Some comments: 1) The theorem holds for any symmetric properties, while nearly all previous works require the property to possess certain forms and be smooth; 2) The theorem trivially implies a weaker result in Acharya et al. [2017] where $\lceil H(\Phi^n) \rceil$ is replaced by $\sqrt{n}$; 3) There is a polynomial-time approximation Charikar et al. [2019a] achieving the same guarantee; 4) We provide a stronger result in Section A.5 of the appendices for general δ.

Besides this theorem, we establish in Section A.6 and A.7 two additional results on PML. The first result addresses sorted distribution estimation, and improves over that established in Hao and Orlitsky [2019a] (Theorem 5) in terms of the lower bound on the accuracy parameter ε. Let $\Sigma_{\mathcal{X}}$ denote the collection of symbol permutations over $\mathcal{X}$. For any $p \in \Delta_{\mathcal{X}}$, denote by $p_\sigma \in \Delta_{\mathcal{X}}$ the permuted distribution satisfying $p_\sigma(x) = p(\sigma(x))$ for all symbols $x \in \mathcal{X}$. Let $\lambda > 0$ be a positive absolute constant that can be made arbitrarily small, e.g., $\lambda = 0.001$.

Lemma 1. *For any $\varepsilon \in (0,1)$, $p \in \mathcal{P} = \Delta_{\mathcal{X}}$, and $\Phi^n \sim p$, if we have $n \geq \Omega(|\mathcal{X}|/(\varepsilon^2 \log |\mathcal{X}|))$ and $\varepsilon \geq 1/n^{1/8-\lambda}$, with probability at least $1 - \mathcal{O}(\exp(-\sqrt{n}))$,*

$$\min_{\sigma \in \Sigma_{\mathcal{X}}} \|p_\sigma - \mathcal{P}_{\Phi^n}\|_1 \leq \mathcal{O}(\epsilon).$$

A few comments in order: 1) The polynomial-time computable variant of PML in Charikar et al. [2019a] satisfies the same guarantee, and the proof for this is also similar to that in Section A.6; 2) Using the existing efficiently computable PML-type methods Charikar et al. [2019a,b], the best possible lower bound on ε is $\Theta(1/n^{1/4})$; 3) Below the $\Theta(1/n^{1/3})$ threshold, the empirical distribution estimator is sample optimal up to constant factors Han et al. [2018].

The second result shows an intriguing connection between the PML method and the task of uniformity testing Goldreich and Ron [2011]. See Section A.7 for details.

Our last result in this section addresses entropy estimation. We show that when plugged into entropy, the estimator in Hao and Orlitsky [2019c] approximates the property as well as the plug-in estimator whose distribution component is the best natural, for every distribution.

Recall that a distribution estimator is natural if it assigns the same probability to symbols of equal multiplicity, and a property estimator is plug-in if it first finds an estimate of the distribution and then evaluates the property at this estimate. As an off-the-shelf method, the plug-in approach is widely used in estimating distribution properties.

If further the property is symmetric, then it suffices to obtain an accurate estimate of the probability multiset, which is intuitively more statistically efficient than recovering the actual distribution. For example, Hao and Orlitsky [2019a] recently show that for any symmetric property that is additively separable and appropriately Lipschitz, the PML multiset estimator Orlitsky et al. [2004] achieves the optimal sample complexity up to small constant factors.

However, the analysis and computation (though efficient) of such multiset-based estimation methods are often involved Valiant and Valiant [2011, 2013, 2016], Han et al. [2018], Charikar et al. [2019b], Hao and Orlitsky [2019a]. For this reason, distribution-based plug-in estimators are still popular in practice, and often, the distribution components are natural.

As an example, for entropy estimation, several widely used distribution-based estimators are *natural plug-in*, such as the empirical estimator plugging in the empirical distribution, James-Stein shrinkage Hausser and Strimmer [2009] that shrinks the distribution estimate towards uniform, and Dirichlet-smoothed Schürmann and Grassberger [1996] that imposes a Dirichlet prior over $\Delta_{\mathcal{X}}$.

The logic behind these estimators is simple: if two distributions are close, then the same is expected to hold for their entropy values. The next theorm shows that for *every* distribution and among all plug-in entropy estimators, the distribution estimator in Hao and Orlitsky [2019c] is as good as the one that performs best in estimating the actual distribution.

Denote by $\mathcal{N}$ the collection of all natural estimators. Write $|H(p) - H(q)|$ as $\ell_H(p,q)$ for compactness and the KL-divergence between $p, q \in \Delta_{\mathcal{X}}$ as $\ell_{\text{KL}}(p,q)$.

Theorem 4 (Competitive entropy estimation)**.** *For any distribution p, sample $X^n \sim p$ with profile $\Phi^n := \varphi(X^n)$, and $\hat{p}^{\mathcal{N}}_{X^n} := \arg\min_{\hat{p} \in \mathcal{N}} \ell_{KL}(p, \hat{p}_{X^n})$, we have*

$$\ell_H(p, \hat{p}^{\star}_{X^n}) - \ell_H(p, \hat{p}^{\mathcal{N}}_{X^n}) \leq \tilde{\mathcal{O}}\left(\sqrt{\frac{\lceil H(\Phi^n) \rceil}{n}}\right).$$

with probability at least $1 - \mathcal{O}(1/n)$.

2.3 Competitive Distribution Estimation

Prior works Competitive estimation calls for an estimator that competes with the instance-by-instance performance of a genie knowing more information, but reasonably restricted. Denote by $\ell_{\text{KL}}(p,q)$ the KL divergence. Introduced in Orlitsky and Suresh [2015], the formulation considers the collection $\mathcal{N}$ of all natural estimators, and shows that a simple variant $\hat{p}^{\text{GT}}$ of the Good-Turing estimator achieves

$$\ell_{\text{KL}}(p, \hat{p}^{\text{GT}}_{X^n}) - \min_{\hat{p} \in \mathcal{N}} \ell_{\text{KL}}(p, \hat{p}_{X^n}) \leq \frac{3 + o(1)}{n^{1/3}},$$

for every distribution p and with high probability. We refer to the left-hand side as the *excess loss* of estimator $\hat{p}_{\text{GT}}$ with respect to the best natural estimator, and note that it vanishes at a rate independent of p. For a more involved estimator in Acharya et al. [2013], the excess loss vanishes at a faster rate of $\tilde{\mathcal{O}}(\min\{1/\sqrt{n}, |\mathcal{X}|/n\})$, optimal up to logarithmic factors for every estimator and the respective worst-case distribution. For the ℓ_1 distance, Valiant and Valiant [2016] derive a similar result.

These estimators track the loss of the best natural estimator for each distribution. Yet an equally important component, the excess loss bound, is still of the worst-case nature. For a fully adaptive guarantee, Hao and Orlitsky [2019c] design an estimator $\hat{p}^\star$ that achieves a $\mathcal{D}_n/n$ excess loss, i.e.,

$$\ell_{\text{KL}}(p, \hat{p}^\star_{X^n}) - \min_{\hat{p}\in\mathcal{N}} \ell_{\text{KL}}(p, \hat{p}_{X^n}) \le \tilde{\mathcal{O}}\left(\frac{\mathcal{D}_n}{n}\right),$$

for every p and $X^n \sim p$, with high probability. Utilizing the adaptiveness of $\mathcal{D}_n$ to the simplicity of distributions, the paper derives excess-loss bounds for several important distribution families, and proves the estimator's optimality under various of classical and modern learning frameworks.

New results While the work of Hao and Orlitsky [2019c] provides an appealing upper bound on the excess loss, it is not exactly clear how good this bound is as a matching lower bound is missing. In this work, we complete the picture by showing that the $\mathcal{D}_n/n$ bound is essential for competitive estimation and optimal up to logarithmic factors of n.

Theorem 5 (Minimal excess loss)**.** *For any $n, \mathcal{D} \in \mathbb{N}$ and distribution estimator $\hat{p}'$, there is a distribution p such that with probability at least $9/10$, we have both*

$$\mathcal{O}(\log n + \mathcal{D}) \ge \mathcal{D}_n$$

and

$$\ell_{KL}(p, \hat{p}'_{X^n}) - \min_{\hat{p}\in\mathcal{N}} \ell_{KL}(p, \hat{p}_{X^n}) \ge \Omega\left(\frac{\mathcal{D}}{n}\right).$$

By Theorem 1, we can replace $\mathcal{D}_n$ by $\tilde{\Theta}(H(\Phi^n))$ in both the upper and lower bounds.

2.4 Optimal Profile Compression

While a labeled sample contains all information, for many modern applications, such as property estimation and differential privacy, it is sufficient Orlitsky et al. [2004] or even necessary to provide only the profile Suresh [2019]. Hence, this section focuses on the lossless compression of profiles.

For any distribution p, it is well-known that the minimal expected codeword length (MECL) for losslessly compressing a sample $X^n \sim p$ is approximately $nH(p)$, which increases linearly in n as long as $H(p)$ is bounded away from zero.

On the other hand, by the Hardy-Ramanujan formula Hardy and Ramanujan [1918], the number $\mathbb{P}(n)$ of integer partitions of n, which happens to equal to the number of length-n profiles, satisfies

$$\log \mathbb{P}(n) = 2\pi\sqrt{\frac{n}{6}}(1 + o(1)).$$

Consequently, the MECL for losslessly compressing the sample profile $\Phi^n \sim p$ is at most $\mathcal{O}(\sqrt{n})$, a number potentially much smaller than $nH(p)$.

By Shannon's source coding theorem, the profile entropy $H(\Phi^n)$ is the *information-theoretic limit* of MECL for the lossless compression of profile $\Phi^n \sim p$. Below, we present explicit block and sequential profile compression schemes achieving this entropy limit, up to logarithmic factors of n.

Block compression The block compression algorithm is intuitive and easy to implement.

Recall that the profile of a sequence x^n is the multiset $\varphi(x^n)$ of multiplicities associated with symbols in x^n. The ordering of elements in a multiset is not informative. Hence equivalently, we can compress $\varphi(x^n)$ into the set $\mathcal{C}(\varphi(x^n))$ of corresponding multiplicity-prevalence pairs, i.e.,

$$\mathcal{C}(\varphi(x^n)) := \{(\mu, \varphi_\mu(x^n)) : \mu \in \varphi(x^n)\}.$$

The number of pairs in $\mathcal{C}(\varphi(x^n))$ is equal to the profile dimension $\mathcal{D}(\varphi(x^n))$. In addition, both a prevalence and its multiplicity are integers in $[0, n]$, and storing the pair takes $2\log n$ nats. Hence,

it takes at most $2(\log n) \cdot \mathcal{D}(\varphi(x^n))$ nats to store the compressed profile. By Theorem 1, for any distribution $p \in \Delta_{\mathcal{X}}$ and $\Phi^n \sim p$,

$$\mathbb{E}[2(\log n) \cdot \mathcal{D}(\Phi^n)] = \tilde{\Theta}(\lceil H(\Phi^n) \rceil).$$

Sequential compression For any sequence x^n, the setting for sequential profile compression is that at time step $t \in [n]$, the compression algorithm knows only $\varphi(x^t)$ and sequentially encodes the new information. This is equivalent to providing the algorithm $\mu_{x_t}(x^{t-1})$ at time step t.

Suppress x, x^t in the expressions for the ease of illustration. For efficient compression, we sequentially encode the profile φ into a *self-balancing binary search tree* $\mathcal{T}$, with each node storing a multiplicity-prevalence pair (μ, φ_μ) and μ being the search key. We present the algorithm details as follows.

Algorithm 1 Sequential Profile Compression

input sequence $(\mu_{x_t}(x^{t-1}))_{t=1}^n$, tree $\mathcal{T} = \varnothing$
output tree $\mathcal{T}$ that encodes the input sequence
  for t = 1 to n **do**
    if $\mu := \mu_{x_t}(x^{t-1}) \in \mathcal{T}$ **then**
      if $\mu + 1 \in \mathcal{T}$ **then**
        $\varphi_{\mu+1} := \mathcal{T}(\mu + 1) \leftarrow \mathcal{T}(\mu + 1) + 1$
      else
        add $(\mu + 1, 1)$ to $\mathcal{T}$
      end if
      if $\varphi_\mu = 1$ **then** delete (μ, φ_μ) from $\mathcal{T}$
      else $\varphi_\mu := \mathcal{T}(\mu) \leftarrow \mathcal{T}(\mu) - 1$ **endif**
    else
      if $1 \notin \mathcal{T}$ **then** add $(1, 1)$ to $\mathcal{T}$
      else $\mathcal{T}(1) \leftarrow \mathcal{T}(1) + 1$ **endif**
    end if
  end for

The algorithm runs for exactly n iterations, with a $\mathcal{O}(\log n)$ per-iteration time complexity. For an i.i.d. sample $X^n \sim p$, the expected space complexity is again $\tilde{\Theta}(\lceil H(\Phi^n) \rceil)$.

2.5 Attributes of Profile Entropy and Dimension

To further our understanding of profile entropy and dimension, we investigate the analytical and statistical attributes of these characteristics concerning their concentration, computation and approximation, monotonicity, and Lipschitzness.

Concentration

Recall that the multiplicity $\mu_y(x^n)$ denotes the number of times symbol y appearing in x^n. Denote by $\bigvee$ the logical OR operator. For any distribution p and $X^n \sim p$, we have

$$\mathcal{D}_n = \sum_{\mu=1}^{n} \bigvee_{x \in \mathcal{X}} \mathbb{1}_{\mu_x(X^n)=\mu}.$$

The statistical dependency landscape of terms in the summation is rather complex, since $\mu_x(X^n)$ and $\mu_y(X^n)$ are dependent for every (x, y) pair due to the fixed sample size; and so are $\mathbb{1}_{\mu_x(X^n)=\mu_1}$ and $\mathbb{1}_{\mu_x(X^n)=\mu_2}$ for every pair of distinct μ_1 and μ_2. To simplify the derivations, we relate this quantity to its variant under the aforementioned Poisson sampling scheme, i.e., making the sample size an independent $N \sim \mathrm{Poi}(n)$. Specifically, define

$$\tilde{\mathcal{D}}_N := \tilde{\mathcal{D}}(X^N) := \sum_{U=1}^{n} \bigvee_{x \in \mathcal{X}} \mathbb{1}_{\mu_x(X^N)=U}.$$

Note that this is not the same as $\mathcal{D}_N$ since the summation index goes up only to n. Denote the expected value of $\tilde{\mathcal{D}}_N$ by $E_n(p)$. Our result shows that the original $\mathcal{D}_n$ satisfies a Chernoff-Hoeffding type bound centered at $E_n(p)$.

Theorem 6. *Under the above conditions and for any $n \in \mathbb{Z}^+$, $p \in \Delta_{\mathcal{X}}$, and $\gamma > 0$,*

$$\Pr\left(\frac{\mathcal{D}_n}{1+\gamma} \geq E_n(p)\right) \leq 3\sqrt{n}e^{-\min\{\gamma^2,\gamma\}E_n(p)/3},$$

and for any $\gamma \in (0,1)$,

$$\Pr\left(\frac{\mathcal{D}_n}{1-\gamma} \leq E_n(p)\right) \leq 3\sqrt{n}e^{-\gamma^2 E_n(p)/2}.$$

As a corollary, the value of $\mathcal{D}_n$ is often close to $E_n(p)$.

Corollary 1. *Under the same conditions as above and for any $n \in \mathbb{Z}^+$ and distribution $p \in \Delta_{\mathcal{X}}$, with probability at least $1 - 6/\sqrt{n}$,*

$$\frac{1}{2}E_n(p) - 4\log n \leq \mathcal{D}_n \leq 2E_n(p) + 3\log n.$$

In addition, we establish an Efron-Stein type inequality.

Theorem 7. *For any distribution p and $\mathcal{D}_n \sim p$,*

$$\mathrm{Var}(\mathcal{D}_n) \leq \mathbb{E}[\mathcal{D}_n].$$

Computation and Approximation

The above results show that $\mathcal{D}_n \sim p$ is often close to $E_n(p)$, with an exponentially small deviation probability. Hence, to approximate $\mathcal{D}_n$, it suffices to accurately compute $E_n(p)$, the expectation of its Poissonized version $\tilde{\mathcal{D}}_N$. By independence and the linearity of expectations,

$$E_n(p) = \sum_{i=1}^{n}\left(1 - \prod_{x\in\mathcal{X}}\left(1 - e^{-np_x}\frac{(np_x)^i}{i!}\right)\right).$$

The expression is exact but does not relate to p in a simple manner. For an intuitive approximation, we partition the unit interval into a sequence of ranges,

$$I_j := \left((j-1)^2\frac{\log n}{n},\ j^2\frac{\log n}{n}\right], 1 \leq j \leq \sqrt{\frac{n}{\log n}}.$$

denote by p_{I_j} the number of probabilities in I_j, and relate $E_n(p)$ to a shape-reflecting quantity

$$H_n^{\mathcal{S}}(p) := \sum_{j\geq 1}\min\left\{p_{I_j}, j\cdot\log n\right\},$$

the sum of the effective number of probabilities lying within each range Hao and Orlitsky [2019c]. To compute $H_n^{\mathcal{S}}(p)$, we simply count the number of probabilities in each I_j. Our main result shows that $H_n^{\mathcal{S}}(p)$ well approximates $E_n(p)$ over the entire $\Delta_{\mathcal{X}}$, up to logarithmic factors of n.

Theorem 8. *For any $n \in \mathbb{Z}^+$ and $p \in \Delta_{\mathcal{X}}$,*

$$\frac{1}{\sqrt{\log n}}\cdot\Omega(H_n^{\mathcal{S}}(p)) \leq E_n(p) \leq \mathcal{O}(H_n^{\mathcal{S}}(p)).$$

Summary The simple expression shows that $H_n^{\mathcal{S}}(p)$ characterizes the variability of ranges the actual probabilities spread over. As Theorem 8 shows, $H_n^{\mathcal{S}}(p)$ closely approximates $E_n(p)$, the value around which $\mathcal{D}_n \sim p$ concentrates (Theorem 6). Henceforth, we use $H_n^{\mathcal{S}}(p)$ as a proxy for both $H(\Phi^n)$ and $\mathcal{D}_n$, and study its attributes and values.

Monotonicity

Among the many attributes that $H_n^{\mathcal{S}}(p)$ possesses, monotonicity is perhaps most intuitive. One may expect a larger value of $H_n^{\mathcal{S}}(p)$ as the sample size n increases, since additional observations reveal more information on the variability of probabilities. Below we confirm this intuition.

Theorem 9. *For any $n \geq m \gg 1$ and $p \in \Delta_{\mathcal{X}}$,*

$$H_n^{\mathcal{S}}(p) \geq H_m^{\mathcal{S}}(p).$$

Besides the above result that lowerly bounds $H_n^{\mathcal{S}}(p)$ with $H_m^{\mathcal{S}}(p)$ for $m \leq n$, a more desirable result is to upperly bound $H_n^{\mathcal{S}}(p)$ with a function of $H_m^{\mathcal{S}}(p)$.

Such a result will enable us to draw a sample of size $m \leq n$, obtain an estimate of $H_m^{\mathcal{S}}(p)$ from $\mathcal{D}_m$, and use it to bound the value of $H_n^{\mathcal{S}}(p)$ and thus of $\mathcal{D}_n$ for a much larger sample size n. With such an estimate, we can perform numerous tasks such as *predicting* the performance of algorithms in Section 2.2 and 2.3 when more observations are available, and the space needed for storing a longer sample profile. The next theorem provides a simple and tight bound on $H_n^{\mathcal{S}}(p)$ in terms of $H_m^{\mathcal{S}}(p)$.

Theorem 10. *For any $n \geq m \gg 1$ and $p \in \Delta_{\mathcal{X}}$,*

$$H_n^{\mathcal{S}}(p) \leq \sqrt{\frac{n \log n}{m \log m}} \cdot H_m^{\mathcal{S}}(p).$$

The aforementioned application of this result is closely related to the recent works on *learnability estimation* Kong and Valiant [2018], Kong et al. [2019].

Lipschitzness

Viewing $H_n^{\mathcal{S}}(p)$ as a distribution property, we establish its Lipschitzness with respect to a weighted Hamming distance and the ℓ_1 distance. Given two distributions $p, q \in \Delta_{\mathcal{X}}$, the vanilla *Hamming distance* is denoted by

$$h(p, q) := \sum_{x \in \mathcal{X}} \mathbb{1}_{p_x \neq q_x}.$$

The distance is suitable for being a statistical distance since there may be many symbols at which the two distributions differ, yet those symbols account for only a negligible total probability and has little effects on many induced statistics. To address this, we propose a *weighted Hamming distance*

$$h_{\mathcal{W}}(p, q) := \sum_{x \in \mathcal{X}} \max\{p_x, q_x\} \cdot \mathbb{1}_{p_x \neq q_x}.$$

The next result measures the Lipschitzness of $H_n^{\mathcal{S}}$ under $h_{\mathcal{W}}$.

Theorem 11. *For any integer n, and distributions p and q, if $h_{\mathcal{W}}(p, q) \leq \varepsilon$ for some $\varepsilon \geq 1/n$,*

$$\left|H_n^{\mathcal{S}}(p) - H_n^{\mathcal{S}}(q)\right| \leq \mathcal{O}(\sqrt{\varepsilon n}).$$

Replacing $\max\{p_x, q_x\}$ with $|p_x - q_x|$ results in a common similarity measure – the ℓ_1 distance. The next theorem is an analog to the above under this classical distance.

Theorem 12. *For any integer n, and distributions p and q, if $\ell_1(p, q) \leq \varepsilon$ for some $\varepsilon \geq 0$,*

$$\left|H_n^{\mathcal{S}}(p) - cH_n^{\mathcal{S}}(q)\right| \leq \mathcal{O}((\varepsilon n)^{2/3}),$$

where c is a constant in $[1/3, 3]$. Note that the inequality is significant iff $\varepsilon \leq \tilde{\Theta}(1/n^{1/4})$, since the value of $H_n^{\mathcal{S}}(p)$ is at most $\mathcal{O}(\sqrt{n \log n})$ for all p.

2.6 Profile Entropy for Structured Families

Following the study of attributes of profile entropy, we derive nearly tight bounds for the $H_n^{\mathcal{S}}(p)$ values of three important structured families, log-concave, power-law, and histogram. These bounds tighten up and significantly improve those in Hao and Orlitsky [2019c], and show the ability of profile entropy in charactering natural shape constraints.

Below, we follow the convention of specifying the structured distributions over $\mathcal{X} = \mathbb{Z}$.

Log-Concave Distributions

We say that $p \in \Delta_{\mathbb{Z}}$ is *log-concave* if p has a contiguous support and $p_x^2 \geq p_{x-1}p_{x+1}$ for all $x \in \mathbb{Z}$.

The log-concave family encompasses a broad range of discrete distributions, such as Poisson, hyper-Poisson, Poisson binomial, binomial, negative binomial, geometric, and hyper-geometric, with wide applications to numerous research areas, including statistics Saumard and Wellner [2014], computer science Lovász and Vempala [2007], economics An [1997], algebra, and geometry Stanley [1989].

The next result upperly bounds the profile entropy of log-concave families, and is *tight* up to logarithmic factors of n.

Theorem 13. *For any $n \in \mathbb{Z}^+$ and distribution $p \in \Delta_{\mathbb{Z}}$, if p is log-concave and has a variance of σ^2,*

$$H_n^{\mathcal{S}}(p) \leq \mathcal{O}(\log n)\left(1 + \min\left\{\sigma, \frac{n}{\sigma}\right\}\right).$$

This upper bound is uniformly better than the $\min\{\sigma, (n^2/\sigma)^{1/3}\}$ bound in Hao and Orlitsky [2019c].

A similar bound holds for t-mixtures of log-concave distributions. More concretely,

Theorem 14. *For any integer n and distribution $p \in \Delta_{\mathbb{Z}}$, if p is a t-mixture of log-concave distributions each has a variance of σ_i^2, where $i = 1, \ldots, t$,*

$$H_n^{\mathcal{S}}(p) \leq \mathcal{O}(\log n)\left(1 + \min\left\{\sum_i \sigma_i, \max_i\left\{\frac{n}{\sigma_i}\right\}\right\}\right).$$

The introduction about log-concave families covers numerous classical discrete distributions, yet leaves many more continuous ones untouched Bagnoli and Bergstrom [2005]. Below, we present a discretization procedure that preserves distribution shapes such as monotonicity, modality, and log-concavity. Applying this procedure to the Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ further shows the optimality of Theorem 13.

Discretization

Let X be a continuous random variable over $\mathbb{R}$ with density function $f(x)$. For any $x \in \mathbb{R}$, denote by $\lceil x \rfloor$ the closest integer z such that $x \in (z - 1/2, z + 1/2]$. The $\lceil X \rfloor$ has a distribution over $\mathbb{Z}$:

$$p(z) := \int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x)dx, \ \forall z \in \mathbb{Z}.$$

We refer to $\lceil X \rfloor$ as the *discretized version* of X.

Shape preservation By definition, one can readily verify that the above transformation preserves several important shape characteristics of distributions, such as monotonicity, modality, and k-modality (possibly yields a smaller k). The following theorem covers log-concavity.

Theorem 15. *For any continuous random variable X over $\mathbb{R}$ with a log-concave density f, the distribution $p \in \Delta_{\mathbb{Z}}$ associated with $\lceil X \rfloor$ is also log-concave.*

Moment preservation Denote by p the distribution of $\lceil X \rfloor$ for $X \sim f$. Let μ and σ^2 be the mean and variance of density f, given that they exist. The theorem below shows that the discrete distribution p has, within small additive absolute constants, a mean of μ and variance of $\Theta(\sigma^2)$.

Theorem 16. *Under the aforementioned conditions, the mean of $\lceil X \rfloor$ satisfies*

$$\mathbb{E}\lceil X \rfloor = \mu \pm \frac{1}{2},$$

and the variance of $\lceil X \rfloor$ satisfies

$$\sigma^2/2 - 1 \leq \mathbb{E}(\lceil X \rfloor - \mathbb{E}\lceil X \rfloor)^2 \leq 2\sigma^2 + 1.$$

Optimality of Theorem 13

By the above formula, the discretized Gaussian $\lceil \mathcal{N} \rfloor(\mu, \sigma^2)$ has a distribution in the form of

$$p_{G}(z) := \frac{1}{\sqrt{2\pi}\sigma} \int_{z-\frac{1}{2}}^{z+\frac{1}{2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx, \ \forall z \in \mathbb{Z}.$$

Consolidating Theorem 15 and 16 shows that p_{G} is a log-concave distribution with a variance of $\Theta(\sigma^2) \pm 1$. Consequently, Theorem 13 yields the following upper bound:

$$H_n^{\mathcal{S}}(p_{G}) \leq \mathcal{O}(\log n)\left(1 + \min\left\{\sigma, \frac{n}{\sigma}\right\}\right).$$

On the other hand, Section B.5 of the appendices shows

Theorem 17. *Under the aforementioned conditions,*

$$H_n^{\mathcal{S}}(p_{G}) \geq \mathcal{O}(\log n)^{-1}\left(1 + \min\left\{\sigma, \frac{n}{\sigma}\right\}\right).$$

The optimality of Theorem 13 follows by these inequalities.

Power-Law Distributions

We say that a discrete distribution $p \in \Delta_{\mathbb{Z}}$ is a *power-law with power* $\alpha > 0$ if p has a support of $[k] := \{1, \ldots, k\}$ for some $k \in \mathbb{Z}^+ \cup \{\infty\}$ and $p_x \propto x^{-\alpha}$ for all $x \in [k]$.

Power-law is a ubiquitous structure appearing in many situations of scientific interest, ranging from natural phenomena such as the initial mass function of stars Kroupa [2001], species and genera Humphries et al. [2010], rainfall Machado and Rossow [1993], population dynamics Taylor [1961], and brain surface electric potential Miller et al. [2009], to man-made circumstances such as the word frequencies in a text Baayen [2002], income rankings Drăgulescu and Yakovenko [2001], company sizes Axtell [2001], and internet topology Faloutsos et al. [1999].

Unlike log-concave distributions that concentrate around their mean values, power-laws are known to possess "long-tails" and always log-convex. Hence, one may expect the profile entropy of power-law distributions to behave differently from that of log-concave ones. The next theorem justifies this intuition and provides tight upper bounds.

Theorem 18. *For a power-law distribution $p \in \Delta_{[k]}$ with power α, we have*

$$H_n^{\mathcal{S}}(p) \leq 7 \log n + e^2 \cdot \min\{k, \mathcal{U}_n^k(\alpha)\},$$

where

$$\mathcal{U}_n^k(\alpha) := \begin{cases} n^{\frac{1}{1+\alpha}} & \text{if } \alpha \geq 1 + \frac{1}{\log k}; \\ \left(\frac{n}{\log n}\right)^{\frac{1}{1+\alpha}} & \text{if } 1 \leq \alpha < 1 + \frac{1}{\log k}; \\ \sqrt{n}\left(\frac{k}{\sqrt{n}} \wedge \left(\frac{\sqrt{n}}{k}\right)^{\frac{1-\alpha}{1+\alpha}}\right) & \text{if } 0 \leq \alpha < 1. \end{cases}$$

The above upper bound fully characterizes the profile entropy of power-laws and surpasses the basic $\{k, \sqrt{n \log n}\}$ bound for both $k \gg \sqrt{n}$ and $k \ll \sqrt{n}$. In comparison, Hao and Orlitsky [2019c] yields a $\mathcal{O}(n^{\min\{1/(1+\alpha), 1/2\}})$ upper bound, which improves over $\sqrt{n \log n}$ for only $\alpha > 1$ and is worse than that above for all $\alpha < 1 + 1/\log k$.

Histogram Distributions

A distribution $p \in \Delta_{\mathcal{X}}$ is a *t-histogram distribution* if there is a partition of $\mathcal{X}$ into t parts such that p has the same probability value over all symbols in each part.

Besides the long line of research on histograms reviewed in Ioannidis [2003], the importance of histogram distributions rises with the rapid growth of data sizes in numerous engineering and science applications in the modern era.

For example, in scenarios where processing the complete data set is inefficient or even impossible, a standard solution is to partition/cluster the data into groups according to the task specifications

and element similarities, and randomly sample from each group to obtain a subset of the data to use. This naturally induces a histogram distribution, with each data point being a symbol in the support.

The work of Hao and Orlitsky [2019c] studies the class of t-histogram distributions and obtains the following upper bound

$$H_n^{\mathcal{S}}(p) \le \tilde{\mathcal{O}}\left(\min\left\{(nt^2)^{\frac{1}{3}}, \sqrt{n}\right\}\right).$$

Our contribution is establishing its optimality.

Theorem 19. *For any $t, n \in \mathbb{Z}^+$, there exists a t-histogram distribution p such that*

$$H_n^{\mathcal{S}}(p) \ge \tilde{\Omega}\left(\min\left\{(nt^2)^{\frac{1}{3}}, \sqrt{n}\right\}\right).$$

Note that uniform distributions correspond to 1-histograms, for which the bounds reduce to $\tilde{\Theta}(n^{1/3})$.

3 Extension and Conclusion

3.1 Multi-Dimensional Profiles

The notion of profile generalizes to the multi-sequence setting.

Let $\mathcal{X}$ be a finite or countably infinite alphabet. For every $\vec{n} := (n_1, \ldots, n_d) \in \mathbb{N}^d$ and tuple $x^{\vec{n}} := (x_1^{n_1}, \ldots, x_d^{n_d})$ of sequences in $\mathcal{X}^*$, the *multiplicity* $\mu_y(x^{\vec{n}})$ of a symbol $y \in \mathcal{X}$ is the vector of its frequencies in the tuple of sequences. The *profile* of $x^{\vec{n}}$ is the multiset $\varphi(x^{\vec{n}})$ of multiplicities of the observed symbols Acharya et al. [2010], Das [2012], Charikar et al. [2019b], and its *dimension* is the number $\mathcal{D}(x^{\vec{n}})$ of distinct elements in the multiset. Drawing independent samples from $\vec{p} := (p_1, \ldots, p_d) \in \Delta_{\mathcal{X}}^d$, the *profile entropy* is simply the entropy of the joint-sample profile.

Many of the previous results potentially generalize to this multi-dimensional setting. For example, the $\sqrt{2n}$ bound on $\mathcal{D}(x^{\vec{n}})$ in the 1-dimensional case becomes

Theorem 20. *For any $\mathcal{X}$, $\vec{n}$, and $x^{\vec{n}} \in \mathcal{X}^{\vec{n}}$, there exists a positive integer r such that*

$$\sum_i n_i \ge d \cdot \binom{d+r-1}{d+1},$$

and

$$\mathcal{D} \le \binom{d+r}{d} - 1.$$

This essentially recovers the $\sqrt{2n}$ bound for $d = 1$.

3.2 Concluding Remarks

The classical view on the entropy of an i.i.d. sample corresponds to the equation

$$H(X^n \sim p) = nH(p),$$

which provides little insight for statistical applications.

This paper presents a different view by decomposing the $nH(p)$ information into three pieces: the labeling of the profile elements, ordering of them, and profile entropy. With no bias towards any symbols and under the i.i.d. assumption, the profile entropy rises as a fundamental measure unifying the concepts of estimation, inference, and compression.

Appendices orgnization In the appendices, we order the results and proofs according to their logical priority. In other words, the proof of a theorem or lemma mainly relies on preceding results. For the ease of reference, the numbering of the theorems is consistent with that in the main paper.

A Dimension and Entropy of Sample Profiles

Consider an arbitrary sequence x^n over a finite or countably infinite alphabet $\mathcal{X}$. The *multiplicity* $\mu_y(x^n)$ of a symbol $y \in \mathcal{X}$ is the frequency of y in x^n. The *prevalence* of an integer μ is the number $\varphi_\mu(x^n)$ of symbols in x^n with multiplicity μ. The *profile* of x^n is the multiset $\varphi(x^n)$ of multiplicities of the symbols in x^n, which we describe as a profile of *length* n.

Let $\Delta_\mathcal{X}$ be a collection of distributions over $\mathcal{X}$. We say that a distribution collection $\mathcal{P} \subseteq \Delta_\mathcal{X}$ is *label-invariant* if for any $p \in \mathcal{P}$, the collection $\mathcal{P}$ contains all its symbol-permuted versions. A *distribution property* over a distribution collection $\mathcal{P} \subseteq \Delta_\mathcal{X}$ is a functional $f : \mathcal{P} \to \mathbb{R}$ that associates with each distribution in $\mathcal{P}$ a real value. For a label-invariant $\mathcal{P} \subseteq \Delta_\mathcal{X}$, we say that a property f over $\mathcal{P}$ is symmetric if it takes the same value for distributions sharing the same probability multiset.

A.1 Profile Dimension and Its Concentration

A profile $\boldsymbol{\phi}$ is said to have length n if there exists $x^n \in \mathcal{X}^n$ satisfying $\boldsymbol{\phi} = \varphi(x^n)$. For any multiset S, we define its *dimension* as the number $\mathcal{D}(S)$ of distinct elements in S. Recall that profiles of sequences are multisets. For notational convenience, we write both $\mathcal{D}(\varphi(x^n))$ and $\mathcal{D}(x^n)$ for the dimension of profile $\varphi(x^n)$.

Viewed as a random variable, the profile of an i.i.d. sample $X^n \sim p$ has a distribution depending only on p and n. Hence, we denote by Φ^n such a profile and write $\Phi^n \sim p$. Analogously, we denote by $\mathcal{D}_n := \mathcal{D}(\Phi^n)$ the *profile dimension* associated with (p, n), and write $\mathcal{D}_n \sim p$.

Recall that the multiplicity $\mu_y(x^n)$ denotes the number of times symbol y appearing in x^n. Denote by $\bigvee$ the logical OR operator. For any distribution p and $X^n \sim p$, we have

$$\mathcal{D}_n = \sum_{\mu=1}^{n} \bigvee_{x\in\mathcal{X}} \mathbb{1}_{\mu_x(X^n)=\mu}.$$

The statistical dependency landscape of terms in the summation is rather complex, since $\mu_x(X^n)$ and $\mu_y(X^n)$ are dependent for every (x, y) pair due to the fixed sample size; and so are $\mathbb{1}_{\mu_x(X^n)=\mu_1}$ and $\mathbb{1}_{\mu_x(X^n)=\mu_2}$ for every pair of distinct μ_1 and μ_2. To simplify the derivations, we relate this quantity to its variant under the aforementioned Poisson sampling scheme, i.e., making the sample size an independent $N \sim \mathrm{Poi}(n)$. Specifically, define

$$\tilde{\mathcal{D}}_N := \tilde{\mathcal{D}}(X^N) := \sum_{U=1}^{n} \bigvee_{x\in\mathcal{X}} \mathbb{1}_{\mu_x(X^N)=U}.$$

Note that this is not the same as $\mathcal{D}_N$ since the summation index goes up only to n. Denote the expected value of $\tilde{\mathcal{D}}_N$ by $E_n(p)$. Our result shows that the original $\mathcal{D}_n$ satisfies a Chernoff-Hoeffding type bound centered at $E_n(p)$.

Theorem 6. *Under the above conditions and for any $n \in \mathbb{Z}^+$, $p \in \Delta_\mathcal{X}$, and $\gamma > 0$,*

$$\Pr\left(\frac{\mathcal{D}_n}{1+\gamma} \geq E_n(p)\right) \leq 3\sqrt{n}e^{-\min\{\gamma^2,\gamma\}E_n(p)/3},$$

and for any $\gamma \in (0, 1)$,

$$\Pr\left(\frac{\mathcal{D}_n}{1-\gamma} \leq E_n(p)\right) \leq 3\sqrt{n}e^{-\gamma^2 E_n(p)/2}.$$

Proof. To simplify our analysis, we first consider an alternative model where the sample size is an independent Poisson random variable N with mean n. A nice attribute of Poisson sampling is that all the multiplicities $\mu_y(X^n)$ are independent of each other. Later, we will relate this model to the fixed-sample-size model and establish our claim rigorously.

For simplicity and clarity, we suppress X^n in $\mu_y(X^n)$ and write ν_y instead of μ_y when the multiplicity is obtained through Poisson sampling. For any $i \in [n]$, denote $G_i(\{\nu_x\}_x) := \bigvee_{x\in\mathcal{X}} \mathbb{1}_{\nu_x=i}$. Instead of analyzing $\mathcal{D}_N$, we consider

$$\tilde{\mathcal{D}}_N := \sum_{i=1}^{n} \bigvee_{x\in\mathcal{X}} \mathbb{1}_{\nu_x=i} = \sum_{i=1}^{n} G_i(\{\nu_x\}_x).$$

Note that for any disjoint $I, J \subseteq [n]$, the functions $\sum_{i\in I} G_i(\{\nu_x\}_x)$ and $\sum_{j\in J} G_j(\{\nu_x\}_x)$ are discordant monotone by each argument, namely, when we increase the value of each ν_x, the increase in the value of one function implies the non-increase of the other. Then, by the results in Lehmann [1966], the values of the two functions, when viewed as random variables, are negatively associated.

Next we show that quantity $\tilde{\mathcal{D}}_N$ satisfies a Chernoff-type bound.

Let γ be an arbitrary positive number. Note that G_i is a Bernoulli random variable with parameter

$$q_i := \mathbb{E}\left[G_i(\{\nu_x\}_x)\right].$$

Then for the expected value of $\tilde{\mathcal{D}}_N$, we have

$$E_n(p) := \mathbb{E}\left[\tilde{\mathcal{D}}_N\right] = \mathbb{E}\left[\sum_{i=1}^{n} G_i(\{\nu_x\}_x)\right] = \sum_i q_i.$$

For simplicity, write $Y := \tilde{\mathcal{D}}_N$ and $\mu := E_n(p)$. By Markov's inequality and the monotonicity of function e^{ty} over $t > 0$,

$$\Pr\left(Y \geq (1+\gamma)\mu\right) = \Pr\left(e^{tY} \geq e^{t(1+\gamma)\mu}\right) \leq \frac{\mathbb{E}[e^{tY}]}{e^{t(1+\gamma)\mu}}.$$

It suffices to bound $\mathbb{E}[e^{tY}]$ by a function of other parameters.

$$\begin{aligned}
\mathbb{E}[e^{tY}] &\overset{(a)}{=} \mathbb{E}\left[\exp\left(t\left(\sum_{i=1}^{n} G_i(\{M_x\}_x)\right)\right)\right] \\
&\overset{(b)}{=} \mathbb{E}\left[\exp\left(tG_1(\{M_x\}_x)\right)\cdot\exp\left(t\left(\sum_{i=2}^{n} G_i(\{M_x\}_x)\right)\right)\right] \\
&\overset{(c)}{\leq} \mathbb{E}\left[\exp\left(tG_1(\{M_x\}_x)\right)\right]\cdot\mathbb{E}\left[\exp\left(t\left(\sum_{i=2}^{n} G_i(\{M_x\}_x)\right)\right)\right] \\
&\overset{(d)}{\leq} \prod_{i=1}^{n}\mathbb{E}\left[\exp\left(tG_i(\{M_x\}_x)\right)\right] \overset{(e)}{=} \prod_{i=1}^{n}\left(1+q_i(e^t-1)\right) \\
&\overset{(f)}{\leq} \prod_{i=1}^{n}\left(\exp\left(q_i(e^t-1)\right)\right) \overset{(g)}{=} \exp\left(\sum_{i=1}^{n} q_i(e^t-1)\right) \\
&\overset{(h)}{=} \exp\left((e^t-1)\mu\right),
\end{aligned}$$

where (a) follows by the definition of Y; (b) follows by $e^{a+b} = e^a \cdot e^b$; (c) follows by the fact that G_1 is negatively associated with $\sum_{i=2}^{n} G_i$; (d) follows by an induction argument via negative association; (e) follows by the fact that G_i is a Bernoulli random variable with mean q_i; (f) follows by the inequality $1 + x \leq e^x, \forall x \geq 0$; (g) follows by $e^a \cdot e^b = e^{a+b}$; and (h) follows by $\mu = \sum_i q_i$.

Applying standard simplifications, we obtain

$$\Pr\left(Y \geq (1+\gamma)\mu\right) \leq e^{-\min\{\gamma^2,\gamma\}\mu/3},\ \forall\gamma > 0,$$

and

$$\Pr\left(Y \leq (1-\gamma)\mu\right) \leq e^{-\gamma^2\mu/2},\ \forall\gamma \in (0,1).$$

The proof will be complete upon noting that: 1) the probability that $N = n$ is at least $1/(3\sqrt{n})$; 2) conditioning on $N = n$ transforms the sampling model to that with a fixed sample size n. □

As a corollary, the value of $\mathcal{D}_n$ is often close to $E_n(p)$.

Corollary 2. *Under the same conditions as above and for any $n \in \mathbb{Z}^+$, $p \in \Delta_{\mathcal{X}}$, with probability at least $1 - 6/\sqrt{n}$,*

$$\frac{1}{2}E_n(p) - 4\log n \le \mathcal{D}_n \le 2E_n(p) + 3\log n.$$

Proof. To establish the lower bound, note that if $E_n(p) \ge 3\log n$, setting $\gamma = 1$ in Theorem 6 yields

$$\Pr\left(\mathcal{D}_n \ge 2E_n(p) + 3\log n\right) \le \Pr\left(\mathcal{D}_n \ge 2E_n(p)\right) \le 3\sqrt{n}e^{-E_n(p)/3} \le \frac{3}{\sqrt{n}},$$

else if $E_n(p) < 3\log n$, setting $\gamma = (3\log n)/E_n(p)$ yields

$$\Pr\left(\mathcal{D}_n \ge 2E_n(p) + 3\log n\right) \le \Pr\left(\mathcal{D}_n \ge E_n(p) + 3\log n\right) \le 3\sqrt{n}e^{-(3\log n)/3} = \frac{3}{\sqrt{n}}.$$

As for the upper bound, if $E_n(p) \ge 8\log n$,

$$\Pr\left(\mathcal{D}_n + 4\log n \le \left(1 - \frac{1}{2}\right)E_n(p)\right) \le \Pr\left(\mathcal{D}_n \le \left(1 - \frac{1}{2}\right)E_n(p)\right) \le 3\sqrt{n}e^{-\mu/8} \le \frac{3}{\sqrt{n}},$$

and for any $E_n(p) < 8\log n$,

$$\Pr\left(\mathcal{D}_n + 4\log n \le \left(1 - \frac{1}{2}\right)E_n(p)\right) \le \Pr\left(\mathcal{D}_n < 0\right) = 0 \le \frac{3}{\sqrt{n}}.$$

Combining these tail bounds through the union bound completes the proof. □

In addition to the above, we establish an Efron-Stein type inequality.

Theorem 7. *For any distribution p and $\mathcal{D}_n \sim p$,*

$$\mathrm{Var}(\mathcal{D}_n) \le \mathbb{E}[\mathcal{D}_n].$$

Proof. First, note that for any $j, t \in [n]$ and $j \ne t$,

$$\begin{aligned}
C_{j,t} &:= \mathrm{Cov}\left(\mathbb{1}_{\varphi_j(X^n)>0}, \mathbb{1}_{\varphi_t(X^n)>0}\right) \\
&= \Pr\left(\varphi_j(X^n), \varphi_t(X^n) > 0\right) - \Pr\left(\varphi_j(X^n) > 0\right) \cdot \Pr\left(\varphi_t(X^n) > 0\right) \\
&= \left(\Pr\left(\varphi_j(X^n) > 0 | \varphi_t(X^n) > 0\right) - \Pr\left(\varphi_j(X^n) > 0\right)\right) \cdot \Pr\left(\varphi_t(X^n) > 0\right) \\
&= \left(\Pr\left(\varphi_j(X^n) > 0 | \varphi_t(X^n) > 0\right) - \Pr\left(\varphi_j(X^n) > 0 | \varphi_t(X^n) = 0\right)\right) \\
&\quad \cdot \Pr\left(\varphi_t(X^n) = 0\right) \cdot \Pr\left(\varphi_t(X^n) > 0\right) \\
&\le 0
\end{aligned}$$

Therefore, the variance of the profile dimension $\mathcal{D}_n$ satisfies

$$\begin{aligned}
\mathrm{Var}\left(\mathcal{D}_n\right) &= \mathrm{Var}\left(\sum_{i=1}^{n} \mathbb{1}_{\varphi_i(X^n)>0}\right) \\
&\le \sum_{i=1} \mathrm{Var}\left(\mathbb{1}_{\varphi_i(X^n)>0}\right) + \sum_{j \ne t} \mathrm{Cov}\left(\mathbb{1}_{\varphi_j(X^n)>0}, \mathbb{1}_{\varphi_t(X^n)>0}\right) \\
&\le \sum_{i=1} \mathbb{E}\left[\mathbb{1}_{\varphi_i(X^n)>0}\right] + \sum_{j \ne t} C_{j,t} \\
&\le \sum_{i=1} \mathbb{E}\left[\mathbb{1}_{\varphi_i(X^n)>0}\right] \\
&= \mathbb{E}\left[\mathcal{D}_n\right].
\end{aligned}$$

□

A.2 Profile Entropy and Its Connection to Dimension

For a distribution $p \in \Delta_{\mathcal{X}}$ and sampling parameter n, the *profile entropy* with respect to (p, n) is the entropy $H(\Phi^n)$ of the sample profile $\Phi^n \sim p$. By Shannon's source coding theorem, profile entropy $H(\Phi^n)$ is the *information-theoretic limit* of the minimal expected codeword length (MECL) for the lossless compression of the sample profile. Hence, characterizing its value is of fundamental importance. But as one may expect, the distribution of Φ^n is sophisticated and over a large alphabet.

More concretely, by the formula of Hardy and Ramanujan [1918], the number $\mathbb{P}(n)$ of integer partitions of n, which happens to equal to the number of length-n profiles, satisfies the equation

$$\log \mathbb{P}(n) = 2\pi\sqrt{\frac{n}{6}}(1 + o(1)).$$

Despite the complex statistical dependency landscape and the exponentially large alphabet size, below we establish that for any distribution and sample size, the profile entropy is often of the same order as the profile size, with high probability. Specifically,

Theorem 1. *For any distribution $p \in \Delta_{\mathcal{X}}$ and $\Phi^n \sim p$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,*

$$\lceil H(\Phi^n) \rceil = \tilde{\Theta}(\mathcal{D}(\Phi^n)),$$

where the notation $\tilde{\Theta}(\cdot)$ hides logarithmic factors of n.

We decompose the proof of the theorem into three steps. First, we show that $\lceil H(\Phi^n) \rceil \leq \tilde{\Theta}(\mathcal{D}(\Phi^n))$ with high probability, which is a simple consequence of Shannon's source coding theorem and Theorem 6 (which shows that $\mathcal{D}(\Phi^n)$ highly concentrates around its expectation). Then, we introduce a simple quantity $H_n^{\mathcal{S}}(p)$ that approximates the expectation of $\mathcal{D}(\Phi^n)$ to within logarithmic factors of n. Finally, leveraging this approximation guarantee, we establish the other direction of the theorem. This step is more involved due to the aforementioned complications.

A. Bounding Profile Entropy by Its Dimension

By the tail bounds (Theorem 6) and trivial lower bound of 1 on the profile dimension, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$, the expectation of $\mathcal{D}(\Phi^n)$ satisfies

$$\mathbb{E}[\mathcal{D}(\Phi^n)] \leq \tilde{\mathcal{O}}(\mathcal{D}(\Phi^n)).$$

By the block profile compression algorithm presented in Section 2.4 of the main paper, storing profile $\Phi^n \sim p$ losslessly takes

$$\mathcal{O}(\log n) \cdot \mathbb{E}[\mathcal{D}(\Phi^n)] + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \cdot \log \mathbb{P}(n) \leq \mathcal{O}(\log n) \cdot \mathbb{E}[\mathcal{D}(\Phi^n)]$$

nats space in expectation. By Shannon's source coding theorem, the expected space to losslessly storing a random variable is at least its entropy. Hence, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,

$$H(\Phi^n) \leq \mathcal{O}(\log n) \cdot \mathbb{E}[\mathcal{D}(\Phi^n)] \leq \tilde{\mathcal{O}}(\mathcal{D}(\Phi^n)).$$

Again, noting that $\mathcal{D}(\Phi^n) \geq 1$ completes the proof.

B. Simple Approximation Formula for Profile Dimension

It remains to show that $\lceil H(\Phi^n) \rceil \geq \tilde{\Omega}(\mathcal{D}(\Phi^n))$, with high probability. To proceed further, we note that $\mathcal{D}(\Phi^n) = \mathcal{D}_n \sim p$ is often close to $E_n(p)$, the expectation of its Poissonized version $\tilde{\mathcal{D}}_N$, with an exponentially small deviation probability. Hence, to approximate $\mathcal{D}_n$, it suffices to accurately compute $E_n(p)$. By independence and the linearity of expectations,

$$E_n(p) = \mathbb{E}[\tilde{\mathcal{D}}_N] = \sum_{i=1}^{n} \left(1 - \prod_{x \in \mathcal{X}} \left(1 - e^{-np_x} \frac{(np_x)^i}{i!}\right)\right).$$

The expression is exact but does not relate to p in a simple manner. For an intuitive approximation, we partition the unit interval into a sequence of ranges,

$$I_j := \left((j-1)^2 \frac{\log n}{n},\ j^2 \frac{\log n}{n}\right], 1 \leq j \leq \sqrt{\frac{n}{\log n}},$$

denote by p_{I_j} the number of probabilities p_x belonging to I_j, and relate $E_n(p)$ to an induced shape-reflecting quantity,

$$H_n^{\mathcal{S}}(p) := \sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\},$$

the sum of the effective number of probabilities lying within each range Hao and Orlitsky [2019c]. To compute $H_n^{\mathcal{S}}(p)$, we simply count the number of probabilities in each I_j. Our main result shows that $H_n^{\mathcal{S}}(p)$ well approximates $E_n(p)$ over the entire $\Delta_{\mathcal{X}}$, up to logarithmic factors of n.

Theorem 8. *For any $n \in \mathbb{Z}^+$ and $p \in \Delta_{\mathcal{X}}$,*

$$\frac{1}{\sqrt{\log n}} \cdot \Omega(H_n^{\mathcal{S}}(p)) \leq E_n(p) \leq \mathcal{O}(H_n^{\mathcal{S}}(p)).$$

Proof. The fact that $\mathcal{O}(H_n^{\mathcal{S}}(p))$ upperly bounds $\mathbb{E}[\tilde{\mathcal{D}}_N]$ simply follows by the concentration of Poisson variables, and is established in Hao and Orlitsky [2019c]. Below we show that the quantity also serves as a lower bound. By construction, for any given sampling parameter n, index j, and symbol x with probability $p_x \in I_j$, the corresponding symbol multiplicity $\mu_x \sim \mathrm{Poi}(np_x)$. Hence, we can express the expectation of $\tilde{\mathcal{D}}_N$ as

$$\begin{aligned}
\mathbb{E}\left[\tilde{\mathcal{D}}_N\right] &= \mathbb{E}\left[\sum_{i=1}^{n} \bigvee_x \mathbb{1}_{\mu_x = i}\right] \\
&= \sum_{i=1}^{n} \mathbb{E}\left[1 - \bigwedge_x \mathbb{1}_{\mu_x \neq i}\right] \\
&= \sum_{i=1}^{n} \left(1 - \mathbb{E}\left[\prod_x \mathbb{1}_{\mu_x \neq i}\right]\right) \\
&= \sum_{i=1}^{n} \left(1 - \prod_x \mathbb{E}\left[\mathbb{1}_{\mu_x \neq i}\right]\right) \\
&= \sum_{i=1}^{n} \left(1 - \prod_x \left(1 - e^{-np_x} \frac{(np_x)^i}{i!}\right)\right).
\end{aligned}$$

This proves the aforementioned formula. Then, for every sufficiently large index j and $i \in S_j := [(j-1)^2, j^2] \log n$, define a sequence of intervals,

$$I_j^i := \frac{i}{n} + [-j, j] \frac{\sqrt{\log n}}{n}.$$

Then for any $i \in S_j$ and $p_x \in I_j^i \cap I_j$, the corresponding Poisson probability satisfies

$$\begin{aligned}
e^{-np_x} \frac{(np_x)^i}{i!} &= e^{-i} \frac{i^i}{i!} \cdot \left(e^{i - np_x} \cdot \frac{(np_x)^i}{i^i}\right) \\
&= e^{-i} \frac{i^i}{i!} \cdot \left(e^{-(np_x - i)} \cdot \left(1 + \frac{np_x - i}{i}\right)^i\right) \\
&= e^{-i} \frac{i^i}{i!} \cdot \exp\left(-(np_x - i) + i \cdot \log\left(1 + \frac{np_x - i}{i}\right)\right) \\
&\geq \frac{1}{3\sqrt{i}} \cdot \exp\left(-\frac{2i}{3} \cdot \left(\frac{np_x - i}{i}\right)^2\right) \\
&\geq \frac{1}{9\sqrt{i}} \geq \frac{1}{9j\sqrt{\log n}}.
\end{aligned}$$

Now we analyze the contribution of indices $i \in S_j$ to the expected value of $\tilde{\mathcal{D}}_N$. For clarity, we divide our analysis into two cases: $p_{I_j} \geq j \log n$ and $p_{I_j} < j \log n$.

Consider the collection $\mathcal{P}_j$ of probabilities $p_x \in I_j$, and the collection $\mathcal{I}_j$ of intervals $I_j^i, i \in S_j$. By construction, each probability in $\mathcal{P}_j$ is contained in at least $j\sqrt{\log n}$ many intervals in $\mathcal{I}_j$. Hence the total number of probabilities (repeatedly counted) included in $\mathcal{I}_j$ is at least $p_{I_j} \cdot j\sqrt{\log n}$. Note that the number of intervals in $\mathcal{I}_j$ is less than $2j \log n$. We claim that there exists one (or more) interval $I_j^{i'} \in \mathcal{I}_j$ containing at least $p_{I_j}/(2\sqrt{\log n})$ probabilities. By construction, there are at least $j\sqrt{\log n}/2$ neighboring intervals of $I_j^{i'}$ that contain at least $p_{I_j}/(4\sqrt{\log n})$ probabilities. The contribution of these these intervals to the expected value of $\tilde{\mathcal{D}}_N$ is at least $j\sqrt{\log n}/2$ times

$$\begin{aligned}
1 - \left(1 - \frac{1}{9j\sqrt{\log n}}\right)^{\frac{p_{I_j}}{4\sqrt{\log n}}} &\geq 1 - \exp\left(\frac{p_{I_j}}{4\sqrt{\log n}} \log\left(1 - \frac{1}{9j\sqrt{\log n}}\right)\right) \\
&\geq 1 - \exp\left(-\frac{p_{I_j}}{40j \log n}\right) \\
&\geq \Theta\left(\frac{p_{I_j}}{j \log n}\right),
\end{aligned}$$

where the last step holds if $p_{I_j} \leq j \log n$. This yields a lower bound of $\Theta(p_{I_j}/\sqrt{\log n})$.

It remains to consider the $p_{I_j} > j \log n$ case. Again, the total number of probabilities included in $\mathcal{I}_j$ is at least $p_{I_j} \cdot j\sqrt{\log n}$. Furthermore, each interval I_j^i contains at most p_{I_j} probabilities and there are less than $2j \log n$ intervals. Therefore, the number of intervals that contain at least $j\sqrt{\log n}/4$ probabilities is at least $j\sqrt{\log n}/2$. Otherwise, the number of probabilities included in $\mathcal{I}_j$ is less than

$$\frac{j\sqrt{\log n}}{4} \cdot 2j \log n + p_{I_j} \cdot \frac{j\sqrt{\log n}}{2} \leq p_{I_j} \cdot j\sqrt{\log n},$$

which leads to a contradiction. Analogously, the contribution of these these intervals to the expected value of $\tilde{\mathcal{D}}_N$ is at least $j\sqrt{\log n}/2$ times

$$\begin{aligned}
1 - \left(1 - \frac{1}{9j\sqrt{\log n}}\right)^{\frac{j\sqrt{\log n}}{4}} &\geq 1 - \exp\left(\frac{j\sqrt{\log n}}{4} \log\left(1 - \frac{1}{9j\sqrt{\log n}}\right)\right) \\
&\geq 1 - \exp\left(-\frac{1}{40}\right) \\
&= \Theta\left(1\right),
\end{aligned}$$

which yields a lower bound of $\Theta(j\sqrt{\log n})$ on the expected value of $\tilde{\mathcal{D}}_N$.

Consolidating the previous results shows that

$$\mathbb{E}\left[\tilde{\mathcal{D}}_N\right] \geq \frac{1}{\sqrt{\log n}} \cdot \Omega(\sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\}).$$ □

C. Bounding Profile Dimension by Its Entropy

Next, we establish that for any distribution $p \in \Delta_{\mathcal{X}}$, $\Phi^n \sim p$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,

$$\lceil H(\Phi^n) \rceil \geq \tilde{\Theta}(\mathcal{D}(\Phi^n)).$$

Let p be an arbitrary distribution in $\Delta_{\mathcal{X}}$. Recall that we partition the interval $(0, 1]$ into a sequence of sub-intervals,

$$I_j := \left((j-1)^2 \frac{\log n}{n}, j^2 \frac{\log n}{n}\right], \quad 1 \leq j \leq \sqrt{\frac{n}{\log n}},$$

and denote by p_{I_j} the number of probabilities p_x in I_j.

Our current objective is to bound $H(\Phi^n \sim p)$ from below by a nontrivial multiple of $H_n^{\mathcal{S}}(p)$. For simplicity of derivations, we will adopt the standard Poisson sampling scheme and make the sample size an independent Poisson variable $N \sim \mathrm{Poi}(n)$. For notational simplicity, we will suppress X^N in all the expressions and write the profile as $\varphi := \Phi^N$ by slightly abusing the notation.

Note that the profile can be equivalently expressed as a length-n vector

$$\varphi = (\varphi_1, \ldots, \varphi_n),$$

where φ_i denotes the number of symbols appearing exactly i times.

For a sufficiently large absolute constant c, decompose φ into c parts according to I_j such that the t-th part ($t = 1, \ldots, c$) consists of φ_i's satisfying $i \in nI_j$ with $j \equiv t \mod c$. Since by definition,

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\{p_{I_j}, j \cdot \log n\},$$

one of the c parts corresponds to a partial sum of at least $H_n^{\mathcal{S}}(p)/c$. Without loss of generality, we assume that it is the second part, i.e.,

$$\sum_{j \equiv 1 \mod c} \min\{p_{I_j}, j \cdot \log n\} \geq \frac{H_n^{\mathcal{S}}(p)}{c}.$$

Apply standard Poisson tail probability bounds, e.g.,

Lemma 2. *Let Y be a Poisson or binomial random variable with mean value λ. Then,*

$$\Pr(X \leq \lambda(1-\delta)) \leq \exp\left(-\frac{\delta^2\lambda}{2}\lambda\right), \quad \forall \delta \in [0,1],$$

and

$$\Pr(X \geq \lambda(1+\delta)) \leq \exp\left(-\frac{\delta^2\lambda}{2+2\delta/3}\right), \quad \forall \delta \geq 0.$$

For any $j \equiv 1 \mod c$ and with probability at least $1 - 1/n^4$, one can express the truncated profile $(\varphi_i)_{i \in nI_j}$ over I_j as a function of μ_x for x satisfying $np_x \in I_{j'}, j' \in (j - c/2, j + c/2)$.

Basically, this says that for every x, the number of its appearance is not too far away from the expected value. By the union bound, this is true for all $j \equiv 1 \mod c$ with probability at least $1 - 1/n^3$, as j can take only n possible values. Denote the last event by A.

To proceed, we recall the formula of Hardy and Ramanujan [1918] on the number $\mathbb{P}(n)$ of integer partitions of n, which happens to equal to the number of length-n profiles:

$$\log \mathbb{P}(n) = 2\pi\sqrt{\frac{n}{6}}(1 + o(1)).$$

Below, we will use a weaker version that works for any n:

$$\log \mathbb{P}(n) \leq \sqrt{3n}.$$

Then, conditioning on A, the truncated profiles $(\varphi_i)_{i \in nI_j}$ for $j \equiv 1 \mod c$ are independent. Since conditioning reduces entropy,

$$\begin{aligned}
H(\varphi) &\geq H((\varphi_i)_{i \in nI_j, j \equiv 1 \bmod c}) \\
&\geq H((\varphi_i)_{i \in nI_j, j \equiv 1 \bmod c} | \mathbb{1}_A) \\
&\geq H((\varphi_i)_{i \in nI_j, j \equiv 1 \bmod c} | \mathbb{1}_A = 1) \cdot \Pr(A) \\
&= \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j} | \mathbb{1}_A = 1) \cdot \Pr(A) \\
&= \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j} | \mathbb{1}_A) - \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j} | \mathbb{1}_A = 0) \cdot (1 - \Pr(A)) \\
&\geq \sum_{j \equiv 1 \bmod c} (H((\varphi_i)_{i \in nI_j}) - H(\mathbb{1}_A)) - \frac{1}{n^3} \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j} | \mathbb{1}_A = 0) \\
&\geq -nH(\mathbb{1}_A) + \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j}) - \frac{1}{n^3} \cdot n \cdot \log(\exp(\Theta(\sqrt{n}))) \\
&= -\mathcal{O}\left(\frac{1}{\sqrt{n}}\right) + \sum_{j \equiv 1 \bmod c} H((\varphi_i)_{i \in nI_j}),
\end{aligned}$$

where the third last step follows by

$$H(X|Y) = H(X) - I(X,Y) = H(X) - H(Y) + H(Y|X) \geq H(X) - H(Y);$$

the second last follows by $H(X) \leq \log k$ for any X with a support size of k, and the fact that there are at most $\exp(3\sqrt{m})$ many profiles of length m, as we explained above; and the last step follows by the elementary inequality

$$H(\mathrm{Bern}(\theta)) \leq 2(\log 2)\sqrt{\theta(1-\theta)},\ \forall \theta \in [0,1].$$

Our new objective is to bound $H((\varphi_i)_{i\in nI_j})$ from below. We will find a sub-interval I_j^s of I_j and bound $H((\varphi_i)_{i\in nI_j^s})$ in the rest of the section, since

$$H((\varphi_i)_{i\in nI_j}) \geq H((\varphi_i)_{i\in nI_j^s}).$$

For all $j \equiv 1 \mod c$, our lower bound is simply

$$H((\varphi_i)_{i\in nI_j^s}) \geq \Omega\left(\frac{1}{\sqrt{\log n}} \min\left\{p_{I_j}, j\cdot\log n\right\}\right),$$

which, together with $\sum_{j\equiv 1 \bmod c} \min\{p_{I_j}, j\cdot\log n\} \geq H_n^{\mathcal{S}}(p)/c$, implies that

$$H(\varphi) \geq -\mathcal{O}\left(\frac{1}{\sqrt{n}}\right) + \sum_{j\equiv 1 \bmod c} H((\varphi_i)_{i\in nI_j}) \geq \Omega\left(\frac{1}{\sqrt{\log n}}\right)\cdot T_n.$$

Henceforth, we assume that j is sufficiently large and denote $L_j := j\sqrt{\log n}$.

For any j and every integer $i \in S_j := [(j-1)^2, j^2]\log n$, define a sequence of intervals,

$$I_j^i := \frac{i}{n} + \frac{L_j}{n}[-1,1].$$

Then for any $i \in S_j$ and $p_x \in I_j^i \cap I_j$, the corresponding Poisson probability satisfies

$$\begin{aligned} e^{-np_x}\frac{(np_x)^i}{i!} &= e^{-i}\frac{i^i}{i!}\cdot\exp\left(-(np_x - i) + i\cdot\log\left(1+\frac{np_x-i}{i}\right)\right) \\ &\geq \frac{1}{3\sqrt{i}}\cdot\exp\left(-\frac{2i}{3}\cdot\left(\frac{np_x-i}{i}\right)^2\right) \\ &\geq \frac{1}{9\sqrt{i}} \geq \frac{1}{9L_j}. \end{aligned}$$

On the other hand, the following upper bound holds.

$$\begin{aligned} e^{-np_x}\frac{(np_x)^i}{i!} &= e^{-i}\frac{i^i}{i!}\cdot\exp\left(-(np_x - i) + i\cdot\log\left(1+\frac{np_x-i}{i}\right)\right) \\ &\leq e^{-i}\frac{i^i}{i!} \leq \frac{1}{\sqrt{2\pi i}} \leq \frac{1}{2L_j}. \end{aligned}$$

In other words, for any $p_x, i/n \in I_j$ that differ by at most L_j/n,

$$\Pr(\mathrm{Poi}(np_x) = i) \in \frac{1}{L_j}\left[\frac{1}{9}, \frac{1}{2}\right].$$

Partition I_j into sub-intervals of equal length L_j/n. The partition has a size of at most $2\sqrt{\log n}$. Assign each probability $p_x \in I_j$ a length-L_j/n interval I_{p_x} centered at p_x. Then, each interval I_{p_x} covers at least one of the sub-intervals in the partition. Since there are exactly p_{I_j} intervals I_{p_x}, one can find a partition sub-interval I_j^s contained in at least $p_{I_j}/(2\sqrt{\log n})$ of them. Denote by $\mathcal{X}_s$ the collection of symbols corresponding to these intervals.

Next, we bound from below the entropy of the truncated profile $(\varphi_i)_{i\in nI_j^s}$ over nI_j^s. Denote by j_s the left end point of nI_j^s. By the chain rule of entropy for multiple random variables,

$$H((\varphi_i)_{i\in nI_j^s}) = \sum_{i=j_s}^{j_s+L_j-1} H(\varphi_i|\varphi_{j_s},\ldots,\varphi_{i-1}).$$

Consider a particular term on the right-hand side with $i \in [j_s, j_s + L_j - 1]$. By the conditional independence and fact that conditioning reduces entropy,

$$\begin{aligned} H(\varphi_i|\varphi_{j_s}, \ldots, \varphi_{i-1}) &\geq H(\varphi_i|\varphi_{j_s}, \ldots, \varphi_{i-1}; \mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \in \mathcal{X}) \\ &= H(\varphi_i|\mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \in \mathcal{X}) \\ &= H(\varphi_i|\mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \in \mathcal{X}_s; \mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \notin \mathcal{X}_s) \end{aligned}$$

To characterize the condition, we define a random variable

$$K_i^s := \sum_{x \in \mathcal{X}_s} \mathbb{1}_{j_s \leq \mu_x \leq i-1}.$$

Note that $\mathbb{E}[\mathbb{1}_{j_s \leq \mu_x \leq i-1}] = \sum_{t=j_s}^{i-1} \Pr(\mathrm{Poi}(np_x) = t) \leq (i - j_s)/(2L_j)$, which is at most $1/10$ for $i \leq j_s + L_j/5$. The following lemma transforms this into a high-probability statement.

Lemma 3. *Let $Y_i, i \in [1, m]$ be independent indicator random variables. Let $Y := \sum_i Y_i$ denote their sum and $\lambda := \mathbb{E}[Y]$ denote the expected sum. Then for $c > 0$, we have*

$$\Pr(Y \geq \lambda(1 + c)) \leq \exp(-\lambda c^2/(2 + 2c/3)).$$

Below we consider only $i \leq j_s + L_j/5$. Note that $c/(2 + 2c/3)$ is increasing for $c > 0$.

Since $\mathbb{E}[K_i^s] = \sum_{x \in \mathcal{X}_s} \mathbb{E}[\mathbb{1}_{j_s \leq \mu_x \leq i-1}] \leq |\mathcal{X}_s|/10$,

$$\Pr(K_i^s \geq |\mathcal{X}_s|/2) \leq \exp(-36/35) < 1/2.$$

where we set $c = 4$ in the above lemma and assume that $|\mathcal{X}_s| \geq 3$ (assuming only $|\mathcal{X}_s| \geq 1$, the upper bound becomes $3/4$). Recall that

$$\begin{aligned} H(\varphi_i|\varphi_{j_s}, \ldots, \varphi_{i-1}) &\geq H(\varphi_i|\mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \in \mathcal{X}_s; \mathbb{1}_{j_s \leq \mu_x \leq i-1}, x \notin \mathcal{X}_s) \\ &= \sum_{(c_x)_{x \in \mathcal{X}} \in \{0,1\}^{\mathcal{X}}} H(\varphi_i|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s) \\ &\qquad \times \Pr(\mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s). \end{aligned}$$

Denote by $V_s \subseteq \{0,1\}^{\mathcal{X}}$ the collection of $(c_x)_{x \in \mathcal{X}}$ satisfying $\sum_{x \in \mathcal{X}_s} c_x < |\mathcal{X}_s|/2$. The above derivation shows that

$$\sum_{(c_x)_{x \in \mathcal{X}} \in V_s} \Pr(\mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s) \geq \frac{1}{2}.$$

By independence, for any $(c_x)_{x \in \mathcal{X}} \in V_s$, we have

$$\begin{aligned} (\varphi_i|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s) &= \sum_{x \in \mathcal{X}: c_x = 0} (\mathbb{1}_{\mu_x = i}|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = 0) \\ &= \sum_{x \in \mathcal{X}_s: c_x = 0} (\mathbb{1}_{\mu_x = i}|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = 0) \\ &\qquad + \sum_{x \notin \mathcal{X}_s: c_x = 0} (\mathbb{1}_{\mu_x = i}|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = 0). \end{aligned}$$

For any $x \in \mathcal{X}_s$ with $c_x = 0$, the corresponding indicator variable satisfies

$$\begin{aligned} \mathbb{E}[\mathbb{1}_{\mu_x = i}|\mathbb{1}_{j_s \leq \mu_x \leq i-1} = 0] &= \frac{\Pr(\mathbb{1}_{\mu_x = i} \text{ and } \mu_x \notin [j_s, i-1])}{\Pr(\mu_x \notin [j_s, i-1])} \\ &= \frac{\Pr(\mathbb{1}_{\mu_x = i})}{1 - \Pr(\mu_x \in [j_s, i-1])} \\ &= \frac{\frac{1}{L_j}\left[\frac{1}{9}, \frac{1}{2}\right]}{1 - \left[0, \frac{L_j}{5}\right] \cdot \frac{1}{L_j}\left[\frac{1}{9}, \frac{1}{2}\right]} \\ &= \frac{1}{L_j}\left[\frac{1}{9}, \frac{5}{9}\right]. \end{aligned}$$

On the other hand, for any $x \notin \mathcal{X}_s$,

$$e^{-np_x}\frac{(np_x)^i}{i!} \le e^{-i}\frac{i^i}{i!} \le \frac{1}{\sqrt{2\pi i}} \le \frac{1}{2L_j}.$$

Therefore, the corresponding indicator variable satisfies

$$\mathbb{E}[\mathbb{1}_{\mu_x=i}|\mathbb{1}_{j_s\le\mu_x\le i-1}=0] = \frac{\Pr(\mathbb{1}_{\mu_x=i})}{1-\Pr(\mu_x\in[j_s,i-1])} \le \frac{\frac{1}{L_j}\left[0,\frac{1}{2}\right]}{1-\left[0,\frac{L_j}{5}\right]\cdot\frac{1}{L_j}\left[0,\frac{1}{2}\right]} \le \frac{5}{9}\cdot\frac{1}{L_j}.$$

To summarize, we have shown that $(\varphi_i|\mathbb{1}_{j_s\le\mu_x\le i-1}=c_x, x\in\mathcal{X}_s)$ is the sum of $|\mathcal{X}|$ independent Bernoulli random variables. Among these Bernoulli variables, at least $|\mathcal{X}_s|/2 \ge p_{I_j}/(2\sqrt{\log n})$ have a bias of $\frac{1}{L_j}\left[\frac{1}{9},\frac{5}{9}\right]$, while others have a bias of at most $\frac{5}{9}\cdot\frac{1}{L_j}$.

The following lemma, recently established by Hillion et al. [2019], shows the relation among the entropy values of sums of independent Bernoulli random variables with different bias parameters.

Lemma 4. *Let $X_t, Y_t, t\in[m]$ be independent indicator random variables. Denote by X and Y the sums of X_t's and Y_t's, respectively. If $\mathbb{E}[X_t]\le\mathbb{E}[Y_t]\le 1/2, \forall t\in[m]$,*

$$H(\sum_t X_t) \le H(\sum_t Y_t).$$

This lemma, together with the previous results, shows that

$$H(\varphi_i|\mathbb{1}_{j_s\le\mu_x\le i-1}=c_x, x\in\mathcal{X}_s) \ge H(\mathrm{bin}(p_{I_j}/(2\sqrt{\log n}), 1/(9L_j)).$$

The next lemma further bounds the entropy of a binomial random variable.

Lemma 5. *For any $m>1$ and $q\in[1/m, 1-1/m]$,*

$$H(\mathrm{bin}(m,q)) \ge \frac{1}{2}\log\left((2\pi)^{1-(1-q)^m-q^m}mq(1-q)\right) - \frac{1}{12m}.$$

Proof. By definition, the left-hand side satisfies

$$\begin{aligned}
H(\mathrm{bin}(m,q)) &= -\sum_{t=0}^{m}\binom{m}{t}q^t(1-q)^{m-t}\log\left(\binom{m}{t}q^t(1-q)^{m-t}\right)\\
&= -\sum_{t=0}^{m}\binom{m}{t}q^t(1-q)^{m-t}(t\log q + (m-t)\log(1-q)\\
&\qquad\qquad + \log m! - \log t! - \log(m-t)!)\\
&= mH(\mathrm{Bern}(q)) - \log m! + \sum_{t=0}^{m}\binom{m}{t}q^t(1-q)^{m-t}(\log t! + \log(m-t)!).
\end{aligned}$$

By Stirling's formula, for any $t\ge 1$,

$$\log t! \ge \left(t+\frac{1}{2}\right)\log t + \frac{1}{2}\log(2\pi) - t.$$

Substituting the right-hand side into the above equation yields

$$\begin{aligned}
S_m(q) := \sum_{t=0}^{m}\binom{m}{t}q^t(1-q)^{m-t}\log t! \ge \frac{1}{2}(1-(1-q)^m)\log(2\pi) - mq\\
+ \sum_{t=1}^{m}\binom{m}{t}q^t(1-q)^{m-t}\left(t+\frac{1}{2}\right)\log t.
\end{aligned}$$

Let $g(x):=0$ for $x\in[0,1)$ and $g(x):=(x+1/2)\log x$ for $x\ge 1$. Simple calculus shows that the function is concave. Applying the concavity of g to the last sum yields

$$\sum_{t=1}^{m}\binom{m}{t}q^t(1-q)^{m-t}\left(t+\frac{1}{2}\right)\log t \ge g\left(\sum_{t=0}^{m}\binom{m}{t}q^t(1-q)^{m-t}\cdot t\right) = \left(mq+\frac{1}{2}\right)\log(mq),$$

where the last step follows by the fact that $mq \geq 1$. A similar inequality holds for the weighted sum of $\log(m-t)!$. Consolidating these inequalities, we obtain

$$\begin{aligned} S_m(q) + S_m(1-q) &\geq \left(mq + \frac{1}{2}\right)\log(mq) + \left(m(1-q) + \frac{1}{2}\right)\log(m(1-q)) \\ &\quad + \frac{1}{2}(1-(1-q)^m)\log(2\pi) - mq + \frac{1}{2}(1-q^m)\log(2\pi) - m(1-q) \\ &= (m+1)\log m - mH(\mathrm{Bern}(q)) + \frac{1}{2}\log(q(1-q)) \\ &\quad + \frac{1}{2}(2-(1-q)^m - q^m)\log(2\pi) - m. \end{aligned}$$

On the other hand, for the $\log m!$ term,

$$\log m! \leq \left(m + \frac{1}{2}\right)\log m + \frac{1}{2}\log(2\pi) - m + \frac{1}{12m}.$$

Substituting the previous term bounds into the $H(\mathrm{bin}(m,q))$ expression yields

$$\begin{aligned} H(\mathrm{bin}(m,q)) &= mH(\mathrm{Bern}(q)) - \log m! + S_m(q) + S_m(1-q) \\ &\geq \frac{1}{2}\log\left((2\pi)^{1-(1-q)^m-q^m} mq(1-q)\right) - \frac{1}{12m}. \end{aligned}$$

□

Before continuing, we remark that the bound in the above lemma has the right dependence on $mq(1-q)$ in the sense that if we fix q and increase m, the lower bound converges to $\frac{1}{2}\log(\Theta(mq(1-q)))$. Another point to mention is that the above bound covers $q \in [1/m, 1-1/m]$, while Lemma 6 appearing later in this section covers $q \notin [1/m, 1-1/m]$. Note that the dependence on $mq(1-q)$ changes from logarithmic to linear, showing an "elbow effect" around $1/m$.

Assume that $p_{I_j}/(2\sqrt{\log n}) \geq 9L_j$, then for any $(c_x)_{x\in\mathcal{X}} \in V_s$,

$$H(\varphi_i | \mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s) \geq H(\mathrm{bin}(p_{I_j}/(2\sqrt{\log n}), 1/(9L_j)) \geq \frac{1}{2}.$$

Consolidating this with the previous results yields that

$$H(\varphi_i | \varphi_{j_s}, \ldots, \varphi_{i-1}) \geq \sum_{(c_x)_{x\in\mathcal{X}} \in V_s} \frac{1}{2} \cdot \Pr(\mathbb{1}_{j_s \leq \mu_x \leq i-1} = c_x, x \in \mathcal{X}_s) \geq \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4},$$

where we utilize $p_{I_j}/(2\sqrt{\log n}) \geq 9L_j \geq 9$ and $(1-q)^m + q^m < 1/e$ for $\forall m \geq 3, q \in [1/m, 1/2]$. We can then bound the quantity of interest as follows.

$$\begin{aligned} H((\varphi_i)_{i\in nI_j^s}) &= \sum_{i=j_s}^{j_s+L_j-1} H(\varphi_i | \varphi_{j_s}, \ldots, \varphi_{i-1}) \\ &\geq \sum_{i=j_s}^{j_s+L_j/5} H(\varphi_i | \varphi_{j_s}, \ldots, \varphi_{i-1}) \\ &\geq \frac{L_j}{5} \cdot \frac{1}{4} = \frac{L_j}{20} \\ &= \frac{1}{20\sqrt{\log n}} \min\left\{p_{I_j}, j \cdot \log n\right\}. \end{aligned}$$

On the other hand, if $9L_j \geq p_{I_j}/(2\sqrt{\log n}) \gg 1$, we can further "compress" the truncated profile $(\varphi_i)_{i\in nI_j^s}$ over nI_j^s to reduce the effective value of L_j. Specifically, for any integer $t < L_j$, we define the t-compressed version of $(\varphi_i)_{i\in nI_j^s}$ as

$$(\varphi_i)^t_{i\in nI_j^s} := \left(\sum_{i=j_s+(\ell-1)t}^{j_s+\ell t-1} \varphi_i\right)_{\ell \in [L_j/t]}.$$

Note that for each t, the length of $(\varphi_i)_{i\in nI_j^s}^t$ is $L_j^t := L_j/t$. For each entry in the compressed version, we can again express the entry as the sum of independent indicator random variables. Specifically,

$$\sum_{i=j_s+(\ell-1)t}^{j_s+\ell t-1} \varphi_i = \sum_{x\in\mathcal{X}} \mathbb{1}_{\mu_x\in[j_s+(\ell-1)t,j_s+\ell t-1]}.$$

Furthermore, for any $x \in \mathcal{X}_s$, the expectation of each indicator variable satisfies

$$\begin{aligned}\mathbb{E}[\mathbb{1}_{\mu_x\in[j_s+(\ell-1)t,j_s+\ell t-1]}] &= \sum_{i=j_s+(\ell-1)t}^{j_s+\ell t-1} e^{-np_x}\frac{(np_x)^i}{i!}\\ &= \frac{t}{L_j}\left[\frac{1}{9},\frac{1}{2}\right] = \frac{1}{L_j^t}\left[\frac{1}{9},\frac{1}{2}\right].\end{aligned}$$

Similarly, for any $x \in \mathcal{X}$, we have $\mathbb{E}[\mathbb{1}_{\mu_x\in[j_s+(\ell-1)t,j_s+\ell t-1]}] \le 1/(2L_j^t)$.

Now, choose t large enough so that $18L_j^t \ge p_{I_j}/(2\sqrt{\log n}) \ge 9L_j^t$. Following the reasoning in the previous case shows that

$$H((\varphi_i)_{i\in nI_j^s}) \ge H((\varphi_i)_{i\in nI_j^s}^t) \ge \Omega\left(\frac{1}{\sqrt{\log n}}\min\left\{p_{I_j}, j\cdot\log n\right\}\right).$$

It remains to consider the case of $\mathcal{O}(\sqrt{\log n}) \ge p_{I_j} \ge 1$, for which we adopt our previous analysis.

Again, partition I_j into sub-intervals of equal length L_j/n. Then, assign each probability $p_x \in I_j$ a length-L_j/n interval I_{p_x} centered at p_x. By construction, each interval I_{p_x} covers at least one of the sub-intervals in the partition. Redefine any of these covered sub-intervals as I_j^s. Denote by $\mathcal{X}_s$ the collection of symbols corresponding to the covering intervals.

Note that $\mathcal{O}(\sqrt{\log n}) \ge p_{I_j} \ge |\mathcal{X}_s| \ge 1$. For any $i \in [j_s, j_s + L_j/5]$, the previous analysis shows that

$$H(\varphi_i|\varphi_{j_s},\ldots,\varphi_{i-1}) \ge H(\mathrm{bin}(|\mathcal{X}_s|, 1/(9L_j))\cdot(1-3/4).$$

We bound the right-hand side with the following lemma.

Lemma 6. *For any $m \ge 1$, and $q \le \min\{1/2, 1/m\}$ or $q \ge \max\{1/2, 1-1/m\}$,*

$$H(\mathrm{bin}(m,q)) \ge \frac{m}{4}\min\{q, 1-q\} \ge \frac{1}{4}mq(1-q).$$

Proof. By symmetry, we need to consider only the case of $q \in [0, 1/m]$.

$$\begin{aligned}H(\mathrm{bin}(m,q)) &\ge H(\mathbb{1}_{\mathrm{bin}(m,q)\ge 1})\\ &= H(((1-q)^m, 1-(1-q)^m))\\ &\ge -(1-q)^m(m\log(1-q))\\ &\ge -\frac{m}{4}\log(1-q)\\ &\ge \frac{m}{4}\cdot q.\end{aligned}$$

□

Consolidating the lemma and the chain rule of entropy yields,

$$\begin{aligned}H((\varphi_i)_{i\in nI_j^s}) &= \sum_{i=j_s}^{j_s+L_j-1} H(\varphi_i|\varphi_{j_s},\ldots,\varphi_{i-1})\\ &\ge \sum_{i=j_s}^{j_s+L_j/5} H(\varphi_i|\varphi_{j_s},\ldots,\varphi_{i-1})\\ &\ge \frac{L_j}{5}\cdot\frac{|\mathcal{X}_s|}{4\cdot 9\cdot L_j}\cdot\left(1-\frac{3}{4}\right) = \frac{|\mathcal{X}_s|}{720}\\ &= \Omega\left(\frac{1}{\sqrt{\log n}}\min\left\{p_{I_j}, j\cdot\log n\right\}\right).\end{aligned}$$

Alternatively, we can use the fact that adding independent random variables does not decrease entropy, i.e., $H(Y+Z) \geq H(Y)$ for any independent variables Y and Z. Note that

$$(\varphi_i)_{i\in nI_j^s}^t = \sum_{x\in\mathcal{X}} (\mathbb{1}_{\mu_x=i})_{i\in I_j^s}.$$

Let y be an arbitrary symbol that belongs to $\mathcal{X}_s$. Then,

$$H((\varphi_i)_{i\in nI_j^s}) \geq H((\varphi_i)_{i\in nI_j^s}^t) \geq H((\mathbb{1}_{\mu_y=i})_{i\in I_j^s}) \geq H((\mathbb{1}_{\mu_y=j_s}, \mathbb{1}_{\mu_y=j_s+1})).$$

By the previous derivations, both $\Pr(\mu_y = j_s)$ and $\Pr(\mu_y = j_s+1)$ belong to $\frac{1}{L_j}[1/9, 1/2]$. Hence,

$$H((\varphi_i)_{i\in nI_j^s}) \geq H\left(\mathrm{Bern}\left(\frac{2}{11}\right)\right) \geq \frac{2}{5} = \Omega\left(\frac{1}{\sqrt{\log n}} \min\left\{p_{I_j}, j\cdot\log n\right\}\right).$$

Note that this argument does not apply to other cases, since

$$H((\mathbb{1}_{\mu_y=i})_{i\in I_j^s}) = \mathcal{O}(\log L_j) = \mathcal{O}(\log n),$$

while $\min\left\{p_{I_j}, j\cdot\log n\right\}$ can be as large as $\tilde{\Theta}(n^{1/3})$ in general.

The proof is complete upon noting that indices with $j = \mathcal{O}(1)$ corresponds to a total contribution of at most $\mathcal{O}(1)$ to $H_n^{\mathcal{S}}(p)$ and $H_n^{\mathcal{S}}(p) = \tilde{\Theta}(\mathbb{E}[\mathcal{D}(\varphi)]) = \tilde{\Theta}(D(\varphi))$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$.

Summary The simple expression shows that $H_n^{\mathcal{S}}(p)$ characterizes the variability of ranges the actual probabilities spread over. As Theorem 8 shows, $H_n^{\mathcal{S}}(p)$ closely approximates $E_n(p)$, the value around which $\mathcal{D}_n \sim p$ concentrates (Theorem 6) and $H(\Phi^n)$ lies (Thoerem 1). Henceforth, we use $H_n^{\mathcal{S}}(p)$ as a proxy for both $H(\Phi^n)$ and $\mathcal{D}_n$, and study its attributes and values.

A.3 Symmetric Property Estimation and Sufficiency of Profiles

The rest of Section A shows that the PML plug-in estimator possesses the amazing ability of adapting to the simplicity of data distributions in inferring all symmetric properties, over any label-invariant classes. For clarity, we divide the full proof into three parts: a) the sufficiency of profiles for estimating symmetric properties; b) the standard "median trick" often used to boost the confidence of learning algorithms; c) the PML method and its competitiveness to the min-max estimators. The proof utilizes several previously established results.

Sufficiency of profiles We first show that profile-based estimators are sufficient for estimating symmetric properties. Recall that a distribution collection $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$ is *label-invariant* if for any $p \in \mathcal{P}$, the collection $\mathcal{P}$ contains all its symbol-permuted versions. Then,

Theorem 2. *Let f be a symmetric functional over a label-invariant distribution collection $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$. For any accuracy $\varepsilon > 0$ and tolerance $\delta \in (0,1)$, if there exists an estimator $\hat{f}$ such that*

$$\Pr_{X^n \sim p}\left(\left|\hat{f}(X^n) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P},$$

there is an estimator $\hat{f}_\varphi$ over Φ satisfying

$$\Pr_{X^n \sim p}\left(\left|\hat{f}_\varphi(\varphi(X^n)) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P}.$$

Note that both estimators can have independent randomness.

Proof. First we show that given estimator $\hat{f}$, there is an estimator $\hat{f}_s$ which is symmetric, i.e., invariant with respect to domain-symbol permutations, and achieves the same guarantee. To see this, consider a random permutation $\tilde{\sigma}$ chosen uniformly randomly from the collection of permutations over the underlying alphabet. Let $\hat{f}_s := \hat{f} \circ \tilde{\sigma}$. Then for any $p \in \mathcal{P}$,

$$\begin{aligned}
\Pr_{X^n \sim p}\left(\left|\hat{f}_s(X^n) - f(p)\right| > \varepsilon\right) &\overset{(a)}{=} \Pr_{X^n \sim p}\left(\left|\hat{f} \circ \tilde{\sigma}(X^n) - f(p)\right| > \varepsilon\right) \\
&\overset{(b)}{=} \sum_{\sigma} \Pr_{X^n \sim p}\left(\left|\hat{f} \circ \sigma(X^n) - f(p)\right| > \varepsilon \,\middle|\, \tilde{\sigma} = \sigma\right) \cdot \Pr\left(\tilde{\sigma} = \sigma\right) \\
&\overset{(c)}{=} \sum_{\sigma} \Pr_{X^n \sim p}\left(\left|\hat{f} \circ \sigma(X^n) - f(p)\right| > \varepsilon\right) \cdot \Pr\left(\tilde{\sigma} = \sigma\right) \\
&\overset{(d)}{=} \sum_{\sigma} \Pr_{X^n \sim \sigma(p)}\left(\left|\hat{f}(X^n) - f(\sigma(p))\right| > \varepsilon\right) \cdot \Pr\left(\tilde{\sigma} = \sigma\right) \\
&\overset{(e)}{<} \sum_{\sigma} \delta \cdot \Pr\left(\tilde{\sigma} = \sigma\right) \\
&\overset{(f)}{=} \delta,
\end{aligned}$$

where (a) follows by the definition of $\hat{f}_s$; (b) follows by the law of total probability; (c) follows by the independence between $\tilde{\sigma}$ and X^n; (d) follows by the symmetry of f and the equivalence of applying σ to X^n and to p; (e) follows by the fact that $\sigma(p) \in \mathcal{P}$ and the guarantee satisfied by the estimator $\hat{f}$; and (f) follows by the law of total probability.

Before we proceed further, we introduce the following definitions. For any sequence x^n, the *sketch* of a symbol x in x^n is the set of indices $i \in [n]$ for which $x_i = x$. The *type* of a sequence x^n is the set $\tau(x^n)$ of sketches of symbols appearing in x^n.

Since $\hat{f}_s$ is symmetric, there exists a mapping $\hat{f}_\tau$ over types satisfying $\hat{f}_s = \hat{f}_\tau \circ \tau$. Due to the i.i.d. assumption on the sample generation process, given the profile of a sample sequence, all the different types corresponding to this profile are equally likely. Let Λ be a mapping that recovers this relation, i.e., Λ maps each profile uniformly randomly to a type having this profile.

Then, for any $p \in \mathcal{P}$ and $X^n \sim p$,

$$\hat{f}_s(X^n) = \hat{f}_\tau \circ \tau(X^n) = \hat{f}_\tau \circ \Lambda \circ \varphi(X^n).$$

Consequently, the mapping $\hat{f}_\varphi := \hat{f}_\tau \circ \Lambda$ is a profile-based estimator that satisfies

$$\Pr_{X^n \sim p}\left(\left|\hat{f}_\varphi(\varphi(X^n)) - f(p)\right| > \varepsilon\right) = \Pr_{X^n \sim p}\left(\left|\hat{f}_s(X^n) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P}. \qquad \square$$

A.4 Median Trick

The following argument is standard and often used to boost the confidence of learning algorithms.

Lemma 7 (Median trick)**.** *Let $\alpha, \beta \in (0,1)$ be real parameters satisfying $1/10 \geq \alpha > \beta$. For an accuracy $\varepsilon > 0$ and a distribution set $\mathcal{P} \subseteq \Delta_\mathcal{X}$, if there exists an estimator $\hat{f}_A$ such that*

$$\Pr_{X^n \sim p}\left(\left|\hat{f}_A(X^n) - f(p)\right| > \varepsilon\right) < \alpha,\ \forall p \in \mathcal{P},$$

we can construct another estimator $\hat{f}_B$ that takes a sample of size $m := \left\lceil \frac{4n}{\log \frac{1}{2\alpha}} \log \frac{1}{\beta} \right\rceil$ and achieves

$$\Pr_{Y^m \sim p}\left(\left|\hat{f}_B(Y^m) - f(p)\right| > \varepsilon\right) < \beta,\ \forall p \in \mathcal{P}.$$

Proof. Given $t \in \mathbb{N}$ i.i.d. copies of $\hat{f}_A(X^n)$, the probability that less than half of them satisfy the inequality in the parentheses is at least

$$\Pr\left(\sum_{i=1}^{t} \mathbb{1}_{A_i} < \frac{t}{2} \text{ for } A_i\text{'s satisfying } \Pr(A_i) < \alpha\right) \geq \Pr\left(\mathrm{bin}\,(t,\alpha) < \frac{t}{2}\right).$$

By the law of total probability, the right-hand side equals to

$$\begin{aligned} 1 - \Pr\left(\mathrm{bin}\,(t,\alpha) \geq \frac{t}{2}\right) &\geq 1 - \exp\left(\left(\left(\frac{1}{2\alpha} - 1\right) - \frac{1}{2\alpha}\log\frac{1}{2\alpha}\right) \cdot \alpha t\right) \\ &\geq 1 - \exp\left(-\frac{t}{4}\log\frac{1}{2\alpha}\right), \end{aligned}$$

where the first step follows by the Chernoff bound of binomial random variables, and the second step follows by $\alpha \leq 1/10$ and the inequality $c - 1 - \frac{c}{2}\log c > 0, \forall c \geq 5$.

Set $t := \left\lceil \frac{4}{\log \frac{1}{2\alpha}} \log \frac{1}{\beta} \right\rceil$, the right-hand side is at least $1 - \beta$.

Therefore, given a sample of size $m = t \cdot n$, we can partition it into t sub-samples of equal size, apply the estimator $\hat{f}_A$ to each subsample, and define the median of the corresponding estimates as $\hat{f}_B$.

By the previous reasoning, this estimator satisfies

$$\Pr_{Y^m \sim p}\left(\left|\hat{f}_B(Y^m) - f(p)\right| > \varepsilon\right) < \beta,\ \forall p \in \mathcal{P}. \qquad \square$$

A.5 Profile Maximum Likelihood and Its Adaptiveness

For every profile $\boldsymbol{\phi}$ of length n and distribution collection $\mathcal{P} \subseteq \Delta_\mathcal{X}$, the *profile maximum likelihood* (PML) estimator Orlitsky et al. [2004] over $\mathcal{P}$ maps $\boldsymbol{\phi}$ to a distribution

$$\mathcal{P}_{\boldsymbol{\phi}} := \arg\min_{p \in \mathcal{P}} \Pr_{X^n \sim p}\left(\varphi(X^n) = \boldsymbol{\phi}\right),$$

that maximizes the probability of observing the profile $\boldsymbol{\phi}$.

For any property f, let $\varepsilon_f(n, \delta, \mathcal{P})$ denote the smallest error that can be achieved by any estimator with a sample size n and tolerance δ on the error probability. This definition is equivalent to that of the sample complexity. In the following, we show that the PML estimator is adaptive to the simplicity of underlying distributions in inferring all symmetric properties, over any label-invariant $\mathcal{P}$.

For brevity, set $\delta = 1/10$ and suppress both δ and $\mathcal{P}$ in $\varepsilon_f(n, \delta, \mathcal{P})$.

Theorem 3 (Adaptiveness of PML)**.** *Let f be a symmetric property and $\mathcal{P} \subseteq \Delta_{\mathcal{X}}$ be a label-invariant distribution collection. For any $p \in \mathcal{P}$ and $\Phi^n \sim p$, with probability at least $1-\mathcal{O}(1/\sqrt{n})$,*

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\varepsilon_f\left(\frac{\tilde{\Omega}(n)}{\lceil H(\Phi^n)\rceil}\right).$$

Proof. For any tolerance $\delta \in (0,1)$ and distribution $p \in \Delta_{\mathcal{X}}$, define the (δ, n)-*typical cardinality of profiles with respect to* p as the smallest cardinality $C_{\delta,n}(p)$ of a set of length-n profiles such that the probability of observing a sample from p with a profile in this set is at least $1-\delta$. The following lemma provides a tight characterization of $C_{\delta,n}(p)$ in terms of the dimension of $\Phi^n \sim p$.

Lemma 8. *For any $p \in \Delta_{\mathcal{X}}$ and $\Phi^n \sim p$, with probability at least $1-6/\sqrt{n}$,*

$$C_{\frac{6}{\sqrt{n}},n}(p) \le n^{8(\mathcal{D}(\Phi^n)+20\log n)}.$$

The proof of the lemma follows by recursively applying Theorem 6. Specifically, let $d := 2E_n(p) + 3\log n$, which is at least $\mathcal{D}_n \sim p$, with probability at least $1-6/\sqrt{n}$. Then,

$$C_{\frac{6}{\sqrt{n}},n}(p) \le \binom{n}{d}\binom{n+d-1}{d-1} \le n^{2d-1} \le n^{2(2E_n(p)+3\log n)} \le n^{8\mathcal{D}(\Phi^n)+20\log n},$$

where the last inequality holds with with probability at least $1-6/\sqrt{n}$.

Now let f be a symmetric functional over $\mathcal{P}$. According to Theorem 2, for any parameters $\varepsilon > 0$ and $\delta \in (0,1)$, if there exists an estimator $\hat{f}$ such that

$$\Pr_{X^n\sim p}\left(\left|\hat{f}(X^n) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P},$$

there is an estimator $\hat{f}_\varphi$ over Φ satisfying

$$\Pr_{X^n\sim p}\left(\left|\hat{f}_\varphi(\varphi(X^n)) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P}.$$

For an arbitrary length-n profile $\boldsymbol{\phi}$ that satisfies $\Pr_{\Phi^n\sim p}(\Phi^n = \boldsymbol{\phi}) \ge 2\delta$, these error bounds yield

$$\Pr\left(\left|\hat{f}_\varphi(\boldsymbol{\phi}) - f(p)\right| > \varepsilon\right) < \frac{1}{2}.$$

and since $\Pr_{\Phi^n\sim \mathcal{P}_{\boldsymbol{\phi}}}(\Phi^n = \boldsymbol{\phi}) \ge \Pr_{\Phi^n\sim p}(\Phi^n = \boldsymbol{\phi}) \ge 2\delta$ by the definition of PML,

$$\Pr\left(\left|\hat{f}_\varphi(\boldsymbol{\phi}) - f(\mathcal{P}_{\boldsymbol{\phi}})\right| > \varepsilon\right) < \frac{1}{2}.$$

By the union bound and triangle inequality,

$$\Pr\left(|f(p) - f(\mathcal{P}_{\boldsymbol{\phi}})| > 2\,\varepsilon\right) < 1 \iff |f(p) - f(\mathcal{P}_{\boldsymbol{\phi}})| \le 2\,\varepsilon \ \text{ surely.}$$

Furthermore, by Lemma 8, with probability at least $1-6/\sqrt{n}$, the total probability of length-n profiles $\boldsymbol{\phi}$ satisfying $\Pr_{\Phi^n\sim p}(\Phi^n = \boldsymbol{\phi}) < 2\delta$ is at most

$$2\delta \cdot C_{\frac{6}{\sqrt{n}},n}(p) + \frac{6}{\sqrt{n}} \le 2\delta \cdot n^{8\mathcal{D}(\Phi^n)+20\log n} + \frac{6}{\sqrt{n}},$$

which basically upperly bounds the probability that

$$|f(p) - f(\mathcal{P}_{\Phi^n})| > 2\,\varepsilon\,.$$

Next we will assume that there exists an estimator $\hat{f}$ satisfying

$$\Pr_{X^m\sim p}\left(\left|\hat{f}(X^m) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P}.$$

By Lemma 7, if $\delta \le 1/10$, we can construct another estimator $\hat{f}'$ that takes a sample of size $n = \frac{4m}{\log\frac{1}{2\delta}}\log\frac{1}{\delta'}$ (n is assumed to be an integer here) and achieves a higher-confidence guarantee

$$\Pr_{X^n\sim p}\left(\left|\hat{f}'(X^n) - f(p)\right| > \varepsilon\right) < \delta',\ \forall p \in \mathcal{P}.$$

Then by the above reasoning, with probability at least $1 - 6/\sqrt{n}$,

$$\Pr_{\Phi^n \sim p}\left(|f(p) - f(\mathcal{P}_{\Phi^n})| > 2\,\varepsilon\right) \le 2\delta' \cdot n^{8\mathcal{D}(\Phi^n)+20\log n} + \frac{6}{\sqrt{n}}$$
$$= 2\exp\left(-\frac{n}{4m}\log\frac{1}{2\delta} + (8\mathcal{D}(\Phi^n) + 20\log n)\log n\right) + \frac{6}{\sqrt{n}}.$$

For the first term on the right hand side to vanish as quickly as $1/\sqrt{n}$, it suffices to have

$$\frac{n}{4m}\log\frac{1}{2\delta} \ge 20\cdot\mathcal{D}(\Phi^n)\log n \iff \frac{n}{\mathcal{D}(\Phi^n)\log n} \ge 80\cdot\frac{m}{\log\frac{1}{2\delta}},$$

and simultaneously have

$$\frac{n}{4m}\log\frac{1}{2\delta} \ge 40\cdot\log^2 n \iff \frac{n}{\log^2 n} \ge 160\cdot\frac{m}{\log\frac{1}{2\delta}}.$$

Simplify the expressions and apply the union bound. It suffices to have both

$$\frac{\tilde{\Theta}(n)}{\mathcal{D}(\Phi^n)} \ge \frac{m}{\log\frac{1}{\delta}} \text{ and } n \ge 8m.$$

If n satisfies these conditions, with probability at least $1 - \Theta(1/\sqrt{n})$,

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\,\varepsilon\,.$$

Summary We have shown the following result, which is a strengthened version of Theorem 3. The result shows that the PML plug-in estimator possesses the amazing ability of adapting to the simplicity of data distributions in inferring all symmetric properties, over any label-invariant classes.

If there exists an estimator $\hat{f}$ such that

$$\Pr_{X^m \sim p}\left(\left|\hat{f}(X^m) - f(p)\right| > \varepsilon\right) < \delta,\ \forall p \in \mathcal{P},$$

for any $p \in \mathcal{P}$ and $\Phi^n \sim p$ where the sample size n satisfies both

$$\frac{\tilde{\Theta}(n)}{\mathcal{D}(\Phi^n)} \ge \frac{m}{\log\frac{1}{\delta}} \text{ and } n \ge 8m,$$

with probability at least $1 - \Theta(1/\sqrt{n})$,

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\,\varepsilon\,.$$

Alternative statement

Fix $\mathcal{P}$ and assume that $\delta \le 1/10$. Recall that $\varepsilon_f(n,\delta)$ denotes the smallest error that can be achieved by the best estimator using a size-n sample with a $(1-\delta)$-confidence guarantee. Below we provide an alternative statement of the above result, which is more compact in its form.

Then, draw a profile $\Phi^n \sim p$. With probability at least $1 - \Theta(1/\sqrt{n})$,

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\varepsilon_f\left(\frac{\tilde{\Theta}(n)\log\frac{1}{\delta}}{\mathcal{D}(\Phi^n)}\Lambda\ \frac{n}{8}, \delta\right).$$

Consolidating this with Theorem 1 yields that

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\varepsilon_f\left(\frac{\tilde{\Theta}(n)\log\frac{1}{\delta}}{\lceil H(\Phi^n)\rceil}\Lambda\ \frac{n}{8}, \delta\right).$$

Setting $\delta = 1/10$ and suppressing it in expressions, we establish the desired guarantee: For any $p \in \mathcal{P}$ and $\Phi^n \sim p$, with probability at least $1 - \mathcal{O}(1/\sqrt{n})$,

$$|f(p) - f(\mathcal{P}_{\Phi^n})| \le 2\varepsilon_f\left(\frac{\tilde{\Omega}(n)}{\lceil H(\Phi^n)\rceil}\right). \qquad \square$$

Below are some comments in order.

1. The theorem holds for any symmetric properties, while nearly all previous works require the property to possess certain forms and be smooth.
2. The theorem trivially implies a weaker result in Acharya et al. [2017] where $\lceil H(\Phi^n) \rceil$ is replaced by $\mathcal{O}(\sqrt{n})$, an upper bound due to the formula of Hardy and Ramanujan [1918].
3. There is a polynomial-time approximation algorithm Charikar et al. [2019a] achieving the same guarantee as that stated in the theorem.

A.6 PML and Sorted Distribution Estimation

The arguments below basically follow by Hao and Orlitsky [2019a] and Hao and Orlitsky [2019d], in which we assume that $|\mathcal{X}| = \mathcal{O}(n \log n)$.

Let f be a function in $\mathcal{L}_1$, the collection of Lipschitz functions over $[0, 1]$. 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

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

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 an integer 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 $z_j := \eta \cdot j/s, j \in \lceil s \rfloor$ where $\lceil s \rfloor := \{0, 1, \ldots, s\}$. Then, for each $j \in \lceil s \rfloor$, 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 $|j^*| \in \lceil j \rfloor$. 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,

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

Therefore, for any $p \in \mathcal{P} := \Delta_\mathcal{X}$, the corresponding properties of f_η and $\tilde{f}_\eta$ satisfy

$$|f_\eta(p) - \tilde{f}_\eta(p)| \le |\mathcal{X}| \cdot \frac{\eta}{s},$$

where we slightly abuse the notation and write $\tilde{f}_\eta(p) := \sum_x \tilde{f}_\eta(p_x)$. Note that $|j^*| \in \lceil j \rfloor$ for all $j \in \lceil s \rfloor$, and $\tilde{f}_\eta(z) = z_{s^*}$ for $z \ge \eta$. While there are infinitely many η-truncations, the cardinality of the discretized η-truncations of functions in $\mathcal{L}_1$ is at most

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

Now we consider the task of estimating $\tilde{f}_\eta(p)$ from $X^n \sim p$. By construction, the real function $\tilde{f}_\eta(z)$ is a constant for $z \ge \eta$. In addition, the function is Lipschitz and hence for an absolute constant C and an arbitrary interval $I := [a, b] \subseteq [0, 1]$, one can construct an explicit polynomial $g(z)$ of degree at most $d \in \mathbb{N}$, satisfying

$$|g(z) - \tilde{f}_\eta(z)| \le C\frac{\sqrt{|b-a|(x-a)}}{d}, \forall x \in I.$$

Combining these facts with Theorem 5 in Hao and Orlitsky [2019d] shows that there exists an estimator $\hat{f}_\eta(X^n)$ that for all $p \in \Delta_\mathcal{X}$ and ε satisfying $n = \Omega(|\mathcal{X}|/(\varepsilon^2 \log |\mathcal{X}|))$ and $\varepsilon > 1/n$,

$$\Pr_{X^n \sim p}\left(\left|\hat{f}_\eta(X^n) - \tilde{f}_\eta(p)\right| > \varepsilon\right) \le \exp(-\Theta(\varepsilon^2/(n^\lambda \eta))),$$

where $\lambda \in (0, 1)$ is an absolute constant bounded away from 0 whose value can be arbitrarily small.

Then, by setting $\eta \le \mathcal{O}(\varepsilon^2/n^{1/2+\lambda})$, where the asymptotic notation hides a sufficiently small absolute constant, the right-hand side is at most $\exp(-4\sqrt{n})$. Then, Theorem 2 in this paper and Theorem 3 in Acharya et al. [2017] imply that the PML distribution $\mathcal{P}_{\varphi(X^n)}$ satisfies

$$\Pr_{X^n \sim p}\left(\left|\tilde{f}_\eta(\mathcal{P}_{\varphi(X^n)}) - \tilde{f}_\eta(p)\right| > 2\varepsilon\right) \le \exp(-\sqrt{n}), \forall p \in \Delta_\mathcal{X}.$$

Consider any $p \in \Delta_{\mathcal{X}}$ and $X^n \sim p$ with a profile $\varphi := \varphi(X^n)$. Consolidate the previous results, and apply the union bound and triangle inequality. With probability at least $1 - \exp\left(3s \log s - \sqrt{n}\right)$, the PML plug-in estimator will satisfy

$$|f_\eta(p) - f_\eta(p_\varphi)| \le |f_\eta(p) - \tilde{f}_\eta(p)| + |\tilde{f}_\eta(p) - \tilde{f}_\eta(p_\varphi)| + |\tilde{f}_\eta(p_\varphi) - f_\eta(p_\varphi)| \le 2|\mathcal{X}| \cdot \frac{\eta}{s} + 2\varepsilon,$$

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 $\bar{f}_\gamma \in \mathcal{L}_1$. To establish the corresponding guarantees, we make use of the following result. Since the profile probability is invariant to symbol permutation, for our purpose, we can assume that $p(y) \le p(z)$ iff $\mathcal{P}_\varphi(x) \le \mathcal{P}_\varphi(y)$, for all $x, y \in \mathcal{X}$. Under this assumption, the following lemma relates $\mathcal{P}_\varphi$ to p.

Lemma 9. *For any $\eta \in (0, \mathcal{O}(1/\sqrt{n}))$, distribution $p \in \Delta_{\mathcal{X}}$ and sample $X^n \sim p$ with profile φ,*

$$\Pr\left(\sum_x |\max\{\mathcal{P}_\varphi(x), \eta\} - \max\{p(x), \eta\}| > \Theta\left(\sqrt{\frac{1}{\eta n}}\right)\right) \le \exp(-\sqrt{n}).$$

The proof of this lemma follows from 1) the fact that empirical distribution satisfies such a guarantee with the probability bound being $\exp(-4\sqrt{n})$, where we ignore labelings and sort the empirical probabilities according to p; 2) the Cauchy-Schwarz inequality applied to bound the expected error of the empirical estimator and McDiarmid's inequality used to bound the error probability; 3) a variant of Theorem 3 in Acharya et al. [2017] that addresses distribution estimation Das [2012].

Hence, with probability at least $1 - \exp(-\sqrt{n})$,

$$\begin{aligned}
|\bar{f}_\eta(p) - \bar{f}_\eta(\mathcal{P}_\varphi)| &= |\sum_x \bar{f}_\eta(p(x)) - \bar{f}_\eta(\mathcal{P}_\varphi(x))| \\
&\le \sum_x |\bar{f}_\eta(\max\{p(x), \eta\}) - \bar{f}_\eta(\max\{\mathcal{P}_\varphi(x), \eta\})| \\
&\le \sum_x |\max\{\mathcal{P}_\varphi(x), \eta\} - \max\{p(x), \eta\}| \\
&\le \Theta\left(\sqrt{\frac{1}{\eta n}}\right),
\end{aligned}$$

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

Consolidate the previous results. By the triangle inequality and the union bound, with probability at least $1 - \exp\left(3s \log s - \sqrt{n}\right) - \exp(-\sqrt{n})$,

$$\begin{aligned}
|f(p) - f(p_\varphi)| &\le |f_\eta(p) - f_\eta(\mathcal{P}_\varphi)| + |\bar{f}_\eta(p) - \bar{f}_\eta(\mathcal{P}_\varphi)| \\
&\le 2|\mathcal{X}| \cdot \frac{\eta}{s} + 2\varepsilon + \Theta\left(\sqrt{\frac{1}{\eta n}}\right),
\end{aligned}$$

for *all* functions f in $\mathcal{L}_1$. Now we can conclude that $\ell_1^<(p, p_\varphi)$ 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

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

It remains to balance the error bounds on the estimation and deviation probability. Recall that we assume $|\mathcal{X}| \le \mathcal{O}(n \log n)$ since otherwise the theorem is trivial to prove. Set $s = \tilde{\Theta}(\sqrt{n})$ such that $3s \log s < \sqrt{n}/2$. Then, the confidence lower bound becomes $1 - \exp(\sqrt{n}/2) - \exp(\sqrt{n})$, and the deviation bound reduces to $\tilde{\mathcal{O}}(\sqrt{n}\eta) + \Theta(\sqrt{1/(\eta n)}) + 2\varepsilon$. The previous derivations also require that $\eta \le \mathcal{O}(\varepsilon^2/n^{1/2+\lambda})$ and $\eta \in (0, \mathcal{O}(1/\sqrt{n}))$. Setting $\eta = \Theta(1/n^{3/4})$ yields the desired result.

A.7 PML and Uniformity Testing

For a finite domain $\mathcal{X}$, denote by $u_{\mathcal{X}}$ the uniform distribution over $\mathcal{X}$. Given an error parameter $\varepsilon > 0$ and a sample X^n from an unknown distribution $p \in \Delta_{\mathcal{X}}$, *uniformity testing* Goldreich and Ron [2011] aims to distinguish between the null hypothesis

$$p = u_{\mathcal{X}},$$

and the alternative hypothesis

$$\|p - u_{\mathcal{X}}\|_1 > \varepsilon.$$

In the work of Hao and Orlitsky [2019a], it is shown that the following simple PML-based algorithm achieves the optimal $\Theta(\sqrt{|\mathcal{X}|}/\varepsilon^2)$ sample complexity Paninski [2008] for uniformity testing, up to logarithmic factors of the alphabet size $|\mathcal{X}|$. Note that we instantiate the distribution collection $\mathcal{P}$ as $\Delta_{\mathcal{X}}$, and use 0 and 1 to indicate whether H_0 or H_1 is accepted.

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

Figure 1: Uniformity tester T_{PML}

In this section, we present another intriguing connection between the PML estimator and the uniformity testing problem. For any profile φ of length n, denote

$$T(\varphi) = \frac{|\mathcal{X}|(\sum_{\mu=1}^{n}\varphi_\mu \cdot \mu^2 - n)}{n^2 - n}.$$

Then for any accuracy $\varepsilon > 0$, the following uniformity tester Diakonikolas et al. [2016a] is sample optimal up to logarithmic factors.

- If $T(\varphi(X^n)) \geq 1 + 3\varepsilon^2/4$, return 1;
- Else, return 0.

The following lemma connects the above algorithm to the PML.

Lemma 10. *Chan et al. [2015] For any profile $\varphi := \varphi(x^n)$ that corresponds to a non-constant sequence $x^n \in \mathcal{X}^*$,*

- *If $T(\varphi) > 1$, then $u_{\mathcal{X}}$ is a local minimum of the PML optimization problem*

$$\mathcal{P}_{\mathcal{X}} = \max_{p \in \mathcal{P} = \Delta_{\mathcal{X}}} \Pr_{Y^n \sim p}(\varphi(Y^n) = \varphi);$$

- *Else, $u_{\mathcal{X}}$ is a local maximum.*

B Attributes of Profile Entropy and Dimension

Let $p \in \Delta_{\mathcal{X}}$ be an arbitrary discrete distribution. Recall that in Section A, we partition the unit interval into a sequence of ranges,

$$I_j := \left((j-1)^2\frac{\log n}{n},\ j^2\frac{\log n}{n}\right], 1 \le j \le \sqrt{\frac{n}{\log n}},$$

denote by p_{I_j} the number of probabilities p_x belonging to I_j, and relate $E_n(p)$ to an induced shape-reflecting quantity,

$$H_n^{\mathcal{S}}(p) := \sum_{j\ge 1} \min\left\{p_{I_j}, j\cdot\log n\right\},$$

the sum of the effective number of probabilities lying within each range.

The simple expression of $H_n^{\mathcal{S}}(p)$ shows that it characterizes the variability of ranges the actual probabilities spread over. As Theorem 8 shows, $H_n^{\mathcal{S}}(p)$ closely approximates $E_n(p)$, the value around which $\mathcal{D}_n \sim p$ concentrates (Theorem 6) and $H(\Phi^n)$ lies (Thoerem 1). In this section, we use $H_n^{\mathcal{S}}(p)$ as a proxy for both $H(\Phi^n)$ and $\mathcal{D}_n$, and study its attributes and values.

To further our understanding of profile entropy and dimension, we investigate the analytical attributes of $H_n^{\mathcal{S}}(p)$ concerning monotonicity and Lipschitzness. Then, we present tight upper and lower bounds on the value of $H_n^{\mathcal{S}}(p)$ for a variety of distribution families.

B.1 Monotonicity

Among the many attributes that $H_n^{\mathcal{S}}(p)$ possesses, monotonicity is perhaps most intuitive. One may expect a larger value of $H_n^{\mathcal{S}}(p)$ as the sample size n increases, since additional observations reveal more information on the variability of probabilities. Below we confirm this intuition.

Theorem 9. *For any $n \ge m \gg 1$ and $p \in \Delta_{\mathcal{X}}$,*

$$H_n^{\mathcal{S}}(p) \ge H_m^{\mathcal{S}}(p).$$

Besides the above result that lowerly bounds $H_n^{\mathcal{S}}(p)$ with $H_m^{\mathcal{S}}(p)$ for $m \le n$, a more desirable result is to upperly bound $H_n^{\mathcal{S}}(p)$ with a function of $H_m^{\mathcal{S}}(p)$. Such a result will enable us to draw a sample of size $m \le n$, obtain an estimate of $H_m^{\mathcal{S}}(p)$ from $\mathcal{D}_m$, and use it to bound the value of $H_n^{\mathcal{S}}(p)$ and thus of $\mathcal{D}_n$ for a much larger sample size n.

With such an estimate, we can perform numerous tasks such as *predicting* the performance of PML in estimating symmetric properties when more observations are available, and the space needed for storing a longer sample profile. These applications are closely related to the recent works on *learnability estimation* Kong and Valiant [2018], Kong et al. [2019].

The next theorem provides a simple and tight upper bound on $H_n^{\mathcal{S}}(p)$ in terms of $H_m^{\mathcal{S}}(p)$.

Theorem 10. *For any $n \ge m \gg 1$ and $p \in \Delta_{\mathcal{X}}$,*

$$H_n^{\mathcal{S}}(p) \le \sqrt{\frac{n\log n}{m\log m}}\cdot H_m^{\mathcal{S}}(p).$$

Implications Before proceeding to the proof, we first present two simple implications.

1. If for $m = \Omega(n^{0.01})$, we have $H_m^{\mathcal{S}}(p) \ll \sqrt{m}$, then $H_n^{\mathcal{S}}(p) \ll \sqrt{n}$.
2. For any two integers $m \le n$ and distribution p,

$$\frac{H_m^{\mathcal{S}}(p)}{\sqrt{m\log m}} \ge \frac{H_n^{\mathcal{S}}(p)}{\sqrt{n\log n}}.$$

   In other words, the sequence $A_m := H_m^{\mathcal{S}}(p)/\sqrt{m\log m}$, $m \le n$, is monotonically decreasing and converges to A_n. As we increase the value of m, $(\sqrt{n\log n}\cdot A_m)$, which can be viewed as our estimate of $H_n^{\mathcal{S}}(p)$, is getting more and more accurate. For the purpose of adaptive estimation, if $n = 2^t$, we can choose the sequence $m = 2^0, 2^1, \ldots, 2^t$.

Proof. For clarity, we denote by $p(m,j)$ the value of p_{I_j} corresponding to $H_m^{\mathcal{S}}(p)$, and $p(n,j)$ the value of p_{I_j} corresponding to $H_n^{\mathcal{S}}(p)$. Furthermore, denote $r := \sqrt{(n/m)((\log m)/\log n)}$, which we assume is an integer. Then by the definition of $H_\cdot^{\mathcal{S}}$,

$$\begin{aligned}
rH_m^{\mathcal{S}}(p) &= r\sum_{j\geq 1}\min\left\{p(m,j), j\cdot\log m\right\} \\
&= \sum_{j\geq 1}\min\left\{r\cdot\sum_{i=rj-r+1}^{rj}p(n,i),\ rj\cdot\log m\right\} \\
&\geq \sum_{j\geq 1}\sum_{t=0}^{r-1}\min\left\{\sum_{i=rj-r+1}^{rj}p(n,i),\ (rj-t)\cdot\log m\right\} \\
&\geq \sum_{j\geq 1}\sum_{t=0}^{r-1}\min\left\{p(n,rj-t),\ (rj-t)\cdot\log m\right\} \\
&= \sum_{i\geq 1}\min\left\{p(n,i),\ i\cdot\log m\right\} \\
&\geq \frac{\log m}{\log n}\cdot H_n^{\mathcal{S}}(p).
\end{aligned}$$

The lower-bound part basically follows by reversing the above inequalities.

$$\begin{aligned}
H_n^{\mathcal{S}}(p) &= \sum_{i\geq 1}\min\left\{p(n,i),\ i\cdot\log n\right\} \\
&= \sum_{j\geq 1}\sum_{t=0}^{r-1}\min\left\{p(n,rj-t),\ (rj-t)\cdot\log n\right\} \\
&\geq \sum_{j\geq 1}\sum_{t=0}^{r-1}\min\left\{p(n,rj-t),\ (rj-r+1)\cdot\log n\right\} \\
&\geq \sum_{j\geq 1}\min\left\{\sum_{t=0}^{r-1}p(n,rj-t),\ (rj-r+1)\cdot\log n\right\} \\
&= \sum_{j\geq 1}\min\left\{p(m,j),(rj-r+1)\cdot\log m\right\} \\
&\geq H_m^{\mathcal{S}}(p).
\end{aligned}$$

This completes the proof of the theorem. □

B.2 Lipschitzness

Viewing $H_n^{\mathcal{S}}(p)$ as a distribution property, we establish its Lipschitzness with respect to a weighted Hamming distance and the ℓ_1 distance. Given two distributions $p, q \in \Delta_{\mathcal{X}}$, the vanilla *Hamming distance* is denoted by

$$h(p,q) := \sum_{x\in\mathcal{X}}\mathbb{1}_{p_x\neq q_x}.$$

The distance is suitable for being a statistical distance since there may be many symbols at which the two distributions differ, yet those symbols account for only a negligible total probability and has little effects on many induced statistics. To address this, we propose a *weighted Hamming distance*

$$h_{\mathcal{W}}(p,q) := \sum_{x\in\mathcal{X}}\max\{p_x,q_x\}\cdot\mathbb{1}_{p_x\neq q_x}.$$

The next result measures the Lipschitzness of $H_n^{\mathcal{S}}$ under $h_{\mathcal{W}}$.

Theorem 11. *For any integer n, and distributions p and q, if $h_{\mathcal{W}}(p,q) \le \varepsilon$ for some $\varepsilon \ge 1/n$,*

$$\left|H_n^{\mathcal{S}}(p) - H_n^{\mathcal{S}}(q)\right| \le \tilde{\mathcal{O}}(\sqrt{\varepsilon n}).$$

Proof. Recall that the quantity of interest is

$$H_n^{\mathcal{S}}(p) := \sum_{j\ge 1} \min\left\{p_{I_j}, j\cdot\log n\right\}.$$

Given the bound of $h_{\mathcal{W}}(p,q) \le \varepsilon$, we denote by $\mathcal{Y} \subseteq \mathcal{X}$ the collection of symbols x at which $p_x \ne q_x$. By definition, we have both $\sum_{x\in\mathcal{Y}} p_x \le \varepsilon$ and $\sum_{x\in\mathcal{Y}} q_x \le \varepsilon$. Below, we show that these symbols modify the value of $H_n^{\mathcal{S}}(p)$ by at most $\tilde{\mathcal{O}}(\sqrt{\varepsilon n})$. By symmetry, the same claim also holds for the distribution q. Combined, these two claims yields the desired result.

First, we consider $x \in \mathcal{Y}$ satisfying $p_x = 0$ or $p_x \in I_1 = (0, (\log n)/n]$. Such a symbol either does not contribute the value of $H_n^{\mathcal{S}}(p)$, or affects only the value of the first term $\min\{p_{I_1}, \log n\}$, which is at most $\log n$. Hence the claim holds for this case.

Next, consider symbols $x \in \mathcal{Y}$ satisfying $p_x \in I_j = ((j-1)^2\frac{\log n}{n},\ j^2\frac{\log n}{n}]$ for some $j \ge 2$ and denote the collection of them by $\mathcal{Z} \subseteq \mathcal{Y}$. By the above assumption, we have $\sum_{x\in\mathcal{Z}} p_x \le \varepsilon$. To maximize their impact on $H_n^{\mathcal{S}}(p)$ under this constraint, we should set their values to be

$$p_j := (j-1)^2\frac{\log n}{n},\ j = 2,\dots J,$$

for some J to be determined, where each p_j repeats exactly $j\log n$ times. Then, the symbols in $\mathcal{Z}$ contributes at most $\sum_{j=2}^{J} j\log n = (\log n)(J-1)(J+2)/2$ to $H_n^{\mathcal{S}}(p)$, and the above constraint on the total probability mass bounds transforms to

$$\varepsilon \ge \sum_{x\in\mathcal{Z}} p_x \ge \sum_{j=2}^{J}(j\log n)\cdot(j-1)^2\frac{\log n}{n} \ge \frac{(\log n)^2}{12n}J(J^2-1)(-2+3J).$$

Therefore in this case, the contribution is again $\tilde{\mathcal{O}}(\sqrt{\varepsilon n})$, which completes the proof. □

Replacing $\max\{p_x, q_x\}$ with $|p_x - q_x|$ results in a common similarity measure – the ℓ_1 distance. The next theorem is an analog to the above under this classical distance.

Theorem 12. *For any integer n, and distributions p and q, if $\ell_1(p,q) \le \varepsilon$ for some $\varepsilon \ge 0$,*

$$\left|H_n^{\mathcal{S}}(p) - cH_n^{\mathcal{S}}(q)\right| \le \mathcal{O}((\varepsilon n)^{2/3}),$$

where c is a constant in $[1/3, 3]$. Note that the inequality is significant iff $\varepsilon \le \tilde{\Theta}(1/n^{1/4})$, since the value of $H_n^{\mathcal{S}}(p)$ is at most $\mathcal{O}(\sqrt{n\log n})$ for all p.

By symmetry, it suffices to prove the following lemma.

Lemma 11. *For any integer n, and distributions p and q, if $\ell_1(p,q) \le \varepsilon$ for some $\varepsilon \ge 0$,*

$$H_n^{\mathcal{S}}(p) \le 3H_n^{\mathcal{S}}(q) + \mathcal{O}((\varepsilon n)^{2/3}).$$

Proof. Consider the task of modifying p by at most ε and maximizing the increase in $H_n^{\mathcal{S}}(p)$. For each j and each probability $p_x \in j$, denote by p'_x the modified value. Depending on the location of p'_x, there are three types of possible modifications as illustrated below.

- For the first type, we still have $p'_x \in I_j$. This does not change the value of p_{I_j} and hence does not increase $H_n^{\mathcal{S}}(p)$.
- For the second type, we have $p'_x \in I_{j-1}$ or $p'_x \in I_{j+1}$. If $p_{I_j} \le j\cdot\log n$, this will decrease the value of $\min\{p_{I_j}, j\cdot\log n\}$ by 1 and increase the value of $\min\{p_{I_{j-1}}, (j-1)\cdot\log n\}$ or $\min\{p_{I_{j+1}}, (j+1)\cdot\log n\}$ by at most one. Hence in this case, the value of $H_n^{\mathcal{S}}(p)$ can only decrease. If $p_{I_j} > j\cdot\log n$, then $\min\{p_{I_j}, j\cdot\log n\} = j\cdot\log n$. For a particular j, all such modifications can increase the value of $H_n^{\mathcal{S}}(p)$ by at most $(j-1)\log n + (j+1)\log n = 2j\log n$, which is twice the value of $\min\{p_{I_j}, j\cdot\log n\}$. Hence, all such modifications, when combined, increase the value of $H_n^{\mathcal{S}}(p)$ by at most $2H_n^{\mathcal{S}}(p)$.

- For the third type, we have $p'_x \in I_i$ and $|i-j| \geq 2$. If $i < j$, we require a probability mass of at least $((j-1)^2 \log n - i^2 \log n)/n \geq (i \log n)/n$, where $j \geq 3$. If $i > j$, we require a probability mass of at least $((i-1)^2 \log n - j^2 \log n)/n \geq (i \log n)/n$. The number of such modifications that could lead to an increase in the value of $H_n^{\mathcal{S}}(p)$ is at most $i \log n$. For each i, let c_i denote the number of such modifications that will lead to an increase of $H_n^{\mathcal{S}}(p)$. Then, the total increase is $\sum_i c_i$, each c_i is at most $i \log n$, and the total required probability mass required is at least $\sum_i c_i \cdot (i \log n)/n \leq \varepsilon$.

  Let $\{c_i\}$ be the optimal solution that maximizes $\sum_i c_i$. Assume that there are two indices $i < j$ satisfying $c_i < i \log n$ and $c_j > 0$. Then, if we replace c_i and c_j by $c_i + 1$ and $c_j - 1$, respectively, $\sum_i c_i$ will not change and $\sum_i c_i \cdot (i \log n)/n$ will decrease. Hence, we can assume that there exists i' satisfying $c_i = i \log n, \forall i < i'$ and $c_i = 0, \forall i > i'$. In addition, assuming $\varepsilon n \geq \log n$ implies that $i' \geq 2$. Hence, we have $\sum_i c_i \leq (\log n) i'(i'+1)/2$ and

$$\sum_i c_i \leq 3.5 \cdot \left(\frac{n\varepsilon}{\sqrt{\log n}}\right)^{2/3}. \qquad \square$$

Profile Entropy for Structured Families

Following the study of attributes of profile entropy, we derive below nearly tight bounds on the $H_n^{\mathcal{S}}(p)$ values of three important structured families, log-concave, power-law, and histogram. These bounds tighten up and significantly improve those in Hao and Orlitsky [2019c], and show the ability of profile entropy in charactering natural shape constraints.

For the remaining sections, we follow the convention and specify structured distributions over $\mathcal{X} = \mathbb{Z}$.

B.3 Log-Concave Distributions

We say a discrete distribution $p \in \Delta_{\mathbb{Z}}$ is *log-concave* if p has a contiguous support over $\mathbb{Z}$ and the inequality $p_x^2 \geq p_{x-1} p_{x+1}$ holds for all symbols $x \in \mathbb{Z}$.

The log-concave family encompasses a broad range of discrete distributions, such as Poisson, hyper-Poisson, Poisson binomial, binomial, negative binomial, geometric, and hyper-geometric, with wide applications to numerous research areas, including statistics Saumard and Wellner [2014], computer science Lovász and Vempala [2007], economics An [1997], algebra, and geometry Stanley [1989].

The next result upperly bounds the profile entropy of log-concave families, and is *tight* up to logarithmic factors of n.

Theorem 13. *For any $n \in \mathbb{Z}$ and distribution $p \in \Delta_{\mathbb{Z}}$, if p is log-concave and has a variance of σ^2,*

$$H_n^{\mathcal{S}}(p) \leq \mathcal{O}(\log n) \left(1 + \min\left\{\sigma, \frac{n}{\sigma}\right\}\right).$$

Proof. The $\mathcal{O}(\log n)(1+\sigma)$ upper bound is established in Hao and Orlitsky [2019c] using some concentration attributes of the log-concave distributions.

For the other component, we can assume that $\sigma \geq \sqrt{n}$ and n is larger than some absolute constant. Then by Diakonikolas et al. [2016b], the maximum probability $p_{\max}$ of p belongs to $[1/(8\sigma), 1/\sigma]$. Hence, the last index J for which $p_{I_J} \neq 0$ satisfies

$$(J-1)^2 \frac{\log n}{n} \leq \frac{1}{\sigma} \iff J \leq \sqrt{\frac{n}{\sigma \log n}} + 1.$$

Hence, we have

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\} \leq \log n + \sum_{j=1}^{\sqrt{n/(\sigma \log n)}+1} j \cdot \log n \leq \mathcal{O}(\log n)\left(1 + \frac{n}{\sigma}\right). \qquad \square$$

This upper bound is uniformly better than the $\min\{\sigma, (n^2/\sigma)^{1/3}\}$ bound in Hao and Orlitsky [2019c]. Theorem 17 further shows that it is optimal up to logarithmic factors of n.

A similar bound holds for t-mixtures of log-concave distributions. More concretely,

Theorem 14. *For any integer n and distribution $p \in \Delta_{\mathbb{Z}}$, if p is a t-mixture of log-concave distributions each has a variance of σ_i^2, where $i = 1, \ldots, t$,*

$$H_n^{\mathcal{S}}(p) \leq \mathcal{O}(\log n) \left(1 + \min\left\{\sum_i \sigma_i, \max_i \left\{\frac{n}{\sigma_i}\right\}\right\}\right).$$

B.4 Discretization of Continuous Distributions

The introduction about log-concave families covers numerous classical discrete distributions, yet leaves many more continuous ones untouched Bagnoli and Bergstrom [2005]. Below, we present a discretization procedure that preserves distribution shapes such as monotonicity, modality, and log-concavity. Applying this procedure to the Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ further shows the optimality of Theorem 13.

Let X be a continuous random variable with density function $f(x)$. For any $x \in \mathbb{R}$, denote by $\lceil x \rfloor$ the closest integer z such that $x \in (z - 1/2, z + 1/2]$. The distribution of $\lceil X \rfloor$ is over $\mathbb{Z}$ and satisfies

$$p(z) := \int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x)dx, \ \forall z \in \mathbb{Z}.$$

We refer to this random variable $\lceil X \rfloor$ as the discretized version of X.

Shape Preservation By definition, one can readily verify that the above transformation preserves several important shape characteristics of distributions, such as monotonicity, modality, and k-modality (possibly yields a smaller k). The following theorem covers log-concavity.

Theorem 15. *For any continuous random variable X over $\mathbb{R}$ with a log-concave density f, the distribution $p \in \Delta_{\mathbb{Z}}$ associated with $\lceil X \rfloor$ is also log-concave.*

To show this, we need the following basic lemma about concave functions.

Lemma 12. *For real numbers x_1, x_2, y_1, and y_2 satisfying $x_1 \leq x_2$, $y_1 \leq y_2$, $x_1 < y_1$, and $x_2 < y_2$,*

$$\frac{f(y_1) - f(x_1)}{y_1 - x_1} \geq \frac{f(y_2) - f(x_2)}{y_2 - x_2}.$$

Following the lemma, for $x, y \in \mathbb{R}$ such that $|x - y| \leq 1$, and any function f that is log-concave,

$$\log f(x+1) - \log f(x) \leq \log f(y) - \log f(y-1) \iff f(x+1)f(y-1) \leq f(x)f(y).$$

Proof. By definition p is log-concave if p has a consecutive support and $p(z)^2 \geq p(z+1)p(z-1), \forall z$. For $\lceil X \rfloor$, the first condition is satisfied since X is has a continuous support on $\mathbb{R}$, and $p(z)$ is positive as long as $f(x) > 0$ for a non-empty sub-interval of $(z - 1/2, z + 1/2]$.

Below we show that p satisfies the second condition. Specifically, for any $z \in \mathbb{Z}$,

$$\begin{aligned}
p(z-1)p(z+1) &= \left(\int_{z-\frac{3}{2}}^{z-\frac{1}{2}} f(x)dx\right)\left(\int_{z+\frac{1}{2}}^{z+\frac{3}{2}} f(x)dx\right) \\
&= \left(\int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x-1)dx\right)\left(\int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x+1)dx\right) \\
&= \int_{z-\frac{1}{2}}^{z+\frac{1}{2}}\int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x-1)f(y+1)dxdy \\
&\leq \int_{z-\frac{1}{2}}^{z+\frac{1}{2}}\int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x)f(y)dxdy \\
&= \left(\int_{z-\frac{1}{2}}^{z+\frac{1}{2}} f(x)dx\right)^2 \\
&= p(z)^2,
\end{aligned}$$

where the inequality follows by the above lemma and its implication. ∎

Moment preservation Denote by p the distribution of $\lceil X \rfloor$ for $X \sim f$. Let μ and σ^2 be the mean and variance of density f, given that they exist. The theorem below shows that distribution p has, within small additive absolute constants, a mean of μ and variance of $\Theta(\sigma^2)$.

Theorem 16. *Under the aforementioned conditions, the mean of $\lceil X \rfloor$ satisfies*

$$\mathbb{E}\lceil X \rfloor = \mu \pm \frac{1}{2},$$

and the variance of $\lceil X \rfloor$ satisfies

$$\sigma^2/2 - 1 \leq \mathbb{E}(\lceil X \rfloor - \mathbb{E}\lceil X \rfloor)^2 \leq 2\sigma^2 + 1.$$

Proof. First consider the mean value of $\lceil X \rfloor$ for $X \sim f$. We have

$$\mathbb{E}\lceil X \rfloor = \mathbb{E}[\lceil X \rfloor - X] + \mathbb{E}[X] = \mu \pm \frac{1}{2}.$$

Next consider the variance of $\lceil X \rfloor$. Applying the inequality $(a+b)^2 \leq 2(a^2+b^2)$ yields

$$\begin{aligned}
\mathbb{E}(\lceil X \rfloor - \mathbb{E}\lceil X \rfloor)^2 &= \int_{-\infty}^{\infty} (\lceil x \rfloor - \mathbb{E}\lceil X \rfloor)^2 \cdot f(x)dx \\
&= \int_{-\infty}^{\infty} (\lceil x \rfloor - x - (\lceil X \rfloor - \mathbb{E}X) + x - \mathbb{E}X)^2 \cdot f(x)dx \\
&\leq 2\int_{-\infty}^{\infty} \left((\lceil x \rfloor - x - (\lceil X \rfloor - \mathbb{E}X))^2 + (x - \mathbb{E}X)^2 \right) \cdot f(x)dx \\
&\leq 2\int_{-\infty}^{\infty} \left(1 + (x - \mathbb{E}X)^2 \right) \cdot f(x)dx \\
&= 2 + 2\mathbb{E}(X - \mathbb{E}X)^2 \\
&= 2(1+\sigma^2).
\end{aligned}$$

By the symmetry in the above reasoning, we also have

$$\sigma^2 = \mathbb{E}(X - \mathbb{E}X)^2 \leq 2(1 + \mathbb{E}(\lceil X \rfloor - \mathbb{E}\lceil X \rfloor)^2).$$

Consolidating these inequalities shows that

$$\sigma^2/2 - 1 \leq \mathbb{E}(\lceil X \rfloor - \mathbb{E}\lceil X \rfloor)^2 \leq 2\sigma^2 + 1.$$ □

B.5 Optimality of Theorem 13

By the above formula, the discretized Gaussian $\lceil \mathcal{N} \rfloor(\mu, \sigma^2)$ has a distribution in the form of

$$p_{G}(z) := \frac{1}{\sqrt{2\pi}\sigma} \int_{z-\frac{1}{2}}^{z+\frac{1}{2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2} \right) dx, \ \forall z \in \mathbb{Z}.$$

Consolidating Theorem 15 and 16 shows that p_{G} is a log-concave distribution with a variance of $\Theta(\sigma^2) \pm 1$. Consequently, Theorem 13 yields the following upper bound:

$$H_n^{\mathcal{S}}(p_{G}) \leq \mathcal{O}(\log n) \left(1 + \min\left\{ \sigma, \frac{n}{\sigma} \right\} \right).$$

In the following, we show that

Theorem 17. *Under the aforementioned conditions,*

$$H_n^{\mathcal{S}}(p_{G}) \geq \mathcal{O}(\log n)^{-1} \left(1 + \min\left\{ \sigma, \frac{n}{\sigma} \right\} \right).$$

The optimality of Theorem 13 follows by these inequalities.

Proof. At it is clear from the context, we will write p instead of p_{G}. Recall that

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\left\{ p_{I_j}, j \cdot \log n \right\},$$

where p_{I_j} denotes the number of probabilities belonging to $I_j = ((j-1)^2, j^2] \cdot (\log n)/n$. Considering part of the distribution can only reduce the value of $H_n^{\mathcal{S}}(p)$. Hence, we focus on the symbols in the range $(\mu+1, \infty) \cap \mathbb{Z}$, over which the probability mass function $p(z)$ is monotone.

We will further assume that $n/\log n \gg \sigma \gg \log n$, since otherwise the right-hand side of the inequality reduces to $\mathcal{O}(1)$, and the result follows by the fact that $H_n^{\mathcal{S}}(p) \geq 1$ for all n and p.

In addition, we focus on $j \gg 1$ in the following argument, as the contribution to $H_n^{\mathcal{S}}(p)$ from these indices is no more than the total $H_n^{\mathcal{S}}(p)$.

Given these assumptions, we have

$$\begin{aligned} p(z) \in I_j &\iff \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(z \pm 1/2 - \mu)^2}{2\sigma^2}\right) \in \left((j-1)^2 \frac{\log n}{n}, j^2 \frac{\log n}{n}\right] \\ &\iff z \pm 1/2 - \mu \in \sqrt{2}\sigma \left[\sqrt{c(\sigma,n) - 2\log j}, \sqrt{c(\sigma,n) - 2\log(j-1)}\right), \end{aligned}$$

where $c(\sigma, n) := \log\left(n/(\sqrt{2\pi}\sigma \log n)\right)$ and the interval is well-defined iff

$$\begin{aligned} c(\sigma,n) \geq 2\log j &\iff \frac{n}{\sqrt{2\pi}\sigma \log n} \geq j^2 \\ &\iff \sqrt{\frac{n}{\sqrt{2\pi}\sigma \log n}} \geq j \\ &\impliedby \sqrt{\frac{n}{\sigma \log n}} \geq 2j. \end{aligned}$$

For clarity, we divide our analysis into two cases: $\sqrt{n} \geq \sigma \gg \log n$ and $n/\log n \gg \sigma > \sqrt{n}$.

For the first case and $j \leq \sqrt{\sigma/\log n}/2 \leq \sqrt{n/(\sigma \log n)}/2$, the length L_j of the above interval, which equals to p_{I_j} up to an additive slack of 2, satisfies

$$\begin{aligned} \frac{L_j}{\sqrt{2}\sigma} &= \sqrt{c(\sigma,n) - 2\log(j-1)} - \sqrt{c(\sigma,n) - 2\log j} \\ &= \frac{2\log(j/(j-1))}{(c(\sigma,n) - 2\log(j-1)) + (c(\sigma,n) - 2\log j)} \\ &= \frac{\log(j/(j-1))}{\log\left(n/(\sqrt{2\pi}j(j-1)\sigma\log n)\right)} \\ &= \Omega\left(\frac{1}{\log n} \log\left(1 + \frac{1}{j-1}\right)\right) \\ &= \Omega\left(\frac{1}{j \log n}\right). \end{aligned}$$

Therefore, we have $L_j = \Omega(\sigma/(j \log n))$. Since $\sigma \gg \log n$ ensures $L_j \geq 3$ and $j \leq \sqrt{\sigma/\log n}/2$ is equivalent to $\sigma \geq 4j^2 \log n$, the lower bound on L_j transforms into $p_{I_j} \geq \Omega(j)$. Hence in this case, $H_n^{\mathcal{S}}(p)$ admits the following bound

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\} \geq \sum_{j=\mathcal{O}(1)}^{\sqrt{\sigma/\log n}/2} \Omega(j) = \Omega\left(\frac{\sigma}{\log n}\right).$$

In the $n/\log n \gg \sigma > \sqrt{n}$ case, we have $\sqrt{\sigma/\log n} > \sqrt{n/(\sigma \log n)}$. Repeating the previous reasoning for $j \leq \sqrt{n/(\sigma \log n)}/2$, we again obtain $L_j = \Omega\left(\sigma/(j \log n)\right)$ and $p_{I_j} \geq \Omega(j)$.

Therefore,

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\} \geq \sum_{j=\mathcal{O}(1)}^{\sqrt{n/(\sigma \log n)}/2} \Omega(j) = \Omega\left(\frac{n}{\sigma \log n}\right).$$

Finally, note that in the first case, $\min\{\sigma, n/\sigma\} = \sigma$, while in the second, $\min\{\sigma, n/\sigma\} = n/\sigma$.

Consolidating these results yields the desired lower bound

$$H_n^{\mathcal{S}}(p) \cdot \mathcal{O}(\log n) \geq 1 + \min\left\{\sigma, \frac{n}{\sigma}\right\}. \qquad \square$$

B.6 Power-Law Distributions

We say that a discrete distribution $p \in \Delta_{\mathbb{Z}}$ is a *power-law with power* $\alpha > 0$ if p has a support of $[k] := \{1, \dots, k\}$ for some $k \in \mathbb{Z}^+ \cup \{\infty\}$ and $p_x \propto x^{-\alpha}$ for all $x \in [k]$.

Power-law is a ubiquitous structure appearing in many situations of scientific interest, ranging from natural phenomena such as the initial mass function of stars Kroupa [2001], species and genera Humphries et al. [2010], rainfall Machado and Rossow [1993], population dynamics Taylor [1961], and brain surface electric potential Miller et al. [2009], to man-made circumstances such as the word frequencies in a text Baayen [2002], income rankings Drăgulescu and Yakovenko [2001], company sizes Axtell [2001], and internet topology Faloutsos et al. [1999].

Unlike log-concave distributions that concentrate around their mean values, power-laws are known to possess "long-tails" and always log-convex. Hence, one may expect the profile entropy of power-law distributions to behave differently from that of log-concave ones. The next theorem justifies this intuition and provides tight upper bounds.

Theorem 18. *For a power-law distribution $p \in \Delta_{[k]}$ with power α, we have*

$$H_n^{\mathcal{S}}(p) \le 7 \log n + e^2 \cdot \min\{k, \mathcal{U}_n^k(\alpha)\},$$

where

$$\mathcal{U}_n^k(\alpha) := \begin{cases} n^{\frac{1}{1+\alpha}} & \text{if } \alpha \ge 1 + \frac{1}{\log k}; \\ \left(\frac{n}{\log n}\right)^{\frac{1}{1+\alpha}} & \text{if } 1 \le \alpha < 1 + \frac{1}{\log k}; \\ \sqrt{n}\left(\frac{k}{\sqrt{n}} \wedge \left(\frac{\sqrt{n}}{k}\right)^{\frac{1-\alpha}{1+\alpha}}\right) & \text{if } 0 \le \alpha < 1. \end{cases}$$

The above upper bound fully characterizes the profile entropy of power-laws and surpasses the basic $\{k, \sqrt{n \log n}\}$ bound for both $k \gg \sqrt{n}$ and $k \ll \sqrt{n}$. In comparison, Hao and Orlitsky [2019c] yields a $\mathcal{O}(n^{\min\{1/(1+\alpha), 1/2\}})$ upper bound, which improves over $\sqrt{n \log n}$ only for $\alpha > 1$ and is worse than that above for all $\alpha < 1 + 1/\log k$.

Proof. For the ease of exposition, write the probability of symbol i assigned by distribution p as $p_i := c_\alpha^{-1} \cdot i^{-\alpha}$, where c_α is a normalizing constant (implicitly depends on k) and k can be infinite. Recall that the quantity of interest is

$$H_n^{\mathcal{S}}(p) = \sum_{j \ge 1} \min\left\{p_{I_j}, j \cdot \log n\right\}.$$

First consider p_{I_j} for a sufficiently large j and note that

$$\begin{aligned} p_i \in I_j &\iff \frac{1}{c_\alpha i^\alpha} \in \left((j-1)^2 \frac{\log n}{n}, j^2 \frac{\log n}{n}\right] \\ &\iff i \in I_j' := \left[\left(j^2 c(\alpha, n)\right)^{-\frac{1}{\alpha}}, \left((j-1)^2 c(\alpha, n)\right)^{-\frac{1}{\alpha}}\right), \end{aligned}$$

where $c(\alpha, n) := (c_\alpha \log n)/n$. Observe that the length L_j of interval I_j', which differs from the value of p_{I_j} by at most 2, is proportional to $(j-1)^{-2/\alpha} - j^{-2/\alpha}$, and hence is a decreasing function of j. Furthermore, each term $\min\{p_{I_j}, j \cdot \log n\} \approx \min\{L_j, j \cdot \log n\}$ is basically the minimum between this decreasing function and $j \log n$, an increasing function of j. This naturally calls for determining the value of j at which the two functions are equal. Concretely,

$$\left((j-1)^2 c(\alpha, n)\right)^{-\frac{1}{\alpha}} - \left(j^2 c(\alpha, n)\right)^{-\frac{1}{\alpha}} = j \log n \Longrightarrow j \ge \left(c(\alpha, n) \cdot (\log n)^\alpha\right)^{\frac{-1}{2+2\alpha}} \ge j - 1.$$

Let J denote the middle quantity on the right-hand side (implicitly depends on α and n). We can decompose the summation $H_n^{\mathcal{S}}(p)$ into two parts. The first part consists of indices $j \le J$,

$$H_{n,1}^{\mathcal{S}}(p) := \sum_{j=1}^{J-1} \min\left\{p_{I_j}, j \cdot \log n\right\} \le (\log n) \sum_{j=1}^{J} j = \frac{\log n}{2}(J+1)J \le \frac{1}{2}\left(\frac{c_\alpha}{n}\right)^{\frac{-1}{1+\alpha}}.$$

Correspondingly, the second part consists of indices $j \geq J$. For these indices j, we have $L_j \leq j \cdot \log n$. Recall that I'_j specifies the range of i satisfying $p_i \in I_j$. Then the second part satisfies

$$H^{\mathcal{S}}_{n,2}(p) := \sum_{j=J}^{n} \min\left\{p_{I_j}, j \cdot \log n\right\} \leq 7\log n + \sum_{j=J\vee 4}^{n} L_j \leq 7\log n + \left(1 - \frac{1}{J \vee 4}\right)^{-\frac{2}{\alpha}} \left(\frac{c_\alpha}{n}\right)^{\frac{-1}{1+\alpha}},$$

where the first inequality follows by the fact that the intervals I'_j are consecutive. Also note that the boundary case $j = J$ is covered in both $H^{\mathcal{S}}_{n,1}$ and $H^{\mathcal{S}}_{n,2}$ under different conditions. Depending on the value of the normalizing constant c_α, the following implications are immediate and apply to all power parameters $\alpha > 0$. If $c_\alpha \leq \tilde{\mathcal{O}}(n)\alpha^{1/(1+\alpha)}$,

$$H^{\mathcal{S}}_n(p) \leq H^{\mathcal{S}}_{n,1}(p) + H^{\mathcal{S}}_{n,2}(p) \leq 7\log n + e\left(\frac{n}{c_\alpha}\right)^{\frac{1}{1+\alpha}}.$$

If $c_\alpha \geq \tilde{\Omega}(n)$, then $H^{\mathcal{S}}_n(p)$ is bounded by $\log n$ since the probabilities p_i are at most $1/c_\alpha$.

These bounds can be refined given the knowledge of k. Note that for $\alpha \leq 1$, the support size k must be finite for the normalizing constant to be well-defined. On the other hand, for $\alpha > 1$, it is common to assume $k = \infty$, which we adopt here.

Then, for $\alpha > 1$, the above bound derivations yield

$$H^{\mathcal{S}}_n(p) = H^{\mathcal{S}}_{n,1}(p) + H^{\mathcal{S}}_{n,2}(p) \leq 7\log n + 2\left(\frac{n}{c_\alpha}\right)^{\frac{1}{1+\alpha}} \leq 7\log n + 2\left(n(\alpha-1)\right)^{\frac{1}{1+\alpha}},$$

where we lower bound c_α by $\max\{1, 1/(\alpha-1)\}$.

Next, we improve the upper bound for $\alpha < 1$. Note that the normalizing constant admits

$$\frac{k^{1-\alpha}}{1-\alpha} + \frac{\alpha}{1-\alpha} \geq 1 + \int_1^k x^{-\alpha}dx \geq c_\alpha = \sum_{i=1}^{k} i^{-\alpha} \geq \int_1^{k+1} x^{-\alpha}dx = \frac{(k+1)^{1-\alpha}}{1-\alpha} - \frac{1}{1-\alpha}.$$

Then for $k \geq \sqrt{n}$, the previous upper bound yields

$$H^{\mathcal{S}}_n(p) - 7\log n \leq e\left(\frac{n}{c_\alpha}\right)^{\frac{1}{1+\alpha}} \leq e\left(\frac{n(1-\alpha)}{(k+1)^{1-\alpha}-1}\right)^{\frac{1}{1+\alpha}} \leq e\left(\frac{en}{n^{\frac{1-\alpha}{2}}}\right)^{\frac{1}{1+\alpha}} \leq e^2\sqrt{n},$$

where we utilize the inequality $((k+1)^{1-\alpha}-1)/(1-\alpha) \geq (k+1)^{1-\alpha}/e$. Furthermore, one can bound $H^{\mathcal{S}}_n(p)$ by k since it is at most the sum of p_{I_j}. Combined, these two results yield

$$H^{\mathcal{S}}_n(p) \leq 7\log n + e^2\sqrt{n}\left(\min\left\{\frac{k}{\sqrt{n}}, \left(\frac{\sqrt{n}}{k}\right)^{\frac{1-\alpha}{1+\alpha}}\right\}\right).$$

To complete the picture, we consider the case of $\alpha = 1$. Note that $c_\alpha = \sum_{i=1}^{k} i^{-1} > \log k$. Hence for $\alpha = 1$, the above reasoning implies that

$$H^{\mathcal{S}}_n(p) \leq \min\left\{k, 2\sqrt{\frac{n}{\log k}}\right\}.$$

Finally, we also analyze the case of $\alpha > 1$ with truncation at k. Our analysis mainly relies on the following inequality that provides a tight lower bound on c_α.

$$c_\alpha = \sum_{i=1}^{k} i^{-\alpha} \geq \int_1^k x^{-\alpha}dx = \frac{1-k^{-(\alpha-1)}}{\alpha-1} \geq \begin{cases} \frac{\log k}{e} & \text{if } \alpha - 1 < \frac{1}{\log k}; \\ \frac{1-1/e}{\alpha-1} & \text{if } \alpha - 1 \geq \frac{1}{\log k}. \end{cases}$$

To summarize, for different values of α, the quantity $H^{\mathcal{S}}_n(p)$ admits

$$H^{\mathcal{S}}_n(p) - 7\log n \leq k \wedge e^2\left(n \cdot U(\alpha, k)\right)^{\frac{1}{1+\alpha}},$$

where

$$U(\alpha, k) := \begin{cases} \alpha - 1 & \text{if } \alpha \geq 1 + \frac{1}{\log k}; \\ \frac{1}{\log k} & \text{if } 1 \leq \alpha < 1 + \frac{1}{\log k}; \\ k^{\alpha-1} & \text{if } 0 \leq \alpha < 1. \end{cases}$$

Note that unless $\alpha < 1$ and $k \approx \sqrt{n}$, all the bounds are better than $\Theta(\sqrt{n})$. In addition, the derivation above already shows that the bounds are tight up to logarithmic factors. Reorganizing the terms yields the desired result: *For any power-law distribution p with power $\alpha \geq 0$,*

$$\frac{H_n^{\mathcal{S}}(p) - 7\log n}{e^2} \leq \begin{cases} n^{\frac{1}{1+\alpha}} & \text{if } \alpha \geq 1 + \frac{1}{\log k}; \\ \left(\frac{n}{\log n}\right)^{\frac{1}{1+\alpha}} & \text{if } 1 \leq \alpha < 1 + \frac{1}{\log k}; \\ \sqrt{n}\left(\frac{k}{\sqrt{n}} \wedge \left(\frac{\sqrt{n}}{k}\right)^{\frac{1-\alpha}{1+\alpha}}\right) & \text{if } 0 \leq \alpha < 1. \end{cases}$$

As a remark, for $\alpha = 0$, the distribution becomes a uniform distribution with support size k. The above result also covers this case, for which the upper bound simplifies to $k \wedge (n/k)$. □

B.7 Histogram Distributions

A distribution $p \in \Delta_{\mathcal{X}}$ is a *t-histogram distribution* if there is a partition of $\mathcal{X}$ into t parts such that p has the same probability value over all symbols in each part.

Besides the long line of research on histograms reviewed in Ioannidis [2003], the importance of histogram distributions rises with the rapid growth of data sizes in numerous engineering and science applications in the modern era.

For example, in scenarios where processing the complete data set is inefficient or even impossible, a standard solution is to partition/cluster the data into groups according to the task specifications and element similarities, and randomly sample from each group to obtain a subset of the data to use. This naturally induces a histogram distribution, with each data point being a symbol in the support.

The work of Hao and Orlitsky [2019c] studies the class of t-histogram distributions and obtains the following upper bound

$$H_n^{\mathcal{S}}(p) \leq \tilde{\mathcal{O}}\left(\min\left\{(nt^2)^{\frac{1}{3}}, \sqrt{n}\right\}\right).$$

Our contribution is establishing its optimality.

Theorem 19. *For any $t, n \in \mathbb{Z}^+$, there exists a t-histogram distribution p such that*

$$H_n^{\mathcal{S}}(p) \geq \tilde{\Omega}\left(\min\left\{(nt^2)^{\frac{1}{3}}, \sqrt{n}\right\}\right).$$

Note that uniform distributions correspond to 1-histograms, for which the bounds reduce to $\tilde{\Theta}(n^{1/3})$.

Proof. Again, recall that the quantity of interest is

$$H_n^{\mathcal{S}}(p) = \sum_{j \geq 1} \min\left\{p_{I_j}, j \cdot \log n\right\}.$$

Our construction depends on the value of t as follows. Let $A \cdot \{B\}$ denote the length-A constant sequence of value B. If $t = 1$, then the distribution p has the following form

$$p := \tilde{\Theta}(n^{1/3}) \cdot \{p_0 \in I_{n^{1/3}}\},$$

where p_0 is a properly chosen probability in $I_{n^{1/3}}$ so that p is well-defined, and the range of support of distribution p is irrelevant for our purpose and hence unspecified. If $2 \leq t < n^{1/4}/(2\sqrt{\log n})$, then for some parameter $s \geq 0$ to be determined, the distribution p has the following form

$$p := L \cdot \left\{\frac{1}{n^2}\right\} \bigcup \left(\bigcup_{j=s+1}^{s+t-1} \left((j \log n) \cdot \left\{j^2 \frac{\log n}{n}\right\}\right)\right),$$

where the probability values are sorted according to the ordering they appear above, and L is a properly chosen to make the probabilities sum to 1. For the distribution to be well-defined, we require

$$\sum_{j=s+1}^{s+t-1} (j \log n) \cdot \left(j^2 \frac{\log n}{n}\right) \leq 1 \impliedby t(s+t)^3 \leq \frac{n}{\log^2 n} \impliedby s \leq \left(\frac{n}{t \log^2 n}\right)^{1/3} - t,$$

where the last inequality is valid given that $t < n^{1/4}/(2\sqrt{\log n})$. Let s be the maximum integer satisfying the inequality above. Then, the quantity $H_n^{\mathcal{S}}(p)$ admits the lower bound

$$H_n^{\mathcal{S}}(p) \geq \sum_{j=s+1}^{s+t-1} (j \log n) \geq \frac{(2s+t)(t-1)}{2} \log n \geq \frac{1}{4} \left(\frac{n}{t \log^2 n} \right)^{1/3} t \log n = \Omega((nt^2 \log n)^{1/3}).$$

Finally, if $t \geq n_0 := n^{1/4}/(2\sqrt{\log n})$, then the distribution p has the following form

$$p := (t - n_0 + 1) \cdot \{p_0\} \bigcup \left(\bigcup_{j=1}^{n_0-1} \left((j \log n) \cdot \left\{ j^2 \frac{\log n}{n} \right\} \right) \right),$$

where p_0 is a properly chosen to make the probabilities sum to 1. By the reasoning for the last case, the distribution p is well-defined. In addition, the quantity $H_n^{\mathcal{S}}(p)$ satisfies

$$H_n^{\mathcal{S}}(p) \geq \sum_{j=1}^{n_0-1} (j \log n) \geq \frac{n_0(n_0-1)}{2} \log n \geq \Omega(\sqrt{n}).$$

Consolidating these results yields the desired lower bound. □

C Competitive Estimation of Distributions and Their Entropy

C.1 Competitive Distribution Estimation

Competitive estimation calls for an estimator that competes with the instance-by-instance performance of a genie knowing more information, but reasonably restricted. Denote by $\ell_{\text{KL}}(p, q)$ the KL divergence. Introduced in Orlitsky and Suresh [2015], the formulation considers the collection $\mathcal{N}$ of all natural estimators, and shows that a simple variant $\hat{p}^{\text{GT}}$ of the Good-Turing estimator achieves

$$\ell_{\text{KL}}(p, \hat{p}^{\text{GT}}_{X^n}) - \min_{\hat{p} \in \mathcal{N}} \ell_{\text{KL}}(p, \hat{p}_{X^n}) \leq \mathcal{O}\left(\frac{1}{n^{1/3}} \right)$$

for every distribution p and with high probability. We refer to the left-hand side as the *excess loss* of estimator $\hat{p}_{\text{GT}}$ with respect to the best natural estimator, and note that it vanishes at a rate independent of p. For a more involved estimator in Acharya et al. [2013], the excess loss vanishes at a faster rate of $\tilde{\mathcal{O}}(\min\{1/\sqrt{n}, |\mathcal{X}|/n\})$, which is optimal up to logarithmic factors for every estimator and the respective worst-case distribution. For the ℓ_1 distance, the work of Valiant and Valiant [2016] derives a similar result.

These estimators track the loss of the best natural estimator for each distribution. Yet an equally important component, the excess loss bound, is still of the worst-case nature. For a fully adaptive guarantee, Hao and Orlitsky [2019c] design an estimator $\hat{p}^\star$ that achieves a $\mathcal{D}_n/n$ excess loss, i.e.,

$$\ell_{\text{KL}}(p, \hat{p}^\star_{X^n}) - \min_{\hat{p} \in \mathcal{N}} \ell_{\text{KL}}(p, \hat{p}_{X^n}) \leq \tilde{\mathcal{O}}\left(\frac{\mathcal{D}_n}{n} \right),$$

for every p and $X^n \sim p$, with high probability. Utilizing the adaptiveness of $\mathcal{D}_n$ to the simplicity of distributions, the paper derives excess-loss bounds for several important distribution families, and proves the estimator's optimality under various of classical and modern learning frameworks.

New results While the work of Hao and Orlitsky [2019c] provides an appealing upper bound on the excess loss, it is not exactly clear how good this bound is as a matching lower bound is missing. In this work, we complete the picture by showing that the $\mathcal{D}_n/n$ bound is essential for competitive estimation and optimal up to logarithmic factors of n.

Theorem 5 (Minimal excess loss)**.** *For any $n, \mathcal{D} \in \mathbb{N}$ and distribution estimator $\hat{p}'$, there is a distribution p such that with probability at least $9/10$, we have both*

$$\mathcal{O}(\log n + \mathcal{D}) \geq \mathcal{D}_n$$

and

$$\ell_{KL}(p, \hat{p}'_{X^n}) - \min_{\hat{p} \in \mathcal{N}} \ell_{KL}(p, \hat{p}_{X^n}) \geq \Omega\left(\frac{\mathcal{D}}{n} \right).$$

According to Theorem 1, we can replace $\mathcal{D}_n$ by multiples $\tilde{\Theta}(H(\Phi^n))$ of the profile entropy in both the upper and lower bounds.

Proof. Denote $s := (D/\log n)^{1/2}$, $I := \{s, s+1, \ldots, 2s\}$, and $P := \cup_{i\in I} P_i := \cup_{i\in I} U_i/n$ where

$$U := \bigcup_{i\in I} U_i := \bigcup_{i\in I} \{i^2 \log^2 n, i^2 \log^2 n + 1, \ldots, i^2 \log^2 n + i \log n\},$$

where $D \lesssim \sqrt{n/\log n}$ for the total to be at most n. Let $A \cdot \{B\}$ denote the length-A constant sequence of value B. Let C be the set of distributions in the form of

$$p := L \cdot \left\{\frac{1}{n^2}\right\} \bigcup \left(\bigcup_i (i \log n) \cdot \left\{ q_i \text{ or } q_i' : nq_i = i^2 \log^2 n, nq_i' = i^2 \log^2 n + i \log n \right\} \right).$$

where the probability values are sorted according to the ordering they appear above, L is a proper variable that makes the probabilities sum to 1, and the range of support of distribution p is irrelevant for our purpose and hence unspecified. Equip a uniform prior over C (equivalently, construct a random distribution). We have several claims in order:

- For any $i \in I$ and $\mu \in U_i$, by the construction and independence,

$$\begin{aligned}\Pr(\varphi_\mu = 1 | q_i \text{ is chosen}) &\approx (i \log n) \cdot \left(\Pr(\mathrm{Poi}(nq_i) = \mu) \cdot (\Pr(\mathrm{Poi}(nq_i) \neq \mu))^{i \log n - 1} \right) \\ &\approx (i \log n) \cdot \left(\frac{1}{\sqrt{nq_i}} \cdot \left(1 - \frac{1}{\sqrt{nq_i}}\right)^{i \log n - 1} \right) \\ &\geq \Omega(1).\end{aligned}$$

  Similarly, we have $\Pr(\varphi_\mu = 1 | q_i' \text{ is chosen}) \geq \Omega(1)$. Hence,

$$\Pr(\varphi_\mu = 1) \geq \Omega(1).$$

- For any $i \in I$ and $\mu \in U_i$, by Bayes' rule,

$$\Pr(q_i \text{ is chosen} | \varphi_\mu = 1) = \frac{\Pr(\varphi_\mu = 1 | q_i \text{ is chosen}) \cdot 0.5}{\Pr(\varphi_\mu = 1)} \geq \Omega(1).$$

  Similarly, we have $\Pr(q_i' \text{ is chosen} | \varphi_\mu = 1) \geq \Omega(1)$.

- For any $i \in I$ and $\mu \in U_i$, the value of M_μ, the total probability of symbols appearing μ times, is q_i if $\varphi_\mu = 1$ and q_i is chosen; and is q_i' if $\varphi_\mu = 1$ and q_i is chosen. Any estimator E_μ will incur an expected absolute error of $\Omega(i(\log n)/n)$ in estimating M_μ given $\varphi_\mu = 1$.

- Note that for any $\alpha \in [0, 1]$ and $x, y > 0$,

$$\alpha(y - z)^2 + (1 - \alpha)(z - x)^2 \geq \alpha(1 - \alpha)(x - y)^2.$$

- Therefore, the expected squared Hellinger distance $H^2(\cdot, \cdot)$ of any estimator E_μ in estimating $(M_\mu)_{\mu \geq 0}$ satisfies

$$\begin{aligned}\frac{1}{2} \sum_{\mu \geq 0} \mathbb{E}\left(\sqrt{E_\mu} - \sqrt{M_\mu}\right)^2 &\geq \frac{1}{2} \sum_{i \in I} \sum_{\mu \in U_i} \mathbb{E}\left[\left(\sqrt{E_\mu} - \sqrt{M_\mu}\right)^2 \Big| \varphi_\mu = 1 \right] \Pr(\varphi_\mu = 1) \\ &= \frac{1}{2} \sum_{i \in I} \sum_{\mu \in U_i} \mathbb{E}\left[\left(\frac{E_\mu - M_\mu}{\sqrt{E_\mu} + \sqrt{M_\mu}} \right)^2 \Big| \varphi_\mu = 1 \right] \Pr(\varphi_\mu = 1) \\ &\geq \sum_{i \in I} (i \log n) \cdot \Omega\left(\frac{(i \log n)/n}{\sqrt{i^2 (\log^2 n)/n}} \right)^2 \\ &\geq s \cdot \Omega\left(\frac{s \log n}{n} \right) \\ &= \Omega\left(\frac{D}{n} \right).\end{aligned}$$

- Consequently, by the inequality $D(P \parallel Q) \geq 2H^2(P,Q)$,

$$\mathbb{E}\left[D(E \parallel M)\right] \geq \mathbb{E}\left[2H^2(E,M)\right] \geq \Omega\left(\frac{D}{n}\right).$$

- Finally, the value of $\mathbb{E}[D(X^n)]$ is at most $\mathcal{O}(\log n + s(s\log n)) = \mathcal{O}(\log n + D)$. □

C.2 Competitive Entropy Estimation

The next theorm shows that for *every* distribution and among all plug-in entropy estimators, the distribution estimator in Hao and Orlitsky [2019c] is as good as the one that performs best in estimating the actual distribution.

Denote by $\mathcal{N}$ the collection of all natural estimators. Write $|H(p) - H(q)|$ as $\ell_H(p,q)$ for compactness and the KL-divergence between $p, q \in \Delta_{\mathcal{X}}$ as $\ell_{\text{KL}}(p,q)$.

Theorem 4 (Competitive entropy estimation)**.** *For any distribution p, sample $X^n \sim p$ with profile $\Phi^n := \varphi(X^n)$, and $\hat{p}^{\mathcal{N}}_{X^n} := \arg\min_{\hat{p}\in\mathcal{N}} \ell_{KL}(p, \hat{p}_{X^n})$, we have*

$$\ell_H(p, \hat{p}^{\star}_{X^n}) - \ell_H(p, \hat{p}^{\mathcal{N}}_{X^n}) \leq \tilde{\mathcal{O}}\left(\sqrt{\frac{\lceil H(\Phi^n)\rceil}{n}}\right).$$

with probability at least $1 - \mathcal{O}(1/n)$.

Proof. Given any natural estimator and a sample $X^n \sim p$, we denote by q the distribution estimate. The entropy of q differs from the true entropy by

$$\begin{aligned}
H(q) - H(p) &= -\sum_x q_x \log q_x + \sum_x p_x \log p_x \\
&= \sum_x p_x \log p_x - \sum_x p_x \log q_x + \sum_x p_x \log q_x - \sum_x q_x \log q_x \\
&= \sum_x p_x \log \frac{p_x}{q_x} + \sum_x (p_x - q_x)\log q_x \\
&= \ell_{\text{KL}}(p,q) + \sum_x (p_x - q_x)\log q_x.
\end{aligned}$$

Denote by $P_\mu(X^n)$ and $Q_\mu(X^n)$ the total probability that distributions p and q assign to symbols with multiplicity μ. Since q is induced by a natural estimator, we also write $q_\mu(X^n)$ for the probability that q assigns to *each* symbol with multiplicity μ in X^n. Recall that prevalence $\varphi_\mu(X^n)$ denotes the number of symbols with multiplicity μ in X^n. Therefore, $Q_\mu(X^n) = \varphi_\mu(X^n) \cdot q_\mu(X^n)$.

Henceforth, whenever it is clear from the context, we suppress X^n in related expressions. Then, the second term on the right-hand side satisfies

$$\begin{aligned}
\sum_x (p_x - q_x)\log q_x &= \sum_x (\sum_\mu \mathbb{1}_{\mu_x=\mu}\cdot p_x - \sum_\mu \mathbb{1}_{\mu_x=\mu}\cdot q_\mu)\log(\sum_\mu \mathbb{1}_{N_x=\mu}\cdot q_\mu) \\
&= \sum_x \sum_\mu \mathbb{1}_{\mu_x=\mu}\cdot (p_x - q_\mu)\log q_\mu \\
&= \sum_\mu (\sum_x \mathbb{1}_{\mu_x=\mu}\cdot p_x - \sum_x \mathbb{1}_{\mu_x=\mu}\cdot q_\mu)\log q_\mu \\
&= \sum_\mu (P_\mu - Q_\mu)\log q_\mu.
\end{aligned}$$

Let q_{min} be the smallest nonzero probability of q. By the triangle inequality and Pinsker's inequality,

$$\begin{aligned}
\left|\sum_\mu (P_\mu - Q_\mu)\log q_\mu\right| &\leq \sum_\mu |(P_\mu - Q_\mu)\log q_\mu| \\
&\leq |\log q_{\text{min}}| \sum_\mu |P_\mu - Q_\mu| \\
&\leq |\log q_{\text{min}}| \sqrt{2\ell_{\text{KL}}(P,Q)}.
\end{aligned}$$

By definition, $\hat{p}^{\mathcal{N}}_{X^n} = \arg\min_{\hat{p}\in\mathcal{N}} \ell_{\mathrm{KL}}(p, \hat{p}_{X^n})$. Now we show that if a symbol x has multiplicity μ, the estimator $\hat{p}^{\mathcal{N}}$ will assign a probability mass of P_μ/φ_μ. In other words, $\hat{P}^{\mathcal{N}}_\mu = P_\mu$ since $p^{\mathcal{N}} \in \mathcal{N}$. Indeed, the corresponding KL-divergence values differ by

$$\begin{aligned}\sum_x p_x \log \frac{p_x}{q_x} - \sum_x \sum_\mu \mathbb{1}_{\mu_x=\mu} \cdot p_x \log \frac{p_x}{P_\mu/\varphi_\mu} &= \sum_x p_x \log \frac{1}{q_x} - \sum_x \sum_\mu \mathbb{1}_{\mu_x=\mu} \cdot p_x \log \frac{\varphi_\mu}{P_\mu} \\ &= \sum_x \sum_\mu \mathbb{1}_{\mu_x=\mu} \cdot p_x \log \frac{P_\mu}{\varphi_\mu q_\mu} \\ &= \sum_\mu P_\mu \log \frac{P_\mu}{Q_\mu} = \ell_{\mathrm{KL}}(P, Q) \geq 0.\end{aligned}$$

Then, the above equalities yield that,

$$H(\hat{p}^{\mathcal{N}}) - H(p) = \ell_{\mathrm{KL}}(p, \hat{p}^{\mathcal{N}}) + \sum_\mu \left(P_\mu - \hat{P}^{\mathcal{N}}_\mu\right) \log p^{\mathcal{N}}_\mu = \ell_{\mathrm{KL}}(p, \hat{p}^{\mathcal{N}}) = \min_{\hat{p}\in\mathcal{N}} \ell_{\mathrm{KL}}(p, \hat{p}_{X^n}).$$

Next consider the other estimator $\hat{p}^\star$, which is also natural. Let $\mathcal{D}_n = \mathcal{D}(\Phi^n)$ be the profile dimension of X^n. By the results in Hao and Orlitsky [2019c], estimator $\hat{p}^\star$ achieves a $\mathcal{D}_n/n$ excess loss, i.e.,

$$\ell_{\mathrm{KL}}(p, \hat{p}^\star_{X^n}) - \min_{\hat{p}\in\mathcal{N}} \ell_{\mathrm{KL}}(p, \hat{p}_{X^n}) = \ell_{\mathrm{KL}}(P(X^n), \hat{P}^\star(X^n)) \leq \tilde{\mathcal{O}}\left(\frac{\mathcal{D}_n}{n}\right),$$

for every p and $X^n \sim p$, with probability at least $1 - \mathcal{O}(1/n)$. In addition, by its construction, the minimum probability $\hat{p}^\star_{\min}(X^n)$ is at least $1/n^4$. Therefore, with probability at least $1 - \mathcal{O}(1/n)$,

$$\left|\sum_x (p_x - \hat{p}^\star_x) \log \hat{p}^\star_x\right| = \left|\sum_\mu \left(P_\mu - \hat{P}^\star_\mu\right) \log \hat{p}^\star_\mu\right| \leq |\log \hat{p}^\star_{\min}| \cdot \sqrt{2\ell_{\mathrm{KL}}(P, \hat{P}^\star_\mu)} \leq \tilde{\mathcal{O}}\left(\sqrt{\frac{\mathcal{D}_n}{n}}\right).$$

Finally, the triangle inequality combines the above results and yields

$$\begin{aligned}\ell_H(p, \hat{p}^\star) - \ell_H(p, \hat{p}^{\mathcal{N}}) &= |H(p) - H(\hat{p}^\star)| - |H(p) - H(\hat{p}^{\mathcal{N}})| \\ &= \left|\ell_{\mathrm{KL}}(p, \hat{p}^\star_x) + \sum_x (p_x - \hat{p}^\star_x) \log \hat{p}^\star_x\right| - \left|\min_{\hat{p}\in\mathcal{N}} \ell_{\mathrm{KL}}(p, \hat{p}_{X^n})\right| \\ &\leq \left|\ell_{\mathrm{KL}}(p, \hat{p}^\star_x) - \min_{\hat{p}\in\mathcal{N}} \ell_{\mathrm{KL}}(p, \hat{p}_{X^n})\right| + \left|\sum_x (p_x - \hat{p}^\star_x) \log \hat{p}^\star_x\right| \\ &= \ell_{\mathrm{KL}}(P, \hat{P}^\star_\mu) + \tilde{\mathcal{O}}\left(\sqrt{\frac{\mathcal{D}_n}{n}}\right) \\ &\leq \tilde{\mathcal{O}}\left(\sqrt{\frac{\mathcal{D}_n}{n}}\right).\end{aligned}$$

This together with Theorem 1 completes the proof. □

D Optimal Profile Compression

While a labeled sample contains all information, for many modern applications, such as property estimation and differential privacy, it is sufficient Orlitsky et al. [2004] or even necessary to provide only the profile Suresh [2019]. Hence, this section focuses on the lossless compression of profiles.

For any distribution p, it is well-known that the minimal expected codeword length (MECL) for losslessly compressing a sample $X^n \sim p$ is approximately $nH(p)$, which increases linearly in n as long as $H(p)$ is bounded away from zero.

On the other hand, by the Hardy-Ramanujan formula Hardy and Ramanujan [1918], the number $\mathbb{P}(n)$ of integer partitions of n, which happens to equal to the number of length-n profiles, satisfies

$$\log \mathbb{P}(n) = 2\pi\sqrt{\frac{n}{6}}(1 + o(1)).$$

Consequently, the MECL for losslessly compressing the sample profile $\Phi^n \sim p$ is at most $\mathcal{O}(\sqrt{n})$, a number potentially much smaller than $nH(p)$.

By Shannon's source coding theorem, the profile entropy $H(\Phi^n)$ is the *information-theoretic limit* of MECL for the lossless compression of profile $\Phi^n \sim p$. Below, we present explicit block and sequential profile compression schemes achieving this entropy limit, up to logarithmic factors of n.

D.1 Block Compression

The block compression algorithm we propose is intuitive and easy to implement.

Recall that the profile of a sequence x^n is the multiset $\varphi(x^n)$ of multiplicities associated with symbols in x^n. The ordering of elements in a multiset is not informative. Hence equivalently, we can compress $\varphi(x^n)$ into the set $\mathcal{C}(\varphi(x^n))$ of corresponding multiplicity-prevalence pairs, i.e.,

$$\mathcal{C}(\varphi(x^n)) := \{(\mu, \varphi_\mu(x^n)) : \mu \in \varphi(x^n)\}.$$

The number of pairs in $\mathcal{C}(\varphi(x^n))$ is equal to the profile dimension $\mathcal{D}(\varphi(x^n))$. In addition, both a prevalence and its multiplicity are integers in $[0, n]$, and storing the pair takes $2 \log n$ nats. Hence, it takes at most $2(\log n) \cdot \mathcal{D}(\varphi(x^n))$ nats to store the compressed profile. By Theorem 1, for any distribution $p \in \Delta_{\mathcal{X}}$ and $\Phi^n \sim p$,

$$\mathbb{E}[2(\log n) \cdot \mathcal{D}(\Phi^n)] = \tilde{\Theta}(\lceil H(\Phi^n) \rceil).$$

D.2 Sequential Compression

For any sequence x^n, the setting for sequential profile compression is that at time step $t \in [n]$, the compression algorithm knows only $\varphi(x^t)$ and sequentially encodes the new information. This is equivalent to providing the algorithm $\mu_{x_t}(x^{t-1})$ at time step t.

Suppress x, x^t in the expressions for the ease of illustration. For efficient compression, we sequentially encode the profile φ into a *self-balancing binary search tree* $\mathcal{T}$, with each node storing a multiplicity-prevalence pair (μ, φ_μ) and μ being the search key. We present the algorithm details as follows.

Algorithm 2 Sequential Profile Compression

input sequence $(\mu_{x_t}(x^{t-1}))_{t=1}^n$, tree $\mathcal{T} = \varnothing$
output tree $\mathcal{T}$ that encodes the input sequence
  for t = 1 to n **do**
    if $\mu := \mu_{x_t}(x^{t-1}) \in \mathcal{T}$ **then**
      if $\mu + 1 \in \mathcal{T}$ **then**
        $\varphi_{\mu+1} := \mathcal{T}(\mu + 1) \leftarrow \mathcal{T}(\mu + 1) + 1$
      else
        add $(\mu + 1, 1)$ to $\mathcal{T}$
      end if
      if $\varphi_\mu = 1$ **then** delete (μ, φ_μ) from $\mathcal{T}$
      else $\varphi_\mu := \mathcal{T}(\mu) \leftarrow \mathcal{T}(\mu) - 1$ **endif**
    else
      if $1 \notin \mathcal{T}$ **then** add $(1, 1)$ to $\mathcal{T}$
      else $\mathcal{T}(1) \leftarrow \mathcal{T}(1) + 1$ **endif**
    end if
  end for

The algorithm runs for exactly n iterations, with a $\mathcal{O}(\log n)$ per-iteration time complexity. For an i.i.d. sample $X^n \sim p$, the expected space complexity is again $\tilde{\Theta}(\lceil H(\Phi^n) \rceil)$.

E Extensions and Additional Results

E.1 Multi-Dimensional Profiles

As we elaborate below, the notion of profile generalizes to the multi-sequence setting.

Let $\mathcal{X}$ be a finite or countably infinite alphabet. For every $\vec{n} := (n_1, \ldots, n_d) \in \mathbb{N}^d$ and tuple $x^{\vec{n}} := (x_1^{n_1}, \ldots, x_d^{n_d})$ of sequences in $\mathcal{X}^*$, the *multiplicity* $\mu_y(x^{\vec{n}})$ of a symbol $y \in \mathcal{X}$ is the vector of its frequencies in the tuple of sequences. The *profile* of $x^{\vec{n}}$ is the multiset $\varphi(x^{\vec{n}})$ of multiplicities of the observed symbols Acharya et al. [2010], Das [2012], Charikar et al. [2019b], and its *dimension* is the number $\mathcal{D}(x^{\vec{n}})$ of distinct elements in the multiset. Drawing independent samples from $\vec{p} := (p_1, \ldots, p_d) \in \Delta_{\mathcal{X}}^d$, the *profile entropy* is simply the entropy of the joint-sample profile.

Many of the previous results potentially generalize to this multi-dimensional setting. For example, the $\sqrt{2n}$ bound on $\mathcal{D}(x^{\vec{n}})$ in the 1-dimensional case becomes

Theorem 20. *For any $\mathcal{X}$, $\vec{n}$, and $x^{\vec{n}} \in \mathcal{X}^{\vec{n}}$, there exists a positive integer r such that*

$$\sum_i n_i \geq d \cdot \binom{d+r-1}{d+1},$$

and

$$\mathcal{D} \leq \binom{d+r}{d} - 1.$$

This essentially recovers the $\sqrt{2n}$ bound for $d = 1$.

Proof. For simplicity, we suppress $x^{\vec{n}}$ in $\mathcal{D}(x^{\vec{n}})$. Let Δ_d denote the standard d-dimensional simplex. As each multiplicity corresponds to a vector in $\mathbb{N}^d$, in the ideal case, the profile that has the maximum dimension $\mathcal{D}$ corresponds to the integer vectors in the scaled simplex $(r \cdot \Delta_d)$, for some properly chosen parameter r. For the minimum value of such a parameter $r \in \mathbb{Z}^+$, we have

$$\begin{aligned}
\sum_i n_i &\geq \sum_{t=0}^{r-1} \binom{t+d-1}{d-1} \cdot t \\
&= d \cdot \sum_{t=1}^{r-1} \binom{t+d-1}{d} \\
&= d \cdot \sum_{(t-1)=0}^{r-2} \binom{(t-1)+d}{(t-1)} \\
&= d \cdot \binom{d+r-1}{d+1},
\end{aligned}$$

and

$$\mathcal{D} \leq \sum_{t=1}^{r} \binom{t+d-1}{t} = \binom{d+r}{d} - 1.$$

Consolidating these two inequalities yields the desired result. ☐

E.2 Discrete Multi-Variate Gaussian

Given a mean vector $\mu \in \mathbb{Z}^d$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ with *eigenvalues at least* 1, the corresponding *discrete d-dimensional Gaussian* is specified by its probability mass function

$$p(x) := \frac{1}{C} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)\right), \forall x \in \mathbb{Z}^d.$$

where $C > 0$ is a normalizing constant. Note that definition is slightly different from that induced by the discretization procedure presented in Section B.4. The reason for adopting this definition (which is also standard in literature) is to simplify the subsequent reasoning. Let $\sigma_1^2 \leq \sigma_2^2 \ldots \leq \sigma_d^2$ be the d eigenvalues of Σ, where $\sigma_1^2 \geq 1$ by assumption. In this section, we show that for $d \geq 9$,

$$H_n^{\mathcal{S}}(p) \leq \mathcal{O}(\log n)\left(1 + \min\left\{\frac{n}{C}, \gamma_d(\alpha_\Sigma \cdot \beta_{d,n})^d \cdot C\right\}\right),$$

where $\alpha_\Sigma := \exp\left(6\sigma_d^2/\sigma_1^2\right)$ and $\beta_{d,n} := \sqrt{(2\log n)/d}$, and γ_d is a constant that depends only on d, which appears in Lemma 14. Note that the above bound resembles that for univariate log-concave distributions (Theorem 13). This result is not included in the main paper due to the different setting.

Lower bound on C First we bound the value of C from below in terms of its eigenvalues and other parameters. By symmetry, we can decompose the matrix Σ as

$$\Sigma = V \Lambda V^T,$$

where Λ is a diagonal matrix with $\Lambda_{ii} = \sigma_i^2$, and V is an orthonormal matrix whose i-th column is the eigenvector v_i associated with σ_i^2.

Partition the real space $\mathbb{R}^d$ into unit cubes whose vertices belong to $\mathbb{Z}^d$. For any two vectors $\tilde{a}, \tilde{b} \in \mathbb{R}^d$ that belong to the same unit cube, we want to bound the ratio between $p(\tilde{a})$ and $p(\tilde{b})$. Denote $a := \tilde{a} - \mu$ and $b := \tilde{b} - \mu$, and express a and b as linear combinations of eigenvectors,

$$a := \sum_{i=1}^{d} x_i \cdot v_i \text{ and } b := \sum_{i=1}^{d} y_i \cdot v_i.$$

The log-ratio between the corresponding probabilities satisfies

$$\begin{aligned}
-2 \log \frac{p(\tilde{a})}{p(\tilde{b})} &= a^T \Sigma^{-1} a - b^T \Sigma^{-1} b \\
&= (a+b)^T \Sigma^{-1} (a-b) \\
&= \left(\sum_i (x_i + y_i) \cdot v_i^T \right) V \Lambda^{-1} V^T \left(\sum_i (x_i - y_i) \cdot v_i \right) \\
&= \left(\sum_i (x_i + y_i) \cdot e_i^T \right) \Lambda^{-1} \left(\sum_i (x_i - y_i) \cdot e_i \right) \\
&= \sum_i \sigma_i^{-2} (x_i^2 - y_i^2).
\end{aligned}$$

Note that $\sum_i (x_i - y_i)^2 = \|a - b\|_2^2 = \sum_i (\tilde{a}_i - \tilde{b}_i)^2 \le d$ since $\tilde{a} - \tilde{b} = a - b$ and $\tilde{a}, \tilde{b}$ belong to the same unit cube. Hence, we bound the absolute value of the ratio by

$$\begin{aligned}
2 \left| \log \frac{p(\tilde{a})}{p(\tilde{b})} \right| &= \left| \sum_i \sigma_i^{-2} (x_i^2 - y_i^2) \right| \\
&\le \sum_i \sigma_i^{-2} \left| x_i^2 - (x_i - (x_i - y_i))^2 \right| \\
&\le 2 \sum_i \sigma_i^{-2} \left(x_i^2 + (x_i - y_i)^2 \right) \\
&\le 2 \sigma_1^{-2} \left(\sum_i x_i^2 + d \right) \\
&= 2 \sigma_1^{-2} \left(\|\tilde{a} - \mu\|_2^2 + d \right).
\end{aligned}$$

Now, consider the hyper-ellipse E induced by

$$(x - \mu)^T \Sigma^{-1} (x - \mu) \le d.$$

For any $x \in E$, simple algebra shows that $\|x - \mu\|_2^2 \le d \sigma_d^2$. Hence by the previous discussion, for any unit cube U with vertices in $\mathbb{Z}^d$, there exists a vertex v_U of U such that for any $x \in U \cap E$,

$$\left| \log \frac{p(x)}{p(v_U)} \right| \le \sigma_1^{-2} \left(\|x - \mu\|_2^2 + d \right) \le \sigma_1^{-2} \left(d \sigma_d^2 + d \right) \le 2d \left(\frac{\sigma_d}{\sigma_1} \right)^2.$$

Note that $x \in E$ is equivalent to $p(x) \ge \exp(-d/2)/C$. The probability mass over E is at least

$$\int_{x \in E} p(x) dx \ge \int_{x \in E} \frac{\exp(-d/2)}{C} = \frac{\exp(-d/2)}{C} \cdot \text{Vol}(E) = \frac{\exp(-d/2)}{C} \cdot \frac{(\pi d)^{d/2}}{\Gamma(d/2+1)} \prod_{i=1}^{d} \sigma_i.$$

On the other hand, this probability mass is at most

$$\int_{x\in E} p(x)dx = \sum_U \int_x p(x)\cdot \mathbb{1}_{x\in E\cap U}dx \quad \le \sum_U p(v_U)\cdot \exp\left(2d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right) \le \exp\left(3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right).$$

Consolidating the lower and upper bounds and multiplying both sides by C yield

$$\begin{aligned}
&C \ge \exp\left(-3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right)\exp\left(-\frac{d}{2}\right)\cdot\frac{(\pi d)^{d/2}}{\Gamma(d/2+1)}\prod_{i=1}^d\sigma_i\\
\implies &C \ge \exp\left(-3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right)\cdot\frac{(\pi d/e)^{d/2}}{\sqrt{e\pi(d/2)}(d/(2e))^{d/2}}\prod_{i=1}^d\sigma_i\\
\implies &C \ge \exp\left(-3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right)\cdot\frac{(2\pi)^{d/2}}{\sqrt{e\pi(d/2)}}\prod_{i=1}^d\sigma_i\\
\implies &C \ge \exp\left(-3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right)\prod_{i=1}^d\sigma_i.
\end{aligned}$$

where the first implication follows by the lemma below.

Lemma 13. *For any integer or semi-integer $x \ge 1/2$,*

$$\sqrt{2\pi x}\left(\frac{x}{e}\right)^x \le \Gamma(x+1) \le \sqrt{e\pi x}\left(\frac{x}{e}\right)^x.$$

Upper bound We proceed to bound $H_n^{\mathcal{S}}(p) = \sum_{j\ge1}\min\left\{p_{I_j}, j\cdot\log n\right\}$.

Below we assume that $C < n/\log n$, since otherwise $p(x) \le (\log n)/n, \forall x$, yielding an $\mathcal{O}(\log n)$ upper bound on $H_n^{\mathcal{S}}(p)$. Then by definition, the last index j such that $p_{I_j} > 0$ satisfies

$$(j-1)^2\frac{\log n}{n} \le \frac{1}{C} \implies j \le 1+\sqrt{\frac{1}{C}\frac{n}{\log n}} \le 2\sqrt{\frac{1}{C}\frac{n}{\log n}}$$

Denote by J the quantity on the right-hand side. Then,

$$\sum_{j\ge1}\min\left\{p_{I_j}, j\cdot\log n\right\} \le \sum_{j=1}^J j\log n \le J^2\log n \le \frac{4n}{C}.$$

Furthermore, by a reasoning similar to that above, the collection of points $x\in\mathbb{Z}^d$ satisfying $p(x) \le 1/(Cn) = p(\mu)/n \le 1/n$ contributes at most $\mathcal{O}(\log n)$ to $H_n^{\mathcal{S}}(p)$. Hence we need to analyze only points x satisfying $p(x) > 1/(Cn)$. Equivalently, points in

$$E^\star := \left\{x\in\mathbb{Z}^d : (x-\mu)^T\Sigma^{-1}(x-\mu) \le 2\log n\right\}.$$

Clearly, these points contribute at most $|E^\star|$ to the sum. Noting that $E^\star$ is a discrete hyper-ellipse, we can bound its cardinality via the following lemma Bentkus and Götze [1997].

Lemma 14. *Let $\mu\in\mathbb{R}^d$ be a mean vector, and $\Sigma\in\mathbb{R}^{d\times d}$ be a real covariance matrix with nonzero eigenvalues $\sigma_1^2 \le \ldots \sigma_d^2$. For any $d \ge 9$ and $t \ge \sigma_d^2$, the discrete ellipsoid*

$$E(t) := \left\{x\in\mathbb{Z}^d : (x-\mu)^T\Sigma^{-1}(x-\mu) \le t\right\}$$

admits the following inequality on its cardinality,

$$|E(t)| \le \left(1+\frac{\gamma_d}{t}\frac{1}{\sigma_d^2}\left(\frac{\sigma_d}{\sigma_1}\right)^{2d+4}\right)\frac{(\pi t)^{d/2}}{\Gamma(d/2+1)}\prod_{i=1}^d\sigma_i,$$

where $\gamma_d > 1$ is a constant that depends only on d.

For simplicity, write $\alpha_\Sigma := \exp\left(6\sigma_d^2/\sigma_1^2\right)$ and $\beta_{d,n} := \sqrt{(2\log n)/d}$. Applying the above lemma to bound $|E^\star|$ (where $t = 2\log n$) and combining the result with our lower bound on C yield

$$\begin{aligned}
|E(2\log n)| &\le \left(1 + \frac{\gamma_d}{2\log n}\frac{1}{\sigma_d^2}\left(\frac{\sigma_d}{\sigma_1}\right)^{2d+4}\right)\frac{(2\pi\log n)^{d/2}}{\Gamma(d/2+1)}\exp\left(3d\left(\frac{\sigma_d}{\sigma_1}\right)^2\right)C \\
&\le \left(1 + \frac{\gamma_d}{2\log n}\frac{1}{\sigma_d^2}\left(\frac{\sigma_d}{\sigma_1}\right)^{2d+4}\right)\frac{1}{\sqrt{\pi d}}\left(4e\pi\frac{\log n}{d}\right)^{d/2} e^{3d(\sigma_d/\sigma_1)^2}C \\
&\le \left(1 + \frac{\gamma_d}{2\log n}\left(\frac{\sigma_d}{\sigma_1}\right)^{3d}\right)\left(\frac{2\log n}{d}\right)^{d/2} e^{5d(\sigma_d/\sigma_1)^2}C \\
&\le \gamma_d\left(\frac{\sigma_d}{\sigma_1}\right)^{3d}\left(\frac{2\log n}{d}\right)^{d/2} e^{5d(\sigma_d/\sigma_1)^2}C \\
&\le \gamma_d\left(\frac{2\log n}{d}\right)^{d/2} e^{6d(\sigma_d/\sigma_1)^2}C \\
&= \gamma_d\left(\alpha_\Sigma\cdot\beta_{d,n}\right)^d C,
\end{aligned}$$

where the second step follows by Lemma 13.

To summarize, we have established the desired bound

$$H_n^{\mathcal{S}}(p) \le \mathcal{O}(\log n)\left(1 + \min\left\{\frac{n}{C}, \gamma_d(\alpha_\Sigma\cdot\beta_{d,n})^d\cdot C\right\}\right).$$

References

Jayadev Acharya, Hirakendu Das, Alon Orlitsky, Shengjun Pan, and Narayana P Santhanam. Classification using pattern probability estimators. In *2010 IEEE International Symposium on Information Theory*, pages 1493–1497. IEEE, 2010.

Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, and Shengjun Pan. Competitive closeness testing. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 47–68, 2011.

Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, Shengjun Pan, and Ananda Suresh. Competitive classification and closeness testing. In *Conference on Learning Theory*, pages 22–1, 2012.

Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Optimal probability estimation with applications to prediction and classification. In *Conference on Learning Theory*, pages 764–796, 2013.

Jayadev Acharya, Hirakendu Das, Alon Orlitsky, and Ananda Theertha Suresh. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. In *International Conference on Machine Learning*, pages 11–21, 2017.

Mark Yuying An. Log-concave probability distributions: Theory and statistical testing. *Duke University Dept of Economics Working Paper*, 95(3), 1997.

Rubén Armañanzas, Iñaki Inza, Roberto Santana, Yvan Saeys, Jose Luis Flores, Jose Antonio Lozano, Yves Van de Peer, Rosa Blanco, Víctor Robles, Concha Bielza, et al. A review of estimation of distribution algorithms in bioinformatics. *BioData mining*, 1(1):6, 2008.

Robert L Axtell. Zipf distribution of US firm sizes. *science*, 293(5536):1818–1820, 2001.

R Harald Baayen. *Word frequency distributions*, volume 18. Springer Science & Business Media, 2002.

Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic theory*, 26(2):445–469, 2005.

Alexander Barvinok. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics*, 16(2):329–342, 2016.

Alexander I Barvinok. Two algorithmic results for the traveling salesman problem. *Mathematics of Operations Research*, 21(1):65–84, 1996.

Vidmantas Bentkus and Friedrich Götze. On the lattice point problem for ellipsoids. *Acta Arithmetica*, 80(2):101–125, 1997.

Peter Bühlmann, Petros Drineas, Michael Kane, and Mark van der Laan. Learning structured distributions Ilias Diakonikolas. In *Handbook of Big Data*, pages 283–300. Chapman and Hall/CRC, 2016.

Chun Lam Chan, Winston Fernandes, Navin Kashyap, and Manjunath Krishnapur. Phase transitions for the uniform distribution in the PML problem and its Bethe approximation. *arXiv preprint arXiv:1506.00753*, 2015.

Moses Charikar, Kirankumar Shiragur, and Aaron Sidford. The Bethe approximation for structured matrices: an improved approximation for the profile maximum likelihood. In *NeurIPS 2019 Workshop on Information Theory and Machine Learning*, 2019a.

Moses Charikar, Kirankumar Shiragur, and Aaron 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, 2019b.

Stanley F Chen and Joshua Goodman. An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13(4):359–394, 1999.

Hirakendu Das. *Competitive tests and estimators for properties of distributions*. PhD thesis, UC San Diego, 2012.

Ilias Diakonikolas, Themis Gouleakis, John Peebles, and Eric Price. Collision-based testers are optimal for uniformity and closeness. *arXiv preprint arXiv:1611.03579*, 2016a.

Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Efficient robust proper learning of log-concave distributions. *arXiv preprint arXiv:1606.03077*, 2016b.

Adrian Drăgulescu and Victor M Yakovenko. Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. *Physica A: Statistical Mechanics and its Applications*, 299(1-2):213–221, 2001.

Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. *ACM SIGCOMM computer communication review*, 29(4):251–262, 1999.

Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.

Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Local moment matching: A unified methodology for symmetric functional estimation and distribution estimation under Wasserstein distance. *arXiv preprint arXiv:1802.08405*, 2018.

Yi Hao and Alon Orlitsky. The broad optimality of profile maximum likelihood. In *Advances in Neural Information Processing Systems*, pages 10989–11001, 2019a.

Yi Hao and Alon Orlitsky. Data amplification: Instance-optimal property estimation. *arXiv preprint arXiv:1903.01432*, 2019b.

Yi Hao and Alon Orlitsky. Doubly-competitive distribution estimation. In *International Conference on Machine Learning*, pages 2614–2623, 2019c.

Yi Hao and Alon Orlitsky. Unified sample-optimal property estimation in near-linear time. In *Advances in Neural Information Processing Systems*, pages 11104–11114, 2019d.

Yi Hao, Alon Orlitsky, Ananda Theertha Suresh, and Yihong Wu. Data amplification: A unified and competitive approach to property estimation. In *Advances in Neural Information Processing Systems*, pages 8834–8843, 2018.

Godfrey H Hardy and Srinivasa Ramanujan. Asymptotic formulaæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.

Jean Hausser and Korbinian Strimmer. Entropy inference and the James-Stein estimator, with application to nonlinear gene association networks. *Journal of Machine Learning Research*, 10(Jul): 1469–1484, 2009.

Erwan Hillion, Oliver Johnson, et al. A proof of the Shepp–Olkin entropy monotonicity conjecture. *Electronic Journal of Probability*, 24, 2019.

Nicolas E Humphries, Nuno Queiroz, Jennifer RM Dyer, Nicolas G Pade, Michael K Musyl, Kurt M Schaefer, Daniel W Fuller, Juerg M Brunnschweiler, Thomas K Doyle, Jonathan DR Houghton, et al. Environmental context explains Lévy and Brownian movement patterns of marine predators. *Nature*, 465(7301):1066–1069, 2010.

Yannis Ioannidis. The history of histograms (abridged). In *Proceedings 2003 VLDB Conference*, pages 19–30. Elsevier, 2003.

Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

Weihao Kong and Gregory Valiant. Estimating learnability in the sublinear data regime. In *Advances in Neural Information Processing Systems*, pages 5455–5464, 2018.

Weihao Kong, Gregory Valiant, and Emma Brunskill. Sublinear optimal policy value estimation in contextual bandits. *arXiv preprint arXiv:1912.06111*, 2019.

Pavel Kroupa. On the variation of the initial mass function. *Monthly Notices of the Royal Astronomical Society*, 322(2):231–246, 2001.

Erich Leo Lehmann. Some concepts of dependence. *The Annals of Mathematical Statistics*, pages 1137–1153, 1966.

László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.

LAT Machado and WB Rossow. Structural characteristics and radiative properties of tropical cloud clusters. *Monthly Weather Review*, 121(12):3234–3260, 1993.

Anne E Magurran. *Measuring biological diversity*. John Wiley & Sons, 2013.

Kai J Miller, Larry B Sorensen, Jeffrey G Ojemann, and Marcel Den Nijs. Power-law scaling in the brain surface electric potential. *PLoS computational biology*, 5(12), 2009.

Michael Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet mathematics*, 1(2):226–251, 2004.

Alon Orlitsky and Ananda Theertha Suresh. Competitive distribution estimation: Why is Good-Turing good. In *Advances in Neural Information Processing Systems*, pages 2143–2151, 2015.

Alon Orlitsky, Narayana P Santhanam, Krishnamurthy Viswanathan, and Junan Zhang. On modeling profiles instead of values. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 426–435. AUAI Press, 2004.

Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.

Adrien Saumard and Jon A Wellner. Log-concavity and strong log-concavity: a review. *Statistics surveys*, 8:45, 2014.

Thomas Schürmann and Peter Grassberger. Entropy estimation of symbol sequences. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 6(3):414–427, 1996.

Richard P Stanley. Log-concave and unimodal sequences in algebra, combinatorics, and geometry. *Ann. New York Acad. Sci*, 576(1):500–535, 1989.

Ananda Theertha Suresh. Differentially private anonymized histograms. In *Advances in Neural Information Processing Systems*, pages 7969–7979, 2019.

Lionel Roy Taylor. Aggregation, variance and the mean. *Nature*, 189(4766):732–735, 1961.

Gregory Valiant and Paul Valiant. Estimating the unseen: an n/log (n)-sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 685–694, 2011.

Gregory Valiant and Paul Valiant. Instance optimal learning of discrete distributions. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 142–155, 2016.

Paul Valiant and Gregory Valiant. Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems*, pages 2157–2165, 2013.

Sergio Verdú. Empirical estimation of information measures: A literature guide. *Entropy*, 21(8): 720, 2019.

Pascal O Vontobel. The Bethe approximation of the pattern maximum likelihood distribution. In *2012 IEEE International Symposium on Information Theory Proceedings*. IEEE, 2012.

Pascal 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.