
Compressed Maximum Likelihood

Yi Hao¹ Alon Orlitsky¹

Abstract

Maximum likelihood (ML) is one of the most fundamental and general statistical estimation techniques. Inspired by recent advances in estimating distribution functionals, we propose *compressed maximum likelihood* (CML) that applies ML to compressed samples. We show that CML is sample-efficient for several fundamental learning tasks over both discrete and continuous domains, including learning structural densities, estimating probability multisets, and inferring symmetric distribution functionals.

1. Introduction

Maximum likelihood (ML) is a strikingly simple yet fundamental statistical inference paradigm for estimating parameters in probabilistic data generalization models. Over the past century, it has been applied to derive numerous important results in mathematics, statistics, and machine learning (Friedman et al., 2001).

At its core, the ML principle favors the model that maximizes the observed data’s probability. Consider for example flipping a coin with unknown heads probability p ten times and observing six heads and four tails. The observed sequence’s probability is $p^6(1-p)^4$, and ML suggests that the best guess of p is $\hat{p} = 0.6$ that maximizes this probability.

ML is so natural and intuitive that it may be stumbled upon without realizing its depth and significance. Yet, it can be shown to be near-optimal under relatively mild and general regularity conditions (Van der Vaart, 2000). In particular, \hat{p} tends to the actual value of p as the number of observations increases and enjoys both consistency and efficiency.

This paper concerns *functional estimation*, where \mathcal{P} is a distribution collection over a domain space \mathcal{Z} , and $f : \mathcal{P} \rightarrow \mathcal{Q}$ is a functional, mapping distributions in \mathcal{P} into a space \mathcal{Q} equipped with a pseudo-metric d .

¹Department of Electrical and Computer Engineering, University of California, San Diego, USA. Correspondence to: Yi Hao <yih179@eng.ucsd.edu>.

An *estimator* for f is a function $\hat{f} : \mathcal{Z} \rightarrow \mathcal{Q}$ such that when $Z \sim p$, the loss $d(\hat{f}(Z), f(p))$ is small. The ML plug-in estimator of f first finds the ML distribution estimate

$$p_z := \arg \max_{p \in \mathcal{P}} p(z)$$

and then estimates the functional’s value as $f(p_z)$.

For example, if \mathcal{P} is a parametric family indexed by Θ , and f maps each $p_\theta \in \mathcal{P}$ to its index $\theta \in \Theta$, the task becomes the classical *parameter estimation* problem, and ML maps sample Z to the parameter θ maximizing $p_\theta(Z)$.

1.1. Statistical Guarantees of ML Methods

ML is particularly significant in modern scientific applications that often involve complex structures in high dimensions. Deriving a bespoke estimator for each application may be an arduous task. Yet ML provides a simple universal methodology that in principle can be applied to any problem.

Unfortunately, it is well-known that for finite samples, ML estimators may be suboptimal, and other estimators often concentrate faster around the true parameters. Hence, establishing *finite-sample guarantees* for ML-based methods is of fundamental importance.

While the ML principle is quite natural, showing its finite-sample efficiency is often not easy. Recent work Acharya et al. (2017a) provided a simple argument for bypassing the difficulty in analyzing ML methods. Their primary lemma, stated next, relies only on the fact that $p_z(z) \geq p(z)$ for any $z \in \mathcal{Z}$. For simplicity, we write $d_f(p, q)$ for $d(f(p), f(q))$.

Lemma 1. *For any \mathcal{Z} , \mathcal{P} , and accuracy $\varepsilon > 0$,*

$$\max_{p \in \mathcal{P}} \Pr_{Z \sim p} (d_f(p, p_Z) > 2\varepsilon) \leq |\mathcal{Z}| \max_{p \in \mathcal{P}} \Pr_{Z \sim p} (d(f(p), \hat{f}(Z)) > \varepsilon).$$

Namely, the ML plug-in estimator is competitive with all other estimators in that if any other estimator achieves loss at most ε with probability $1 - \delta$ then the ML plug-in estimator will achieve a 2ε loss with probability at least $1 - |\mathcal{Z}| \cdot \delta$.

Leveraging this lemma, the paper showed that the ML-based *profile maximum likelihood (PML)* probability-multiset estimator in Orlitsky et al. (2004) is sample-optimal for several symmetric functionals, including entropy and support size.

The *profile* of a sample represents the number of elements appearing any given number of times, and hence is a sufficient statistic for symmetric functionals. Therefore instead of maximizing the sample’s probability, Acharya et al. (2017a) maximized the probability of the sample’s profile.

1.2. Compressed Maximum Likelihood

Motivated by the above rationale, our paper develops a general framework for applying and analyzing ML for several statistical and machine-learning problems. It does so by adding a *compressor* between the sample and ML learner.

Specifically, given domain \mathcal{Z} , a *compressor* is a function φ that maps each $z \in \mathcal{Z}$ to an element $\varphi(z)$ in some co-domain Φ . In general, we will allow φ to possess randomness independent of the samples.

Given compressor φ , we define the φ -compressed maximum likelihood (CML) estimator as

$$p_\varphi := \arg \max_{p \in \mathcal{P}} p(\varphi(Z)),$$

where for notational brevity, we suppress Z in $p_{\varphi(Z)}$.

2. Main Results

2.1. Compression for Learning

Consider an arbitrary compressor φ , mapping elements in \mathcal{Z} to a co-domain Φ . We evaluate the compressor’s quality for our learning objective via the following two criteria.

Typicality A compressor is (m, γ) -typical for an integer m and probability threshold $\gamma \in (0, 1)$, if for every $p \in \mathcal{P}$, there is an m -element subset $\mathcal{T} \subseteq \Phi$ with probability

$$p(\mathcal{T}) \geq 1 - \gamma.$$

Intuitively, smaller m corresponds to a better compressor.

Learnability Given error parameters ε, δ , we say that the compressor *enables* (ε, δ) -learning under pseudo-metric d if there is an algorithm $\mathcal{A} : \Phi \rightarrow \mathcal{Q}$ satisfying

$$\Pr_{Z \sim p} (d(f(p), \mathcal{A}(\varphi(Z))) > \varepsilon) \leq \delta, \forall p \in \mathcal{P}.$$

Intuitively, the smaller the error parameters, the better the compressor is for our learning tasks.

The theorem below shows that for any “good quality” compressor, the respective CML plug-in estimate will also be accurate, with high probability.

Theorem 1. *For any compressor φ that is (m, γ) -typical and enables (ε, δ) -learning, distribution $p \in \mathcal{P}$, and $Z \sim p$,*

$$\Pr (d_f(p, p_{\varphi(Z)}) > 2\varepsilon) \leq \gamma + m \cdot \delta.$$

Proof. We prove the theorem by classifying all possible patterns $\phi \in \Phi$ into three categories.

Given any distribution $p \in \mathcal{P}$, let \mathcal{T} be the smallest set with $p(\mathcal{T}) \geq 1 - \gamma$. By definition, $|\mathcal{T}| \leq m$.

For any pattern $\phi \in \mathcal{T}$ that satisfies $p(\phi) > \delta$, since the compressor enables (ε, δ) -learning with an algorithm \mathcal{A} , we must have $d(f(p), \mathcal{A}(\phi)) \leq \varepsilon$. By the definition of CML, $p_\varphi(\phi) \geq p(\phi) > \delta$, hence, $d(f(p_\varphi), \mathcal{A}(\phi)) \leq \varepsilon$. The triangle inequality combines both and yields $d_f(p_\varphi, p) \leq 2\varepsilon$.

Consider $\phi \in \mathcal{T}$ satisfying $p(\phi) \leq \delta$. By $|\mathcal{T}| \leq m$, the total probability of all such patterns is at most $m \cdot \delta$. In addition, the total probability of patterns $\phi \notin \mathcal{T}$ is at most γ since $p(\mathcal{T}) \geq 1 - \gamma$. Therefore, the probability that $d_f(p_\varphi, p) > 2\varepsilon$ is at most $m \cdot \delta + \gamma$, which completes the proof. \square

Similar to Acharya et al. (2017a), it suffices to obtain a CML approximation for the competitive guarantees to hold.

Definition 1 (Approximate CML (ACML)). *For any $\beta \leq 1$, $z \in \mathcal{Z}$, and compressor φ , a distribution $\tilde{p}_{\varphi(z)} \in \mathcal{P}$ is a β -approximate CML if $\tilde{p}_{\varphi(z)}(\varphi(z)) \geq \beta \cdot p_{\varphi(z)}(\varphi(z))$.*

Corollary 1. *For any (m, γ) -typical compressor φ that enables (ε, δ) -learning, distribution $p \in \mathcal{P}$, $Z \sim p$, and a β -approximate CML $\tilde{p}_{\varphi(Z)}$,*

$$\Pr (d_f(p, \tilde{p}_{\varphi(Z)}) > 2\varepsilon) \leq \gamma + m \cdot \delta / \beta.$$

Building on this framework, we design and analyze CML estimators for various applications. For each, we add a compressor between the samples and the estimator. The main challenges are finding a good compressor φ that works well with ML and an effective algorithm \mathcal{A} for the task.

2.2. CML Estimators

The previous sections presented the general CML framework, and Theorem 1 demonstrated its statistical competitiveness. The remaining sections describe concrete CML estimators for four statistical inference tasks, over both continuous and discrete domains, described below.

Continuous distributions Section 4 applies the CML method to learn structured continuous *i.i.d.* distributions.

We consider a broad distribution class \mathcal{P} where the difference $p - q$ between any two distributions $p, q \in \mathcal{P}$ has essentially at most s sign changes. This class encompasses numerous essential distributions such as log-concave, piecewise polynomials, and Gaussian mixtures.

Theorem 2 shows that for any $p \in \mathcal{P}$, with sample size $\Theta(s \log(s/\varepsilon)/\varepsilon^3)$, the CML estimator achieves

$$\Pr(\|p - p_\varphi\|_1 > \varepsilon) \leq \mathcal{O}(\varepsilon).$$

Note that Yatracos’ method, described in Section 3, finds the approximate minimizer to the empirical distribution in \mathcal{A}_k -distance, and achieves better sample efficiency. However, in general, no efficient algorithm is known for finding such a minimizer, e.g., Theorem 6.4 in Devroye & Lugosi (2012).

By contrast, the CML approach transforms the learning task to a maximum likelihood problem. This establishes the efficiency of ML methods for learning structured distributions, and enables the use of numerous ML optimization algorithms. Some of these algorithms, like the EM algorithm (Bishop, 2006) are heuristic, while others are rigorous.

Learning distributions with bounded crossings was also considered in Acharya et al. (2017b). However, they assume access to an \mathcal{A}_k -projection oracle that finds the approximate \mathcal{A}_k -distance minimizers to the empirical distribution.

Discrete distributions Section 5 applies CML to learn discrete *i.i.d.* distributions where as above, every two distributions cross values at most s times.

We present two different CML formulations.

The first formulation in Section 5.1 applies the compressor in Section 4 along with a random mapping that transforms the sample from a discrete distribution to a continuous analog while maintaining the structural properties.

The second formulation constructs CML directly from the discrete sample. The approach essentially performs maximum likelihood on the quantized empirical distribution, with a well-tuned quantization level. The resulting ML objective, presented in Section 5.2, resembles that of the PML estimator in Acharya et al. (2017a).

Theorem 3 shows that for any distribution p with support $\{1, \dots, N\}$ satisfying $N \geq s/\varepsilon$, given a sample from p of size $n = \Theta(s \log(N)/\varepsilon^3)$, with probability at least $1 - e^{-s}$,

$$\|p - p_\varphi\|_1 \leq \mathcal{O}(\varepsilon).$$

Probability multisets Section 6 is geared towards the original PML method, a special CML whose compressor φ maps samples to *profiles* (Orlitsky et al., 2004; 2011; Das, 2012; Acharya et al., 2012; 2017a; Hao & Orlitsky, 2019a; 2020b; Charikar et al., 2019a,b; Han & Shiragur, 2021).

PML yields a natural estimate for the distribution probability multiset. Given support bound N and desired accuracy ε , it is necessary (in the worst case) to obtain $\Theta(N/(\varepsilon^2 \log N))$ observations from the distribution to estimate its probability multiset under the sorted ℓ_1 -distance (Valiant & Valiant, 2011c; 2013; Han et al., 2018; Hao & Orlitsky, 2019a).

Over the years, a sequence of works established the optimality of PML for multiset learning in sorted ℓ_1 distance with different accuracy levels (Das, 2012; Hao & Orlitsky,

2019a; Han & Shiragur, 2021). A different multiset learning guarantee stated in terms of the earth-mover’s distance was considered in Valiant & Valiant (2016).

Theorem 4 shows that PML also enjoys this learning guarantee, hence is sample-optimal for both sorted ℓ_1 -distance and τ -truncated relative earth-mover distance. For any level $\tau \in [0, 1]$, the latter metric is

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

where $\Gamma_{p,q}$ represents all possible couplings of p and q . Section 6.2 shows that with probability at least $1 - \mathcal{O}(1/n)$, for any $w \in [1, \log n]$, the compressed estimator p_φ achieves

$$R_{\frac{w}{n \log n}}(p_\varphi, p) = \mathcal{O}\left(\frac{1}{\sqrt{w}}\right).$$

Distribution functionals Section 7 explores the application of CML to functional estimation. The compressor maps each sample sequence to the multiset of multiplicities that are small relative to n , leading to a unified algorithm that optimally learns several symmetric functionals.

3. Preliminaries

General notation The rest of the paper considers two types of *univariate domains* \mathcal{X} , continuous, \mathbb{R} , and discrete, $[N] := \{1, \dots, N\}$. We consider distributions p over \mathcal{X} . Given a domain partition \mathcal{I} , the *quantized distribution* $p_\mathcal{I}$ assigns to each part $I \in \mathcal{I}$ probability $p(I)$.

More generally, \mathcal{I} can be any collection of disjoint measurable subsets of $\mathcal{X} = \mathbb{R}$, and f can be any real measurable function over \mathbb{R} . Then, $f_\mathcal{I}$ becomes the *quantized function* that assigns any $I \in \mathcal{I}$ a value of $\int_{x \in I} f(x) dx$ and the remaining set $\mathcal{X} \setminus \bigcup_{I \in \mathcal{I}} I$ a probability mass of 0. Slightly abusing notation, we use $f_\mathcal{I}$ also to represent the *flattened distribution* that assigns to every x in any $I \in \mathcal{I}$ of positive measure $|I| > 0$, the value

$$f_\mathcal{I}(x) := \frac{f_\mathcal{I}(I)}{|I|}.$$

Finally, write $a \wedge b$ for $\min\{a, b\}$, and $a \vee b$ for $\max\{a, b\}$.

Function norms We utilize a few norms of functions. The ℓ_1 -norm that evaluates $\|f\|_1 := \int_{-\infty}^{\infty} |f(x)| dx$. Another norm that will come in handy is the \mathcal{A}_k -norm. Specifically, for any positive integer k , let \mathcal{I}_k denote the collection of all unions of k disjoint intervals. Then, for any measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, the \mathcal{A}_k -norm of f is

$$\|f\|_{\mathcal{A}_k} := \sup_{I \in \mathcal{I}_k} \|f_I\|_1.$$

Note that the above notations *naturally extend* to $\mathcal{X} = [N]$ if we replace the integrals with the respective finite sums.

VC inequality A well-known convergence bound associated with the \mathcal{A}_k -norm is the VC inequality (Devroye & Lugosi, 2012). Often, the inequality appears as an upper bound on the expected \mathcal{A}_k -norm of the difference between the empirical and actual distributions.

We leverage the following variant in Acharya et al. (2017b), making the error probability explicit.

Lemma 2 (VC inequality). *For any $\varepsilon, \delta > 0$, and continuous distribution p , draw a sample $X^n \sim p$ and let \hat{p} denote the empirical distribution. For $n = \Theta((k + \log(1/\delta))/\varepsilon^2)$, with probability at least $1 - \delta$,*

$$\|p - \hat{p}\|_{\mathcal{A}_k} \leq \varepsilon.$$

4. Structured Continuous Distributions

A distribution collection \mathcal{C} is an ε -cover for \mathcal{P} under ℓ_1 -distance if every distribution in \mathcal{P} is within ℓ_1 -distance ε from some distribution in \mathcal{C} . The *crossing number* of two real functions f and g is the number of times they cross each other, or equivalently, the number of sign changes of $f(x) - g(x)$. The crossing number of a distribution collection is the largest crossing number of any pair of distributions in the collection.

Numerous important distribution families have ε -covers whose crossing number grows moderately as ε decreases. For example, t -piecewise, degree- d polynomials have 0-covers with crossing number $\mathcal{O}(t(d+1))$, and t -mixtures of univariate Gaussians have ε -covers with crossing number $\mathcal{O}(t \log(1/\varepsilon))$ growing only logarithmically in ε .

4.1. Compressor

Let \mathcal{P} have an ε -cover with crossing number s . Draw a sample Y^n from an unknown $p \in \mathcal{P}$ and let \hat{q} be the empirical distribution. Partition the real line into $t := s/\varepsilon$ intervals, $\mathcal{I} := (I_i)_{i=1}^t$, such that $\hat{q}_{\mathcal{I}}$ is uniform, where for simplicity, we assume that s/ε is an integer.

Next, draw an independent sample $X^n \sim p$. Let \hat{p} denote the empirical distribution, and let the compressor φ be its flattened version over \mathcal{I} ,

$$\varphi(X^n) := \hat{p}_{\mathcal{I}}.$$

Note that $\varphi(X^n)$ depends implicitly on Y^n , and its definition consists of both the Y^n partition and the empirical probabilities of each part according to X^n .

4.2. CML Estimator

Given a sample Y^n and the corresponding interval partition \mathcal{I} , the respective CML estimator is

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \Pr_{X^n \sim p} (\varphi(X^n) \mid \mathcal{I}).$$

Note that the probability term corresponds to a multinomial distribution. We can rewrite the CML estimator as

$$p_\varphi = \arg \min_{p \in \mathcal{P}} H(\hat{p}_{\mathcal{I}}, p_{\mathcal{I}}).$$

Hence, CML minimizes the cross entropy between the quantized empirical distribution and the actual one.

Theorem 2. *For any ε and $p \in \mathcal{P}$, draw $(X^n, Y^n) \sim p$. If $n = \Theta(t \log(t/\varepsilon)/\varepsilon^2)$, with probability at least $1 - 2\varepsilon$,*

$$\|p - p_\varphi\|_1 \leq 34\varepsilon.$$

Note that by traditional VC theory, the empirical \mathcal{A}_k -norm minimizer based on Y^n alone achieves sample complexity of $\Theta(t/\varepsilon)$. However, this method is usually not used for learning structured densities. The theorem on the other hand shows that the practically ubiquitous ML technique can also be applied to learn structured continuous distributions with a rigorous guarantee. It thereby demonstrates a different, continuous, application of Theorem 1 beyond the original application to discrete domains. The use of the samples X^n , while not strictly necessary, facilitates the application of the competitive argument.

The rest of the section is devoted to the theorem's proof.

4.3. Typical Events

First, we leverage standard concentration inequalities to establish several claims, each holds with high probability. The subsequent analysis will assume all claims hold.

By construction, each interval $I_i \in \mathcal{I}$ carries an empirical probability mass of $\hat{q}(I_i) = \varepsilon/s$. The respective probability mass under the actual density, $p(I_i)$, turns out to follow a $\text{Beta}(n/t, n(1 - 1/t))$ distribution. The following inequality (Hao et al., 2020) shows the typical values of $p(I_i)$.

Lemma 3. *For any $\alpha \in [0, 1]$ and $1 \leq t \leq n$,*

$$\Pr \left(\left| p(I_i) - \frac{1}{t} \right| \geq \frac{\alpha}{\sqrt{t}} \right) \leq e^{-\frac{n\alpha^2}{2}} + e^{-\frac{n\alpha^2}{2(1+\alpha\sqrt{t})}}.$$

For any $\gamma \in (0, 1)$ and $n \geq 12t \log(2t/\gamma)$, choosing $\alpha = 1/(2\sqrt{t})$ in the lemma yields

$$\Pr \left(\exists i \in [t], p(I_i) \notin \left(\frac{1}{2t}, \frac{3}{2t} \right) \right) \leq \gamma.$$

Henceforth, we assume that $n \geq 12t \log(2t/\gamma)$ and for each i , the actual probability mass of I_i falls in $(1/(2t), 3/(2t))$.

Next, consider the second sample X^n and an arbitrary index $i \in [t]$. By independence, $n\hat{p}(I_i) \sim \text{bin}(n, p(I_i))$, where $p(I_i) > 1/(2t)$ as above. From the binomial Chernoff and union bounds, for all $i \in [k]$, with probability at least $1 - \gamma$,

$$|\hat{p}(I_i) - p(I_i)| \leq \sqrt{\frac{3p(I_i)}{n} \log \frac{2t}{\gamma}}.$$

Our lower bounds on n and $p(I_i)$ also imply that for all i ,

$$\hat{p}(I_i) \leq p(I_i) + \sqrt{\frac{3p(I_i)}{n} \log \frac{2t}{\gamma}} \leq \frac{1 + \sqrt{2}}{\sqrt{2}} \cdot p(I_i) < \frac{3}{t}.$$

Below we *assume* that \hat{p} satisfies these inequalities.

4.4. Guarantees of CML

Given Y^n 's empirical distribution \hat{q} , hence also \mathcal{I} , there are finitely many flattened X^n -distributions \hat{p}_x over \mathcal{I} . Given a probability threshold δ , we consider two disjoint cases according to the probability of a particular distribution \hat{p}_x .

Likely \hat{p}_x : The probability $p(\hat{p}_x) > \delta$. We show that for such \hat{p}_x , with high probability, the CML distribution p_φ is close to p . The following argument assumes that the claims made in Section 4.3 hold.

First, since $\hat{p}(I_i) = \hat{p}_x(I_i) < 3/t$ for all i , then for any distribution p satisfying $\|p - \hat{p}_x\|_{\mathcal{A}_s} > 7\varepsilon$,

$$\|p - \hat{p}\|_{\mathcal{A}_s} \geq \|p - \hat{p}_x\|_{\mathcal{A}_s} - \|\hat{p}_x - \hat{p}\|_{\mathcal{A}_s} > 7\varepsilon - 2s \cdot \frac{3}{t} = \varepsilon.$$

By the VC inequality, this event happens with probability at most $\delta/2$ for a sample size of $n = \Theta((s + \log(2/\delta))/\varepsilon^2)$, contradicting our assumption that $p(\hat{p}_x) > \delta$. Therefore,

$$\|p - \hat{p}_x\|_{\mathcal{A}_s} \leq 7\varepsilon.$$

On the other hand, a fundamental attribute of the ML method is that $p_\varphi(\hat{p}_x) \geq p(\hat{p}_x) \geq \delta$. Hence,

$$\|p_\varphi - \hat{p}_x\|_{\mathcal{A}_s} \leq 7\varepsilon.$$

For any distribution $q \in \mathcal{P}$, we denote by q' the closest distribution in the ε -cover mentioned above under ℓ_1 distance. From the above and the triangle inequality,

$$\begin{aligned} \|p_\varphi - p\|_1 &\leq \|p'_\varphi - p'\|_1 + 2\varepsilon \\ &\leq 2(\|p - \hat{p}_x\|_{\mathcal{A}_s} + \|p_\varphi - \hat{p}_x\|_{\mathcal{A}_s}) + 6\varepsilon \\ &\leq 34\varepsilon, \end{aligned}$$

where the second inequality follows as

$$\|q_1 - q_2\|_1 = 2 \sup_{A \in \mathcal{R}} |q_1(A) - q_2(A)|.$$

Unlikely \hat{p}_x : The probability $p(\hat{p}_x) \leq \delta$. Note that the number of possible \hat{p}_x 's is at most

$$\binom{t+n-1}{t} \leq \left(\frac{e(t+n-1)}{t} \right)^t \leq e^{t \log(e(1+\frac{n}{t}))}.$$

Denote by $L_{t,n}$ the exponent of the last term. Setting $\delta = e^{-2L_{t,n}}$ and $n = \Theta(t \log(t/\varepsilon)/\varepsilon^2)$, the total probability of \hat{p}_x 's falling into this category is at most

$$e^{L_{t,n}} \cdot e^{-2L_{t,n}} = e^{-L_{t,n}} \leq \varepsilon^{-t},$$

where the last step follows by choosing a sufficiently large absolute constant in the above asymptotic expression for n .

Summary Choosing $\gamma = \varepsilon/2$ in Section 4.3, if $n = \Theta(t \log(t/\varepsilon)/\varepsilon^2)$, the union bound implies that with probability at least $1 - 2\varepsilon$,

$$\|p_\varphi - p\|_1 \leq 34\varepsilon.$$

5. Structured Discrete Distributions

Moving from \mathbb{R} to the discrete domain $[N]$, we apply CML to learn structured discrete distributions. Again, we assume that \mathcal{P} has an ε -cover \mathcal{C} with a crossing number at most s .

5.1. Reduction from Continuous CML

First, we present a concrete *random mapping* from discrete to continuous domains that enables the use of the continuous CML estimator.

For any distribution p over $[k]$, we define by \tilde{p} its *flattened version*, a continuous distribution that assigns

$$\tilde{p}(x) := p(\lceil x \rceil), \quad \forall x \in (0, k],$$

and $\tilde{p}(x) := 0$ for $x \notin (0, k]$. This notation yields a bijective mapping that maintains the number of sign changes of the difference between any two distributions.

Note that we only have sample access to the underlying distribution p . To apply the CML estimator in Section 4, we need to simulate a sample from \tilde{p} from $X^n \sim p$. A random mapping \mathcal{S} for this purpose counts the number of times each symbol $i \in [N]$ observed in X^n , say n_i , and respectively draws n_i independent sample points from the uniform distribution over $(i-1, i]$.

Now, it is straightforward to apply the continuous CML. Specifically, draw samples $(X^n, Y^n) \sim p$, define the compressor φ based on $\mathcal{S}(Y^n)$, and let the CML estimate be

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \Pr_{X^n \sim p} (\varphi(\mathcal{S}(X^n)) \mid \mathcal{I}).$$

5.2. Discrete CML

Next, we present an alternative CML method that applies directly to discrete samples.

Compressor Draw a size- n sample $X^n \sim p$, and denote by \hat{p} the empirical distribution. Then, sort elements in X^n in a non-descending order, say, $X_{(1)}, \dots, X_{(n)}$, such that

$$X_{(i)} \leq X_{(j)}, \quad \forall i < j.$$

Without loss of generality, assume that $t := s/\varepsilon$ is an integer and n is divisible by t . Sequentially partition the sequence into t sub-sequences, such that the i -th sub-sequence contains observations with indices from $(i-1)n/t + 1$ to in/t .

The compressor φ maps every sample X^n to the boundary symbols that of these sub-sequences. Specifically,

$$\varphi(X^n) := (X_{(in/t)})_{i=1}^t,$$

which is *essentially* \hat{p}_T with $\mathcal{I} := (I_i)_{i=1}^t$ and

$$I_i := [X_{((i-1)n/t)} : X_{(in/t)}], \forall i \in [t].$$

CML estimator Consequently, the CML estimator for the above compressor takes the form

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \Pr_{X^n \sim p} (\varphi(X^n) \mid \mathcal{I}).$$

Equivalently, the CML estimator can be written as

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \sum_{y^n: \varphi(y^n) = \varphi(X^n)} \prod_{i=1}^n p(y_i).$$

The estimator often performs well, as described below.

Theorem 3. For $p \in \mathcal{P}$ and $X^n \sim p$, if $n = \Theta((t \wedge N) \log(t \vee N)/\varepsilon^2)$, with probability at least $1 - e^{-t \wedge N}$,

$$\|p - p_\varphi\|_1 = \mathcal{O}(\varepsilon).$$

5.3. Hypothesis Selection Algorithm

First, we utilize the assumptions and the standard VC inequality to show the *existence* of a hypothesis selection algorithm that, with high probability, finds an accurate estimate of the underlying distribution.

By the VC inequality, for any $\varepsilon, \delta > 0$, and $n = \Theta((s + \log(2/\delta))/\varepsilon^2)$, with probability at least $1 - \delta/2$,

$$\|\hat{p} - p\|_{\mathcal{A}_s} \leq \varepsilon.$$

Denote by \mathcal{I}_s the set of unions of at most s disjoint intervals in \mathcal{I} . By our construction, for any set $I \in \mathcal{I}_s$,

$$\|\hat{p}_I\|_1 \leq \frac{1}{t} \cdot s = \varepsilon.$$

For any $q \in \mathcal{P}_{s,\varepsilon}$, let q' be a distribution in the ε -cover \mathcal{C} with minimal ℓ_1 -distance to q . For any distribution pair $q_1, q_2 \in \mathcal{P}_{s,\varepsilon}$, define the following *variant* of the Scheffé set (Devroye & Lugosi, 2012) as

$$S'_{12} := \{x \in \mathbb{R} : q'_1(x) > q'_2(x)\},$$

and the original Scheffé set as S_{12} .

Given the previous inequality and distribution \hat{p}_T , we can approximate $\hat{p}(q_1 \geq q_2) := \hat{p}(S_{12})$ to a 3ε additive error by summing up $\hat{p}_T(I_i)$ over indices i satisfying $I_i \in S'_{12}$. Note that we assumed that $q'_1 - q'_2$ has at most s sign changes.

Next, we show that if the VC inequality holds, there exists a selection algorithm that uses \hat{p}_T to find a density $p^* \in \mathcal{C}$ satisfying $\|p^* - p\| \leq \mathcal{O}(\varepsilon)$. In addition, this algorithm is a variant of that in Theorem 6.4 of Devroye & Lugosi (2012).

Let \mathcal{S}' denote the collection of all the modified Scheffé sets induced by distributions in \mathcal{C} . Let p^* be the distribution q in \mathcal{C} minimizing $\sup_{S \in \mathcal{S}'} |\hat{p}_T(S) - q(S)|$ up to an additive 4ε error, where we enlarged the error bound by ε to guarantee the existence of such a distribution. Then, the triangle inequality yields

$$\|p - p^*\|_1 \leq \|p - p'\|_1 + \|p' - p^*\|_1.$$

For any pair (q_1, q_2) of distributions, define

$$D_S(q_1, q_2) := \sup_{S \in \mathcal{S}'} |q_1(S) - q_2(S)|.$$

Consider the last term in the above inequality. By a recursive application of the triangle inequality and properties embedded in the prior constructions,

$$\begin{aligned} \|p' - p^*\|_1 &= 2D_S(p', p^*) \\ &\leq 2D_S(p', \hat{p}) + 2D_S(\hat{p}, p^*) \\ &\leq 2D_S(p', \hat{p}_T) + 2D_S(\hat{p}_T, p^*) + \mathcal{O}(\varepsilon) \\ &\leq 4D_S(p', \hat{p}_T) + \mathcal{O}(\varepsilon) \\ &\leq 4D_S(p', \hat{p}) + \mathcal{O}(\varepsilon) \\ &\leq 4D_S(p', p) + 4D_S(p, \hat{p}) + \mathcal{O}(\varepsilon) \\ &\leq 2\|p' - p\|_1 + 4D_S(p, \hat{p}) + \mathcal{O}(\varepsilon). \end{aligned}$$

Recall that \mathcal{C} is an ε -cover of $\mathcal{P}_{s,\varepsilon}$, we have $\|p' - p\|_1 \leq \varepsilon$. In addition, the collection \mathcal{S}' of modified Scheffé sets has a VC dimension of at most $2s + 2$. Therefore by the VC inequality, for a sample size of $n = \Theta((s + \log(2/\delta))/\varepsilon^2)$, with probability at least $1 - \delta/2$,

$$D_S(p, \hat{p}) \leq \varepsilon.$$

Consolidating these results shows that with probability at least $1 - \delta$, the selected hypothesis p^* achieves

$$\|p - p^*\|_1 = \mathcal{O}(\varepsilon).$$

Note, however, that the above selection algorithm does not provide a method for finding p^* .

5.4. Guarantees of CML

There are only finitely many possible \hat{p}_T , or equivalently, $\varphi(X^n)$. Below, we consider two disjoint cases according to the probability of observing a particular *pattern* of $\varphi(X^n)$ under the actual distribution.

Likely pattern: Pattern ϕ whose probability $p(\phi) := \Pr_{X^n \sim p}(\varphi(X^n) = \phi) > \delta$. By the argument in Section 5.3,

there exists a selection algorithm that maps this pattern to a distribution p^* whose error probability in achieving

$$\|p - p^*\|_1 = \mathcal{O}(\varepsilon)$$

is at most δ . Since $p(\phi) > \delta$ by assumption, the selected hypothesis must satisfy the error bound above.

It is straightforward to combine these error bounds through Theorem 1 and the triangle inequality,

$$|p - p_\varphi| \leq |p_\varphi - p^*| + |p - p^*| = \mathcal{O}(\varepsilon).$$

Unlikely pattern: Pattern ϕ whose probability $p(\phi) \leq \delta$. Since ϕ can only be a sorted integer sequence with values in $[N]$, the number of possible patterns equals the number of ways to pick t elements from $[N]$ with repetition, which is

$$\begin{aligned} \binom{t+N-1}{t} &\leq \left(\frac{e(t+N)}{t}\right)^t \bigwedge \left(\frac{e(N+t)}{N}\right)^k \\ &\leq \exp\left((t \wedge N) \log\left(e\left(1 + \frac{N}{t} \vee \frac{t}{N}\right)\right)\right). \end{aligned}$$

Denote by $L_{t,n}$ the exponent of the last term. Setting $\delta = e^{-2L_{t,n}}$ and $n = \Theta(L_{t,n}/\varepsilon^2)$, the total probability of patterns falling into this category is at most

$$e^{L_{t,n}} \cdot e^{-2L_{t,n}} = e^{-L_{t,n}} \leq e^{-t \wedge N},$$

where the last step follows by choosing a sufficiently large absolute constant in the expression for n .

6. Probability Multisets

In this section, we address probability multiset estimation under permutation invariance.

6.1. Compressor and PML

For any discrete distribution p and sample $X^n \sim p$, let compressor φ map the sequence to its *profile* $\varphi(X^n)$, defined as the *multiplicity multiset* of symbols appearing in X^n . The respective CML estimator, introduced in Orlitsky et al. (2004) as the *PML estimator*. Specifically, PML computes

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \sum_{y^n: \varphi(y^n) = \varphi(X^n)} \prod_{i=1}^n p(y_i).$$

For any $\tau \in [0, 1]$, define the τ -truncated relative earth-mover distance (Valiant & Valiant, 2015) between two discrete distributions p and q as

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

where $\Gamma_{p,q}$ represents all the possible couplings of these two distributions. In the following, we show that the PML distribution estimator satisfies

Theorem 4. For any discrete distribution p , draw a sample $X^n \sim p$. With probability at least $1 - 4 \exp(-\Omega(n^{\frac{1}{3}}))$, for any $w \in [1, \log n]$,

$$R_{\frac{w}{n \log n}}(p_\varphi, p) = \mathcal{O}\left(\frac{1}{\sqrt{w}}\right).$$

For any $\tau \in [0, 1]$, define the τ -truncated sorted ℓ_1 -distance between two distributions $p, q \in \Delta_{\mathcal{X}}$ as

$$\tilde{\ell}_\tau(p, q) := \min_{p' \in \Delta_{\mathcal{X}}: \{p'\} = \{p\}} \sum_x |p'(x) \vee \tau - q(x) \vee \tau|,$$

where $\{p\}$ denotes the probability multiset of p . By Fact 1 in Valiant & Valiant (2015), for any distributions $p, q \in \Delta_{\mathcal{X}}$, $\tilde{\ell}_\tau(p, q) \leq 2R_\tau(p, q)$, implying

Corollary 2. Under the same conditions as Theorem 4,

$$\tilde{\ell}_{\frac{w}{n \log n}}(p_\varphi, p) = \mathcal{O}\left(\frac{1}{\sqrt{w}}\right).$$

6.2. Proof of Theorem 4

This section provides a sketch for the proof of Theorem 4. We relegate most technical details to Appendix A. Part of the proof is adapted from Theorem 2 in Valiant & Valiant (2015). The original reasoning is not sufficient for our purpose as the error probability derived is too large to invoke the competitiveness of PML. For this reason, we modify the linear program used in the paper, carefully separate the analysis of the estimators for large and small probabilities, and provide a refined analysis with tighter probability bounds by reducing union-bound arguments.

First, we define histograms and the relative earth-moving cost and give operational meaning to R_τ . For a distribution p , the *histogram* of a multiset $\mathcal{A} \subseteq \{p\}$ is a mapping, denoted by $h_{\mathcal{A}} : (0, 1] \rightarrow \mathbb{Z}_{\geq 0}$, that maps each number $y \in (0, 1]$ to the number of times it appears in \mathcal{A} . Note that every y corresponds to a probability mass of $y \cdot h_{\mathcal{A}}(y)$. More generally, we also allow *generalized histograms* h with non-integral values $h(y) \in \mathbb{R}_{\geq 0}$.

For any $y_1, y_2 \in (0, 1]$, generalized histogram h , and non-negative $m < y_1 \cdot h(y_1)$, we can *move* a probability mass from location y_1 to y_2 by reassigning $h(y_1) - m/y_1$ to y_1 , and $h(y_2) + m/y_2$ to y_2 . Given $\tau \in [0, 1]$, we define the *cost* associated with this operation as

$$c_{\tau,m}(y_1, y_2) := m \cdot \left| \log \frac{y_1 \vee \tau}{y_2 \vee \tau} \right|,$$

and term it as τ -truncated earth-moving cost. The cost of multiple operations is additive. Note that $R_\tau(p, q)$ is the minimal total τ -truncated earth-moving cost associated with any operation schemes of moving $h_{\{p\}}$ to yield $h_{\{q\}}$.

For simplicity, we suppress X^n in $\varphi_i(X^n)$ and $\mu_s(X^n)$, representing respectively the *number of symbols appearing i times* and the *number of times symbol s appears*.

For any absolute constants B and C satisfying $0.1 > B > C > \frac{B}{2} > 0$, define $x_n := \frac{n^B + n^C}{n}$ and $S := \{\frac{1}{n^2}, \frac{2}{n^2}, \dots, x_n\}$. Consider the following linear program.

For each $x \in S$, define the associated variable v_x

Minimize $\sum_{i=1}^{n^B} \left| \varphi_i - \sum_{x \in S} \text{bin}(n, x, i) \cdot v_x \right|$

s.t. $\sum_{x \in S} x \cdot v_x = \sum_{i \leq n^B + 2n^C} \frac{i}{n} \cdot \varphi_i$

and $\forall x \in S, v_x \geq 0$

Figure 1. Linear program (LP)

EXISTENCE OF A GOOD FEASIBLE POINT

Let p be the underlying distribution and h be its histogram. First, we show that with *high* probability, the linear program LP has a feasible point (v_x) that is *good* in the following sense: 1) the corresponding objective value is relatively *small*; 2) for $\tau \geq n^{-3/2}$, the generalized histogram $h_0 : x \rightarrow v_x$ is *close* to $h_n : y \rightarrow h(y) \cdot \mathbb{1}_{y \leq x_n}$, admitting a low τ -truncated earth-mover cost.

In the appendix, we leverage the Chernoff bound and union bound to show that with probability at least $1 - \exp(-n^{\frac{1}{3} + \kappa})$, the objective value of the feasible point (v_x) is at most $n^B \cdot \mathcal{O}(n^{\frac{2}{3} + \kappa} + 1) = \mathcal{O}(n^{B + \frac{2}{3} + \kappa})$.

For any $\tau \geq n^{-3/2}$, the minimal τ -truncated earth-moving cost of moving the generalized histogram h_0 corresponding to (v_x), and the histogram $h_n : y \rightarrow h(y) \cdot \mathbb{1}_{y \leq x_n}$, so that they differ from each other only at $x = x_n$, is at most

$$\log \left(\frac{n^{-3/2} + n^{-2}}{n^{-3/2}} \right) + \mathcal{O} \left(\frac{\log n}{n^{\frac{1}{3} - \kappa}} \right) = \mathcal{O}(n^{-\frac{1}{3} + 2\kappa}).$$

ALL SOLUTIONS ARE GOOD SOLUTIONS

Let (v_x) be the solution described above. The appendix then proceeds to show that for any solution (v'_x) to LP whose objective value is $\mathcal{O}(n^{B + \frac{2}{3} + \kappa})$, the generalized histogram h_1 corresponding to (v'_x) is close to h_0 .

Specifically, the proof establishes a $\mathcal{O}(1/\sqrt{\log n})$ discrepancy bound whenever $B = 1.5C = 10\kappa = 0.01$.

Consolidate the previous results. For $w \in [1, \log n]$ and $\tau = w/(n \log n)$, with probability at least $1 - \exp(-n^{\frac{1}{3} + \kappa})$, the solution to LP will yield a generalized histogram h_1 , such that the minimal τ -truncated earth-moving cost of moving h_1 and h_n so that they differ only at x_n , is $\mathcal{O}(1/\sqrt{w})$.

COMPETITIVENESS OF PML

For the PML distribution associated with a sample satisfying our prior assumptions, denote by h_n^{PML} the histogram corresponding to its entries that are at most x_n .

By a recent result in (Han & Shiragur, 2021), for any $w \in [1, \log n]$ and $\tau = w/(n \log n)$, the minimal τ -truncated cost of moving h_n^{PML} and h_n so that they differ only at x_n , is $\mathcal{O}(1/\sqrt{w})$, with probability at least $1 - 2 \exp(-\Omega(n^{\frac{1}{3}}))$.

PROPERTIES OF THE EMPIRICAL HISTOGRAM

Denote by h^{EMP} the empirical histogram. By the Chernoff bound, for any symbol s , the probability that

$$|n \cdot p(s) - \mu_s| \geq \mu_s^{3/4} \text{ and } \mu_s > n^B + 2n^C$$

is at most $2np(s) \exp(-\Omega(n^{2C-B}))$, and similarly,

$$n \cdot p(s) \geq n^B + 4n^C \text{ and } \mu_s \leq n^B + 2n^C$$

will happen with probability at most $2 \exp(-\Omega(n^{2C-B}))$. Hence, we assume that $|n \cdot p(s) - \mu_s| < \mu_s^{3/4}$ for all symbols s appearing more than $n^B + 2n^C$ times, and that any symbol s with probability $p(s) \geq (n^B + 4n^C)/n$ appears more than $n^B + 2n^C$ times. By the union bound, we will be correct with probability at least $1 - 4n \exp(-\Omega(n^{2C-B}))$.

Let $y_n := (n^B + 4n^C)/n$ for notational convenience. If for each symbol s satisfying $\mu_s \geq n^B + 2n^C$, we move a μ_s/n probability mass of h^{EMP} from μ_s/n to $p(s)$, then at all locations $y \geq y_n$, the total discrepancy between the resulting generalized histogram and the actual one is at most

$$\sum_{j > n^B + 2n^C} \varphi_j \frac{j^{\frac{3}{4}}}{n} = \frac{1}{n} \sum_{j > n^B + 2n^C} \varphi_j^{\frac{1}{4}} (\varphi_j j)^{\frac{3}{4}} \leq n^{-\frac{B}{4}},$$

where the last step follows by Hölder's inequality. Moreover, the associated total earth-moving cost is at most

$$\sum_{j > n^B + 2n^C} \varphi_j \frac{j}{n} \log \left| \frac{j}{j \pm j^{\frac{3}{4}}} \right| \leq \sum_{j > n^B + 2n^C} \varphi_j \frac{j^{\frac{3}{4}}}{n} \leq n^{-\frac{B}{4}}.$$

The y_n -truncated earth-mover distance between h^{EMP} and h is thus at most $2n^{-\frac{B}{4}} \log n = \mathcal{O}(n^{-\frac{B}{5}})$, which, together with the above error probability bound, upper bounds the expected value of $R_{y_n}(h^{\text{EMP}}, h)$ by $\mathcal{O}(n^{-\frac{B}{5}}) + 4n \exp(-\Omega(n^{2C-B})) \log n = \mathcal{O}(n^{-\frac{B}{5}})$.

Moreover, changing any element in the sample sequence changes the value of $R_{y_n}(h^{\text{EMP}}, h)$ by at most $(\log n)/n$. Hence, by McDiarmid's inequality, with probability at least $1 - 2 \exp(-2\sqrt{n})$, the value of $R_{y_n}(h^{\text{EMP}}, h)$ is less than $\mathcal{O}(n^{-\frac{B}{5}}) + n^{-\frac{1}{4}} \log n = \mathcal{O}(n^{-\frac{B}{5}})$.

COMPETITIVENESS OF PML

Consider the PML histogram h^{PML} and its y_n -truncated earth-mover distance to h . Since there are at most $\exp(3\sqrt{n})$ different profiles (Hardy & Ramanujan, 1918), with probability at least $1 - 2\exp(-\Omega(\sqrt{n}))$,

$$R_{y_n}(h^{\text{PML}}, h) \leq \mathcal{O}(n^{-\frac{B}{5}}).$$

PERFORMANCE OF PML

We consolidate the previous results. For any $w \in [1, \log n]$ and $\tau = w/(n \log n)$, we design an earth-moving scheme that moves h^{PML} to h .

First, with probability at least $1 - 2\exp(-\Omega(n^{\frac{1}{3}}))$, we can move the probability mass of h_n^{PML} so that it differs from h_n only at $x_n = (n^B + n^C)/n$ while incurring a τ -truncated earth-mover cost of at most $\mathcal{O}(1/\sqrt{w})$.

Second, by the empirical-histogram argument and $x_n < y_n$, with probability at least $1 - 2\exp(-\Omega(n^{\frac{1}{3}}))$, we can further move the probability mass of h^{PML} at locations $y \geq x_n$ to coincide with h above y_n while incurring a loss of $\mathcal{O}(n^{-\frac{B}{5}})$.

After the previous two steps, the modified PML histogram differs from the actual histogram h only at locations $y \in I_n = [x_n, y_n]$. Note that the cost of moving a unit mass within I_n is at most $3n^{C-B}$, implying that with probability at least $1 - 2\exp(-\Omega(n^{\frac{1}{3}}))$,

$$R_\tau(h^{\text{PML}}, h) \leq \mathcal{O}\left(\frac{1}{\sqrt{w}} + n^{-\frac{B}{5}}\right) + 3n^{C-B} = \mathcal{O}\left(\frac{1}{\sqrt{w}}\right).$$

7. Symmetric Functionals

In this section, we apply the idea of CML to estimate multi-sets of low probabilities and leverage the respective plug-in estimator to approximate several distribution functionals.

The study of functional estimation dates back more than half a century (Carlton, 1969; Good, 1953; Good & Toulmin, 1956) and has steadily grown over the years. While the empirical distribution plug-in estimator performs well in the large-sample regime, modern data science applications often study high-dimensional data, for which more sophisticated methods lead to estimators that possess better guarantees (Jiao et al., 2015; Orlitsky et al., 2016; Valiant & Valiant, 2011a;b; Wu & Yang, 2016; Hao et al., 2018; Wu & Yang, 2019; Hao & Orlitsky, 2019b; 2020a;b).

Recently, a line of research works studied the use of PML in functional estimation (Orlitsky et al., 2004; 2011; Das, 2012; Acharya et al., 2012; 2017a; Hao & Orlitsky, 2019a; 2020b; Charikar et al., 2019a;b; Han & Shiragur, 2021).

For functionals addressed in this section, the plain PML is known to be sample-optimal only for additive errors

$\varepsilon \geq 1/n^{1/3}$. Concurrently with our result, Charikar et al. (2019c) proposed a different PML-type functional estimation method with optimal sample complexity down to $\varepsilon \gg 1/\sqrt{n}$ accuracy. We note that some specialized estimators also achieve better accuracy of order $1/\sqrt{n}$.

Recently, an efficient algorithm (Anari et al., 2020) was proposed to compute the PML-type estimators in Charikar et al. (2019c). The techniques also apply to the algorithms presented in this section.

7.1. Compressor

For any distribution p over \mathcal{X} , sample $X^n \sim p$, and multiplicity i , recall that $\varphi_i(X^n)$ denotes the number of symbols appearing exactly i times.

Our objective is to extract information about low probabilities of p , regardless of symbol permutations. It is natural to consider, for an integer $t \leq n$, the compressor

$$\varphi(X^n) := (\varphi_i(X^n))_{i=1}^t,$$

where t determines the inference horizon.

7.2. CML Estimator

Similar to previous sections, for fixed t and sample $X^n \sim p$, the CML estimator for the above compressor takes the form

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \Pr_{X^n \sim p}(\varphi(X^n)).$$

Equivalently, the CML estimator can be written as

$$p_\varphi := \arg \max_{p \in \mathcal{P}} \sum_{y^n: \varphi(y^n) = \varphi(X^n)} \prod_{i=1}^n p(y_i).$$

For space considerations, we relegate the rest of this section and the paper's appendix to the supplementary material.

Conclusion

The paper proposes a simple, novel, and unified compressed maximum likelihood (CML) approach for several fundamental tasks. The new technique bridges algorithms and results in several research directions over both discrete and continuous domains.

Acknowledgements

We thank the reviewers for their helpful comments and are grateful to the National Science Foundation for supporting this work through grants CIF-1564355 and CIF-1619448.

We also thank an anonymous reviewer for an excellent summary of our contribution, incorporated into Section 2.1, and another anonymous reviewer for suggesting the approximate CML estimator (Definition 1).

References

- Acharya, J., Das, H., Jafarpour, A., Orlitsky, A., and Pan, S. Estimating multiple concurrent processes. In *Proceedings 2012 IEEE International Symposium on Information Theory (ISIT)*, pp. 1628–1632. IEEE, 2012.
- Acharya, J., Das, H., Orlitsky, A., and Suresh, A. T. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. In *International Conference on Machine Learning*, pp. 11–21, 2017a.
- Acharya, J., Diakonikolas, I., Li, J., and Schmidt, L. Sample-optimal density estimation in nearly-linear time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1278–1289. SIAM, 2017b.
- Anari, N., Charikar, M., Shiragur, K., and Sidford, A. Instance based approximations to profile maximum likelihood. *Advances in neural information processing systems*, 2020.
- Bishop, C. M. *Pattern recognition and machine learning*. springer, 2006.
- Carlton, A. G. On the bias of information estimates. *Psychological Bulletin*, 71(2):108, 1969.
- Charikar, M., Shiragur, K., and Sidford, A. 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.
- Charikar, M., Shiragur, K., and Sidford, A. Efficient profile maximum likelihood for universal symmetric property estimation. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 780–791, 2019b.
- Charikar, M., Shiragur, K., and Sidford, A. A general framework for symmetric property estimation. *Advances in neural information processing systems*, 2019c.
- Das, H. *Competitive tests and estimators for properties of distributions*. PhD thesis, UC San Diego, 2012.
- Devroye, L. and Lugosi, G. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.
- Friedman, J., Hastie, T., Tibshirani, R., et al. *The elements of statistical learning*. Number 10 in 1. Springer series in statistics New York, 2001.
- Good, I. J. The population frequencies of species and the estimation of population parameters. *Biometrika*, 40(3-4): 237–264, 1953.
- Good, I. J. and Toulmin, G. H. The number of new species, and the increase in population coverage, when a sample is increased. *Biometrika*, 43(1-2):45–63, 1956.
- Han, Y. and Shiragur, K. On the competitive analysis and high accuracy optimality of profile maximum likelihood. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1317–1336. SIAM, 2021.
- Han, Y., Jiao, J., and Weissman, T. Local moment matching: A unified methodology for symmetric functional estimation and distribution estimation under wasserstein distance. In *Conference On Learning Theory*, pp. 3189–3221, 2018.
- Hao, Y. and Orlitsky, A. The broad optimality of profile maximum likelihood. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 10989–11001, 2019a.
- Hao, Y. and Orlitsky, A. Unified sample-optimal property estimation in near-linear time. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 11106–11116, 2019b.
- Hao, Y. and Orlitsky, A. Data amplification: Instance-optimal property estimation. In *International Conference on Machine Learning (ICML)*, pp. 4049–4059. PMLR, 2020a.
- Hao, Y. and Orlitsky, A. Profile entropy: A fundamental measure for the learnability and compressibility of discrete distributions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020b.
- Hao, Y., Orlitsky, A., Suresh, A. T., and Wu, Y. Data amplification: A unified and competitive approach to property estimation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 8834–8843, 2018.
- Hao, Y., Jain, A., Orlitsky, A., and Ravindrakumar, V. SURF: A simple, universal, robust, fast distribution learning algorithm. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Hardy, G. H. and Ramanujan, S. Asymptotic formulae in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.
- Jiao, J., Venkat, K., Han, Y., and Weissman, T. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015.
- Orlitsky, A., Santhanam, N. P., Viswanathan, K., and Zhang, J. On modeling profiles instead of values. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 426–435. AUAI Press, 2004.

Orlitsky, A., Santhanam, N. P., Viswanathan, K., and Zhang, J. On estimating the probability multiset. *Online Draft*, 2011. URL <http://alon.ucsd.edu/papers/pml1.pdf>.

Orlitsky, A., Suresh, A. T., and Wu, Y. Optimal prediction of the number of unseen species. *Proceedings of the National Academy of Sciences*, 113(47):13283–13288, 2016.

Valiant, G. and Valiant, P. Estimating the unseen: An $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings 43rd Annual ACM Symposium on Theory of Computing (STOC)*, pp. 685–694. ACM, 2011a.

Valiant, G. and Valiant, P. The power of linear estimators. In *Proceedings 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 403–412. IEEE, 2011b.

Valiant, G. and Valiant, P. Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 685–694, 2011c.

Valiant, G. and Valiant, P. Instance optimal learning. *arXiv preprint, arXiv: 1504.05321*, 2015.

Valiant, G. and Valiant, P. Instance optimal learning of discrete distributions. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 142–155, 2016.

Valiant, P. and Valiant, G. Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems*, pp. 2157–2165, 2013.

Van der Vaart, A. W. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Wu, Y. and Yang, P. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, 2016.

Wu, Y. and Yang, P. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *The Annals of Statistics*, 47(2):857–883, 2019.