

Approximate Profile Maximum Likelihood

Dmitri S. Pavlichin

DMITRIP@STANFORD.EDU

*Department of Applied Physics
Stanford University
Stanford, CA 94305, USA*

Jiantao Jiao

JIANTAO@EECS.BERKELEY.EDU

*Department of Electrical Engineering and Computer Sciences
University of California
Berkeley, CA 94720, USA*

Tsachy Weissman

TSACHY@STANFORD.EDU

*Department of Electrical Engineering
Stanford University
Stanford, CA 94305, USA*

Editor: Sara van de Geer

Abstract

We propose an efficient algorithm for approximate computation of the profile maximum likelihood (PML), a variant of maximum likelihood maximizing the probability of observing a sufficient statistic rather than the empirical sample. The PML has appealing theoretical properties, but is difficult to compute exactly. Inspired by observations gleaned from exactly solvable cases, we look for an approximate PML solution, which, intuitively, clumps comparably frequent symbols into one symbol. This amounts to lower-bounding a certain matrix permanent by summing over a subgroup of the symmetric group rather than the whole group during the computation. We extensively experiment with the approximate solution, and the empirical performance of our approach is competitive and sometimes significantly better than state-of-the-art performances for various estimation problems.

Keywords: Profile maximum likelihood, dynamic programming, sufficient statistic, partition of multi-partite numbers, integer partition

1. Introduction

The maximum likelihood principle, proposed by Ronald Fisher, has proved to be a powerful and versatile method used throughout nearly all scientific fields. However, it is still the source of considerable controversy in the statistics and machine learning community. Indeed, quoting Efron (Efron, 1982):

“The controversy centers on the relationship between decision theory and maximum likelihood. Beginning with the Neyman-Pearson lemma, decision theory has reshaped the theory and practice of hypothesis testing...”

Indeed, Wald’s statistical decision theory (Wald, 1950) provided a comprehensive framework for evaluating and proposing statistical procedures. The maximum likelihood approach, which can be proved to be asymptotically efficient under mild conditions in the

Hájek-Le Cam theory (Van der Vaart, 2000, Chap. 9), lacks a finite-sample justification. Indeed, as Le Cam (Le Cam, 1979) argued, even in the asymptotic regime the maximum likelihood approach may have weird behavior and is not always consistent.

The recent work (Acharya et al., 2017) provided an elegant justification of the maximum likelihood approach which is intimately connected to decision theory. It was shown in (Acharya et al., 2017) that plugging-in the *profile maximum likelihood* (PML) distribution estimators into a variety of functionals achieves the optimal sample complexity in estimating those functionals, which the generic *sequence maximum likelihood* (SML) fails to achieve. We also refer the readers to (Anevski et al., 2017) for statistical properties of PML.

To compute the profile maximum likelihood estimator we must solve the following optimization problem: given n samples with empirical distribution $\hat{p} = (\hat{p}_1, \hat{p}_2, \dots)$, maximize the probability to observe \hat{p} up to *relabeling* $\sigma\hat{p} \triangleq (\hat{p}_{\sigma(1)}, \hat{p}_{\sigma(2)}, \dots)$ of the empirical distribution, where σ is some permutation. This amounts to computing the PML distribution p^* :

$$p^* = \operatorname{argmax}_p \sum_{\sigma} \mathbb{P}_p(\sigma\hat{p}) = \operatorname{argmax}_p \sum_{\sigma} e^{-nD(\sigma\hat{p}||p)} \quad (1)$$

where the max is over all discrete distributions p with a known support set size (we treat the unknown case later; it is similar), the sum is over all permutations σ of the support set of distribution p , and where $D(\cdot||\cdot)$ is the Kullback-Leibler divergence.

The profile maximum likelihood estimator is computationally challenging to find – as can be shown to be equivalent to optimizing a certain matrix permanent – and the best known algorithm has running time exponential in the support size of the unknown discrete distribution. Our main contributions are the following:

1. We present an efficiently computable approximation for the PML distribution. The approximation idea is motivated by observations for the exact PML for small alphabets, where the PML distribution tends to put symbols whose empirical counts are close (within $O(\sqrt{n})$) into the same level set¹. The idea is to lower bound the sum in (1) by summing over only the permutations that contribute “a lot” – namely the subgroup of permutations that only mix symbols within the level sets of p . This leads to the objective function in the argmax below (a lower bound to the objective function in (1)) to define the approximate PML (APML) distribution \bar{p}^* :

$$\bar{p}^* = \operatorname{argmax}_p \left(e^{-nD(\hat{p}||p)} \prod_{\alpha \in \mathcal{A}(p)} |\alpha|! \right) \quad (2)$$

where $\mathcal{A}(p)$ is the partition of the support set of p into the level sets α of p . The second factor of the objective function rewards clumping many symbols together into a few large level sets, while the first factor encourages similarity to the empirical distribution and dominates as the sample size $n \rightarrow \infty$.

We present a dynamic programming approach to compute the approximate PML distribution \bar{p}^* . Compared with existing approximations to the PML (Orlitsky et al., 2004b; Vontobel, 2012), our algorithm has no tuning parameters, is deterministic, runs

1. A level set of distribution p is $\{x : p_x = u\}$ for some value u .

in at most linear time in the sample size (usually computing the empirical histogram is the slowest part), and achieves state-of-the-art empirical performance in various estimation tasks that we detail below.

For the case of unknown support set size, we state the appropriate generalization of the PML distribution and its approximation. We modify our dynamic programming algorithm slightly to handle this case, preserving a linear worst-case runtime. It may happen that the PML and our approximate PML distributions have a discrete part with finite support and a “continuous part” in the terminology of (Orlitsky et al., 2004b); we are able to detect this case and find the discrete part of the approximate PML distribution and the probability mass of the continuous part.

2. Given the result in (Acharya et al., 2017) that plugging in the PML distribution into various functionals achieves the optimal sample complexity, we extensively experiment with the plug-in estimator using our APML distribution. We estimate entropy, Rényi entropy, support set size, L_1 distance to uniformity, and the sorted probability vector of a discrete distribution and compare with the state-of-the-art approaches with available code for estimating those functionals. We find that the performance of plugging in our APML distribution into those functionals is consistently competitive, and sometimes much better than state-of-the-art packages.
3. We extend our approximation scheme to the multi-dimensional PML problem, which is intimately connected to estimating divergence functions of discrete distributions. Utilizing results on partitions of bipartite numbers, we show that solving the two-dimensional PML problem leads to estimators that achieve the optimal sample complexity for estimating a variety of functionals such as the KL divergence, the L_1 distance, the squared Hellinger distance, and the χ^2 divergence. The multi-dimensional APML distribution turns out to be harder to compute than the one-dimensional APML, so we settle for a greedy heuristic to approximate the multi-dimensional APML distribution.
4. We provide extensive experimental results on the plug-in approach for estimating divergence functions using our approximation to the two-dimensional PML. This achieves competitive and sometimes much better results than state-of-the-art approaches in KL divergence and L_1 distance estimation.
5. We generalize the PML idea to general group actions, leading to the generalized PML approach for mutual information estimation. We analyze several candidates for the PML solution and show that one candidate makes the key arguments in (Acharya et al., 2017) fail, while the other succeeds.

We provide code at <https://github.com/dmitrip/PML> and (Pavlichin et al., 2017) for computing the approximate PML distributions.

This work is organized as follows. Section 2 reviews some results on profile maximum likelihood and functional estimation. Section 3 specializes the discussion of Section 2 to the setting of estimating functionals of discrete distributions, and defines the PML and multi-dimensional PML distributions. Section 4 presents our APML approach and observations about the exact PML that inspired it. Section 5 shows the results of numerical experiments

using the APML to estimate symmetric functionals of a distribution: the sorted probability vector, entropy, Rényi entropy, L_1 distance to uniformity, and support set size. Section 6 shows the results of numerical experiments using the APML to estimate symmetric functionals of multiple distributions: the KL divergence and L_1 distance. Appendices contain proofs, examples, most of the notation and algorithms related to the multi-dimensional APML (Appendix H), and application of the PML to estimation of mutual information (Appendix I).

2. The principle of profile maximum likelihood and partial sufficiency

We rephrase the key result in (Acharya et al., 2017) regarding the maximum likelihood principle in Theorem 3; similar ideas appeared in (Acharya et al., 2011). We define a statistical model

$$\mathcal{E} = \{\mathcal{X}, P_\theta, \theta \in \Theta\}, \quad (3)$$

where the observation $X \sim P_\theta$ for some $\theta \in \Theta$, $X \in \mathcal{X}$. Let $F(\theta) : \Theta \mapsto \mathcal{Y}$ be a measurable function. Let $d(F, \hat{F}) : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_{\geq 0}$ be the loss function, which is also assumed to be a metric.

Given observation x , upon asserting that a statistic T is *partially sufficient* (Basu, 2011) for estimating $F(\theta)$, the general profile maximum likelihood approach aims at maximizing the probability that $t = T(x)$ appears rather than the probability of x , which is the aim of traditional maximum likelihood. Then, we simply plug-in the PML estimator into the function $F(\cdot)$ to obtain an estimate of the functional $F(\theta)$. The profile maximum likelihood estimator of θ is defined as follows.

Definition 1 (Profile maximum likelihood) *The profile maximum likelihood estimator of θ is defined as*

$$\hat{\theta}_T(t) \triangleq \underset{\theta}{\operatorname{argmax}} P_\theta(T(X) = t). \quad (4)$$

Remark 2 *The notion of profile likelihood has appeared in the statistics literature before (Murphy and Van der Vaart, 2000), bearing a different meaning from the one we adopt. The principle of profile maximum likelihood we introduced in Definition 1 is known in the statistics community as restricted maximum likelihood, which was introduced in (Patterson and Thompson, 1971) and surveyed in (Harville, 1977).*

The following theorem provides a general performance guarantee for the PML algorithm introduced in Definition 1 in terms of estimating any functional $F(\theta)$.

Theorem 3 (Acharya et al., 2017) *We fix a statistical model $\mathcal{E} = \{\mathcal{X}, P_\theta : \theta \in \Theta\}$. Let $T = T(X)$ be a statistic such that $T : \mathcal{X} \mapsto \mathcal{T}$ and $|\mathcal{T}| < \infty$. Let $F(\theta) : \Theta \mapsto \mathcal{Y}$ be a measurable function. Suppose there exists an estimator $\hat{F} : \mathcal{T} \mapsto \mathcal{Y}$ such that*

$$\sup_{\theta \in \Theta} P_\theta \left(d(F(\theta), \hat{F}(T)) > \epsilon \right) < \delta. \quad (5)$$

Then,

$$\sup_{\theta \in \Theta} P_\theta \left(d(F(\theta), F(\hat{\theta}_T)) > 2\epsilon \right) \leq \delta \cdot |\mathcal{T}|. \quad (6)$$

The usefulness of Theorem 3 would rely on two factors: the existence of a good estimator $\hat{F}(T)$ in terms of *worst case risk*, and the small cardinality of the set \mathcal{T} that the statistic T lies in. As we discussed above, it would be sensible to employ the PML approach if we can assert that the statistic T is “sufficient” for estimating $F(\theta)$. However, what is the right definition of sufficiency in this context?

The definition of sufficiency in this context is one of the most basic problems in decision theory, which turns out to be a highly non-trivial question. It is usually termed “partial sufficiency”, and we refer the readers to (Basu, 2011) for a survey. To illustrate the subtleties of this definition, we review some historical developments below. This question was considered by Kolmogorov (Kolmogoroff, 1942), who proposed the following definition:

Definition 4 (Kolmogorov’s definition of partial sufficiency) *(Kolmogoroff, 1942)* A statistic $T = T(X)$ is called *partially sufficient* for $F(\theta)$, if the posterior distribution of $F(\theta)$ given $X = x$, depends only on $T = t$ and on the prior distribution of θ .

It was later shown by Hájek (Hájek, 1967) that the Kolmogorov definition is void in the sense that if $F(\theta)$ is not a constant, and T is partially sufficient for $F(\theta)$ in the sense of Kolmogorov, then T is sufficient for θ in the usual sense as well.

In this context, Hájek proposed another definition of partial sufficiency.²

Definition 5 (Hájek, 1967, Def. 2.2) A statistic $T = T(X)$ is called *partially sufficient* for $F(\theta)$ if there exists some other functional $R(\theta)$ such that $F(\theta) = F_1(R(\theta))$, and the following is satisfied:

1. The distribution of T will depend on $R(\theta)$ only, that is

$$P_\theta(dT) = P_{R(\theta)}(dT); \tag{7}$$

2. There exists a distribution $Q_R \in \mathcal{P}_R$ such that T is sufficient for the family $\{Q_R\}$, where \mathcal{P}_R is the convex hull of the distributions $\{P_\theta: R(\theta) = R\}$.

It was shown in (Hájek, 1967) that under the partial sufficiency in Def. 5, an analogue of the Rao-Blackwell theorem in terms of the worst case risk can be proved.³

It begs the question: in general how could one prove a certain statistic T is partially sufficient for $F(\theta)$? The following equivalence relation induced by group transformations is common in practice. Let $G = \{g\}$ be a group of one-to-one transformations of the \mathcal{X} space on itself. We shall say that an event is G -invariant, if $gA = A$ for all $g \in G$. The set of G -invariant events is a sub- σ -algebra \mathcal{B} , and a measurable function f is \mathcal{B} -measurable if and only if $f(gx) = f(x)$ for all $g \in G, x \in \mathcal{X}$.

Theorem 6 (Hájek, 1967) Suppose parameter θ can be represented as $\theta = (\tau, g)$, where $g \in G$, and G is a group of one-to-one transformations of the \mathcal{X} space on itself. Suppose there exists a family of probability distributions $\{P_\tau\}$ indexed by τ such that

$$P_\theta(A) = P_\tau(gX \in A), \tag{8}$$

2. For other definitions of sufficiency, see (Le, 1964; Birnbaum, 1962; Blackwell and Girshick, 1979).

3. Note that it is weaker than the usual Rao-Blackwell theorem, which holds pointwise for the risk function.

where $X \sim P_\tau$. Then, under mild conditions⁴, the sub- σ -algebra of G -invariant events is partially sufficient for τ under statistical model $\{\mathcal{X}, P_\theta\}$ in the sense of Def. 5.

Theorem 6 can be interpreted in the following way. We introduce the equivalence relation $x_1 \sim x_2$ if and only if there exists some $g \in G$ such that $x_1 = gx_2$. Then, Theorem 6 states that it is partially sufficient for τ to look at the equivalence classes $[a] = \{x \in \mathcal{X} | x \sim a\}$.

3. Discrete distributions up to relabeling

In this section, we specialize Theorem 6 to several different settings, centered around the problem of estimating functionals of discrete distributions.

3.1. Permutation group and single sorted probability vector

The standard PML problem can be viewed as a special case of Theorem 6 with the group being the permutation group. Concretely, let G be the permutation group $\mathcal{S}_\mathcal{X}$ on \mathcal{X} , let p be a distribution supported on set \mathcal{X} – that is, $p_x > 0 \forall x \in \mathcal{X}$ – and let τ be the sorted non-increasing probability vector $(p_{(1)}, p_{(2)}, \dots, p_{(|\mathcal{X}|)})$ of p . It is clear that the label-invariant properties of a distribution, such as the entropy, Rényi entropy, and support set size, depend on p only through τ .

Suppose we observe n i.i.d. samples $x_1^n = (x_1, \dots, x_n)$ with distribution p . Denote by $\hat{p} = \hat{p}(x_1^n)$ the empirical distribution:

$$\hat{p} = (\hat{p}_x)_{x \in \mathcal{X}} = \left(\frac{1}{n} \sum_{i=1}^n \mathbb{1}(x_i = x) \right)_{x \in \mathcal{X}} \quad (9)$$

It is well known that the empirical distribution is the minimum complete sufficient statistic for p (Lehmann and Casella, 1998). The probability of observing a specific empirical distribution is given by

$$\mathbb{P}_p(\hat{p}) = \binom{n}{n\hat{p}} \prod_{x \in \mathcal{X}} p_x^{n\hat{p}_x} \quad (10)$$

where the prefactor in (10) is a multinomial coefficient⁵.

Applying Theorem 6 with $\theta = p = (\tau, g)$, we obtain that the *fingerprint* statistic (Valiant and Valiant, 2011b, 2013)⁶ is partially sufficient for τ . Concretely, the fingerprint $\mathcal{F} = (\mathcal{F}_i)_{i \geq 0}$ is defined so that \mathcal{F}_i is the number of symbols observed exactly i times in x_1^n :

$$\mathcal{F} = \mathcal{F}(\hat{p}(x_1^n)) = (\mathcal{F}_i)_{i \geq 0} \triangleq (|\{x \in \mathcal{X} : n\hat{p}_x = i\}|)_{i \geq 0} \quad (11)$$

4. Precisely, Condition 3.1 in (Hájek, 1967).

5. $\binom{n}{n\hat{p}} \triangleq n! \prod_{x \in \mathcal{X}} \frac{1}{(n\hat{p}_x)!}$.

6. The fingerprint is also called the profile (Orlitsky et al., 2004b), histogram order statistics (Paninski, 2003), and histogram of histograms (Batu et al., 2000).

Below, we compute the probability of a specific fingerprint for the case that p is supported on finite alphabet $\mathcal{X} = \text{Supp}(p)$ with

$$K \triangleq |\mathcal{X}| \tag{12}$$

and the empirical distribution \hat{p} is supported on empirical alphabet $\hat{\mathcal{X}} \triangleq \text{Supp}(\hat{p}) = \{x \in \mathcal{X} : \hat{p}_x > 0\} \subset \mathcal{X}$ with

$$\hat{K} \triangleq |\hat{\mathcal{X}}|. \tag{13}$$

Then $\mathcal{F}_0 = |\mathcal{X} \setminus \hat{\mathcal{X}}| = K - \hat{K}$ counts the number of “unseen” symbols, and is thus unknown if the support set size K is unknown.

In Appendix A we derive the following expression for the probability of a specific fingerprint for the case of finite \mathcal{X} :

$$\mathbb{P}_p(\mathcal{F}) = \left(\prod_{i \geq 0} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \text{perm} \left(\underbrace{\left(\binom{n\hat{p}_{x'}}{p_x} \right)_{x,x' \in \mathcal{X}}}_Q \right) \tag{14}$$

where $\text{perm}(A)$ denotes the matrix permanent of the $K \times K$ matrix A :

$$\text{perm}(A) = \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \prod_{x \in \mathcal{X}} A_{x,\sigma(x)} \tag{15}$$

where the sum is over the symmetric group $\mathcal{S}_{\mathcal{X}}$ on \mathcal{X} . To simplify notation, we let Q denote the $K \times K$ matrix in (14). Note that to evaluate $\text{perm}(Q)$ in (14) we need to know the support set \mathcal{X} , both to evaluate $\text{perm}(Q)$ and $\mathcal{F}_0 = |\mathcal{X} \setminus \hat{\mathcal{X}}|$. Appendix B shows examples of the computation of $\mathbb{P}_p(\mathcal{F})$ over small alphabets.

For completeness, Appendix A also contains expression (87) for $\mathbb{P}_p(\mathcal{F})$ for the case of infinite countable \mathcal{X} , though this expression is not used elsewhere in this work for the reason stated in the last paragraph of Section 3.1.

For a given collection of distributions \mathcal{P} , the *profile maximum likelihood* distribution is defined as

$$p^* \triangleq \underset{p \in \mathcal{P}}{\text{argmax}} \mathbb{P}_p(\mathcal{F}) = \underset{p \in \mathcal{P}}{\text{argmax}} \frac{\text{perm}(Q)}{(K - \hat{K})!}, \tag{16}$$

where in the second equality we discarded all p -independent factors of $\mathbb{P}_p(\mathcal{F})$ (14). $\mathcal{F}_0 = K - \hat{K} = |\mathcal{X} \setminus \hat{\mathcal{X}}|$ depends on p through its support set size. Note that $\mathbb{P}_p(\mathcal{F})$ is invariant under relabeling of the components of p , so we can choose p^* to be non-increasing in the same ordering as we choose for the support set \mathcal{X} . Note that the set \mathcal{P} is not necessarily the same as the set of all discrete distributions.

If the collection of distributions \mathcal{P} includes distributions with different support set sizes (for example, all finite support set sizes at least as large as \hat{K}), then we can estimate the support set size by breaking up the optimization in (16) into two steps:

$$K^* \triangleq \underset{K}{\text{argmax}} \left(\frac{1}{(K - \hat{K})!} \max_{p \in \mathcal{P}_K} \text{perm}(Q) \right) \tag{17}$$

whenever the max over K exists, where $\mathcal{P}_K \triangleq \{p \in \mathcal{P} : |\text{Supp}(p)| = K\}$. The maximizer K^* usually exists because increasing K makes the first factor in (17) smaller, but makes the $K \times K$ matrix Q bigger, boosting the number $K!$ of permutations to sum over in computing the permanent.

It may happen that K^* does not exist⁷, in which case we are still able to define a PML distribution in terms of a discrete part and continuous part in the terminology of (Orlitsky et al., 2004b). (Orlitsky et al., 2004b) showed that for any fingerprint \mathcal{F} , there exists a distribution p^* maximizing $\mathbb{P}_p(\mathcal{F})$, but this distribution may assign some symbols to the continuous part, and others to the discrete part with finite support. The intuition is that if there are sufficiently many symbols that appear exactly once in the sample, then the PML distribution p^* assigns discrete probability 0 to infinitely many symbols to maximize the probability that each of them is seen only once. Then we can define K_d^* to be the optimal support set size of the discrete part of p^* . (Orlitsky et al., 2004b) further derived conditions lower- and upper-bounding K_d^* in terms of $\max_x \hat{p}_x$ and $\min_x \hat{p}_x$ and thus showing K_d^* exists, and upper-bounding the size of the continuous part in terms of \mathcal{F}_1 (the numbers of symbols occurring once). (Acharya et al., 2009) computed p^* for all sample sizes up to $n = 7$.

Note that either K^* exists or the PML distribution p^* has a continuous part, but it is never the case that p^* has infinite countable support. For this reason we never use expression (87) for the probability of a fingerprint when \mathcal{X} is countably infinite.

3.2. Permutation group and multiple probability vectors

Suppose we have two distributions p, q on the same alphabet \mathcal{X} , such that $p_x + q_x > 0 \forall x \in \mathcal{X}$. This condition ensures that there are no symbols x such that $p_x = q_x = 0$, which simplifies the expressions below. Draw n samples i.i.d. from p with empirical distribution \hat{p} and draw m samples i.i.d. from q with empirical distribution \hat{q} . It is clear that the label-invariant properties of two distributions, such as the Kullback-Leibler divergence, L_1 distance, and the general family of divergences $\sum_{x \in \mathcal{X}} f(p_x, q_x)$, are invariant to the permutation group $\mathcal{S}_{\mathcal{X}}$ acting on pairs of distributions as $\sigma(p, q) \triangleq (\sigma p, \sigma q)$ for all $\sigma \in \mathcal{S}_{\mathcal{X}}$. Denote the 2-D fingerprint (Raghunathan et al., 2017) by $\mathcal{F}(\hat{p}, \hat{q})$, analogous to (11):

$$\mathcal{F} = \mathcal{F}(\hat{p}, \hat{q}) = (\mathcal{F}_{i,j})_{i,j \geq 0} \triangleq (|\{x \in \mathcal{X} : n\hat{p}_x = i, m\hat{q}_x = j\}|)_{i,j \geq 0} \quad (18)$$

The probability to draw the 2-D fingerprint under p, q is:

$$\mathbb{P}_{p,q}(\mathcal{F}) = \left(\prod_{i,j \geq 0} \frac{1}{\mathcal{F}_{i,j}!} \right) \binom{n}{n\hat{p}} \binom{m}{m\hat{q}} \text{perm} \left(\left(\begin{matrix} n\hat{p}_{x'} & m\hat{q}_{x'} \\ p_x & q_x \end{matrix} \right)_{x,x' \in \mathcal{X}} \right) \quad (19)$$

Following the general profile maximum likelihood methodology, for a collection \mathcal{P} of pairs of distributions on alphabet \mathcal{X} , the PML distributions p^*, q^* are:

$$(p^*, q^*) \triangleq \underset{(p,q) \in \mathcal{P}}{\text{argmax}} \mathbb{P}_{p,q}(\mathcal{F}) \quad (20)$$

Note that finding the PML pair (p^*, q^*) is not equivalent to solving two independent optimization problems since the matrix permanent in (19) does not factor into two terms. As

7. For example, if $n \geq 2$ and each symbol occurs exactly once in the sample, so $\hat{p}_x = \frac{1}{n}$ for all $x \in \mathcal{X}(\hat{p})$.

in the 1D PML case, if the support set size $K = |\mathcal{X}|$ is unknown, then we can attempt to estimate it analogously to (17). See Appendix H for a precise statement.

Appendix H generalizes the argument above in a straightforward manner to D -dimensional fingerprints (Raghunathan et al., 2017). This allows estimation of functionals of more than two distributions.

4. An approximation to the profile maximum likelihood distribution

4.1. Our approximate PML algorithm: an overview

It seems computationally challenging to compute the profile maximum likelihood distribution in all the cases discussed in Section 3, though the computational complexity of this problem is not known. Indeed, it involves computing the matrix permanent in (16), and computing permanents is $\#P$ -hard (Valiant, 1979). Since the computation of permanents is exponentially slow in the support set size $|\mathcal{X}|$ with the best-known algorithms, one may lower the target and hope to find an approximate efficient algorithm to all the profile maximum likelihood problems in Section 3. Some algorithms on approximate PML have been developed: (Orlitsky et al., 2004a) proposes an expectation maximization algorithm and (Vontobel, 2012) proposes a Bethe approximation to the PML distribution. None of these works, including ours, provides guarantees on the nearness of the approximate and exact PML distributions.

Here we develop a PML distribution approximation inspired by some empirical observations in the small examples where this distribution is computable exactly in reasonable time. We provide a fast algorithm for computing this approximation. The algorithm has no tunable parameters and runs in time at most linear in the sample size in the standard PML problem (Section 3.1), its run time usually dominated by the computation of the empirical distribution \hat{p} . It is also efficient in solving divergence estimation problems as described in Section 3.2.

Observations from solving the PML exactly in small cases (see Section 4.2) yield the intuition that the PML distribution p^* assigns equal mass to (“clumps together”) symbols whose empirical counts are close, differing on the order of \sqrt{n} , where n is the sample size. For such a “clumped” distribution p^* , we further guess that the value of the matrix permanent in $\mathbb{P}_{p^*}(\mathcal{F}) \sim \text{perm}(Q)/\mathcal{F}_0!$ (14) is dominated by a subset of terms in the summation over the symmetric group $\mathcal{S}_{\mathcal{X}}$ – namely the subgroup of permutations that only mix symbols within the same “clump” – a level set of p^* – but not between different clumps. Armed with this intuition, we replace maximization of the matrix permanent, which seems hard, with maximization of a lower bound involving a sum over only this subgroup of permutations, which is much easier. Section 4.3 gives a precise statement of these remarks and further motivates our approach. The remainder of this section is devoted to an overview of our approximation algorithm.

Let $\mathcal{A}(p) = \{\alpha\}$ be the partition of \mathcal{X} into level sets of distribution p :

$$\mathcal{A}(p) \triangleq \{\{x \in \mathcal{X} : p_x = u\} : u \in \{p_x : x \in \mathcal{X}\}\} \quad (21)$$

That is, $x, x' \in \alpha \in \mathcal{A}(p)$ if and only if $p_x = p_{x'}$.

Then our lower bound \bar{V} to the log permanent is defined by summing over a subgroup of the symmetric group, and satisfies (see derivation in Section 4.3):

$$\log(\text{perm}(Q)) \geq \bar{V}(p) = \sum_{\alpha \in \mathcal{A}(p)} \log(|\alpha|!) - n(D(\hat{p}||p) + H(\hat{p})) \quad (22)$$

where $D(\hat{p}||p) = \sum_{x \in \mathcal{X}} \hat{p}_x \log(\hat{p}_x/p_x)$ is the Kullback-Leibler divergence and $H(\hat{p}) = -\sum_{x \in \mathcal{X}} p_x \log(p_x)$ is the entropy of the empirical distribution. The approximate PML (APML) distribution is then the maximizer of the lower bound \bar{V} :

$$\bar{p}^* \triangleq \operatorname{argmax}_{p \in \mathcal{P}} \bar{V}(p) \quad (23)$$

The right hand side of (22) yields some intuition for the approximate PML distribution. Suppose \mathcal{P} contains only distributions with support set size $K = |\mathcal{X}|$, so $\mathcal{F}_0 = K - |\hat{\mathcal{X}}|$ is fixed in (16), and the first term on the right hand side of (22) encourages \bar{p}^* to clump many symbols together into a few large clumps (since $\log(|\alpha|!) \sim |\alpha| \log(|\alpha|)$ is superlinear in $|\alpha|$), while the second term encourages \bar{p}^* to be similar to the empirical distribution \hat{p} . As the sample size $n \rightarrow \infty$, the second term dominates, and we have $\bar{p}^* \rightarrow \hat{p}$, consistent with our intuition for the large sample size limit and consistent with the result of (Orlitsky et al., 2004b) (Theorem 16) for the exact PML distribution.

It turns out that computation of the approximate PML distribution \bar{p}^* is equivalent to optimization over all partitions of \mathcal{X} . This is because, as we show, any distribution \bar{p}^* maximizing \bar{V} satisfies an averaging property with respect to the empirical distribution \hat{p} for all $x \in \mathcal{X}$: Let $\alpha(x, \bar{p}^*) = \{x' \in \mathcal{X} : \bar{p}_{x'}^* = \bar{p}_x^*\} \in \mathcal{A}(\bar{p}^*)$ be the level set of \bar{p}^* containing symbol x , then

$$\bar{p}_x^* = \frac{1}{|\alpha(x, \bar{p}^*)|} \sum_{x' \in \alpha(x, \bar{p}^*)} \hat{p}_{x'} \quad (24)$$

Therefore \bar{p}^* is determined by its partition of the alphabet \mathcal{X} into level sets, so we can replace maximization of \bar{V} over p with maximization over all partitions \mathcal{A} . Let $\bar{\mathcal{A}}^*$ denote the optimal partition:

$$\bar{\mathcal{A}}^* \triangleq \operatorname{argmax}_{\mathcal{A}: \text{partition of } \mathcal{X}} \bar{V}(\mathcal{A}) \quad (25)$$

where $\bar{V}(\mathcal{A}) \triangleq \bar{V}(p_{\mathcal{A}}(\hat{p}))$ and $p_{\mathcal{A}}(\hat{p})$ is the distribution obtained by averaging the empirical distribution over the partition elements – that is, satisfying property (24) for all $x \in \mathcal{X}$. The approximate PML distribution is then $\bar{p}^* = p_{\bar{\mathcal{A}}^*}(\hat{p})$.

We derive constraints on the optimal partition $\bar{\mathcal{A}}^*$ that enable an efficient⁸ dynamic programming algorithm to compute $\bar{\mathcal{A}}^*$ and the approximate PML distribution \bar{p}^* . This algorithm is presented in Section 4.5.

For the case of $D \geq 2$ probability vectors and D -dimensional fingerprints – for example, when estimating the KL divergence, corresponding to $D = 2$ as in Section 3.2 – we are unable to give an efficient algorithm for maximizing our lower bound \bar{V} suitably generalized to the D -dimensional case. The difficulty is that there is no natural ordering on \mathbb{N}^D for

8. Running in time $O(|\text{Supp}(\mathcal{F})|^2)$. Any empirical distribution \hat{p} satisfies $|\text{Supp}(\mathcal{F}(\hat{p}))| \leq (\sqrt{8n+1}+1)/2$ (see discussion at end of Section 4.5), where n is the sample size, so the running time is $O(n)$. The run time of our algorithm is usually dominated by computation of the empirical distribution \hat{p} .

$D \geq 2$, so the dynamic programming approach for $D = 1$ does not work here. In this case we settle for another layer of approximation, using a greedy heuristic to iteratively merge clumps of symbols: we enlarge the first term in (22) until the second term becomes too large. This algorithm is presented in Appendix H.

Figures 1 and 2 show the approximate PML distributions for the cases $D = 1$ and $D = 2$, respectively. For $D = 1$, the approximate PML distribution is as defined in (23). For $D \geq 2$, we approximately maximize our lower bound \bar{V} to the log matrix permanent as described in Appendix H. Figure 14 in Appendix H shows the case $D = 3$.

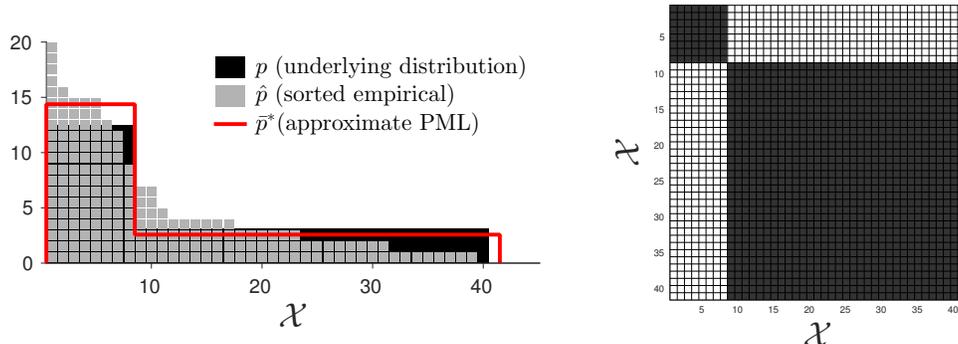


Figure 1: (left) An empirical histogram \hat{p} sorted in non-increasing order (gray squares) drawn from underlying distribution p (black area) and the approximate PML (APML) distribution \bar{p}^* (23) (red line), scaled by the sample size $n = 200$. The alphabet is $\mathcal{X} = \{1, \dots, 40\}$, assumed unknown in computing \bar{p}^* . (right) Computation of the log permanent lower bound $\bar{V}(\bar{p}^*)$ (31) involves summing over all permutations that mix symbols only within the level sets of \bar{p}^* , corresponding to all permutation matrices with nonzero entries within the black blocks.

4.2. Intuitions from solving the PML exactly

For small alphabets, we can numerically and sometimes analytically optimize the permanent in (16) and (20). We observe that the PML distribution tends to assign equal mass to symbols whose empirical counts are close, differing on the order of \sqrt{n} , where n is the sample size. This observation motivates our approximate PML scheme in Section 4.3.

4.2.1. EXACT SOLUTION FOR SIZE 2 ALPHABET

Suppose the alphabet size is $|\mathcal{X}| = 2$. Given sample size n and empirical distribution $\hat{p} = (\hat{p}_1, \hat{p}_2) = (\hat{p}_1, 1 - \hat{p}_1)$, without loss of generality one may assume $\hat{p}_1 \geq \frac{1}{n}$ (otherwise the PML is given by $\text{Bern}(1)$) and $\hat{p}_1 \leq \frac{1}{2}$. The exact PML distribution \hat{p}^* was computed in (Orlitsky et al., 2004b):

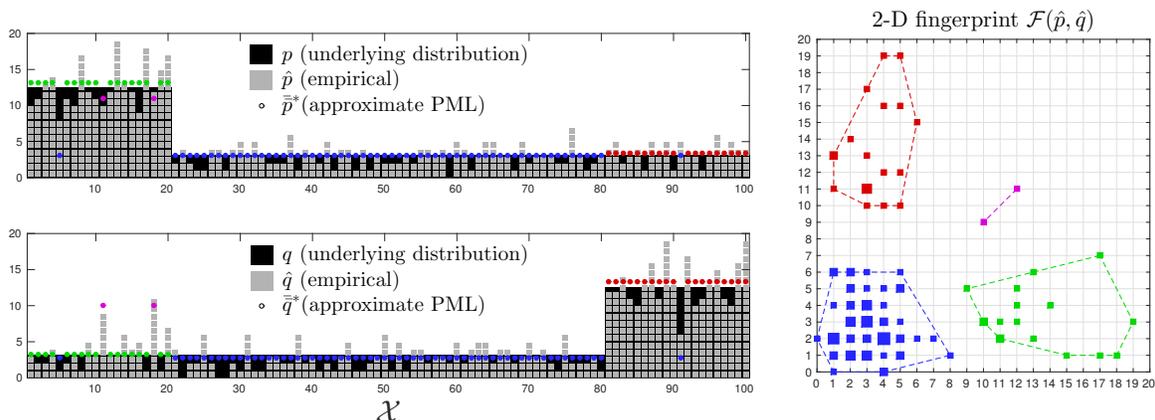


Figure 2: (Left) empirical distributions \hat{p} and \hat{q} (gray squares) drawn from underlying distributions p and q (black area) with $n = m = 400$ samples. The approximate PML distributions (\bar{p}^* , \bar{q}^*) are computed by clumping entries of a 2-D fingerprint according to the greedy heuristic described in Appendix H. The approximate PML distributions are ordered in the same way as the empirical distributions (with a distinct color for each level set corresponding to the subplot on the right). All distributions are plotted scaled by the sample size. (Right) the 2-D fingerprint $\mathcal{F}_{i,j}$ at coordinates (i, j) . Marker areas are roughly proportional to $\mathcal{F}_{i,j}$. All points within the colored convex hulls correspond to a level set of the approximate PML distributions. We can see three large level sets corresponding to the three distinct values of (p_x, q_x) for $x \in \mathcal{X}$.

Theorem 7 ((Orlitsky et al., 2004b)) For all $\frac{1}{n} \leq \hat{p}_1 \leq \frac{1}{2}$,

$$\hat{p}^* = \begin{cases} \left(\frac{1}{2}, \frac{1}{2}\right) & : |\hat{p}_1 - \hat{p}_2| \leq \frac{1}{\sqrt{n}} \\ \left(\frac{1}{1+p}, \frac{p}{1+p}\right) & : |\hat{p}_1 - \hat{p}_2| > \frac{1}{\sqrt{n}} \end{cases}, \quad (26)$$

where p is the unique root in $(0, 1)$ of the polynomial

$$\hat{p}_1 p^{n(1-2\hat{p}_1)+1} - \hat{p}_2 p^{n(1-2\hat{p}_1)} + \hat{p}_2 p - \hat{p}_1. \quad (27)$$

Theorem 7 states that $|\hat{p}_1 - \hat{p}_2| \leq \frac{1}{\sqrt{n}} \Leftrightarrow p^*$ is uniform, which confirms our intuition that the relative ranking of two bins is “resolvable” if their empirical counts are more than about \sqrt{n} apart, since the empirical counts are marginally binomially-distributed with standard deviation proportional to \sqrt{n} . Figure 3 summarizes these observations for the size-2 alphabet.

4.2.2. SIZE 3 ALPHABET

Next let’s consider the case of $|\mathcal{X}| = 3$ bins. Given empirical distribution $\hat{p} = (\hat{p}_1, \hat{p}_2, \hat{p}_3)$, we use an expectation-maximization (EM) algorithm to numerically search for the PML distribution. The EM algorithm starts from an initial point $p \leftarrow \hat{p}$ and converges to a local maximum of $\text{perm}(Q)$ (14). We sort the components of this local maximum in non-

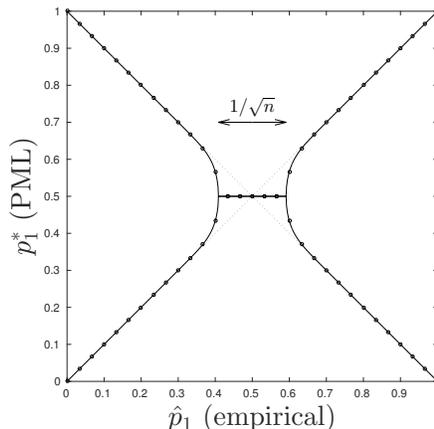


Figure 3: “Phase diagram” for the PML distribution $p^* = (p_1^*, 1 - p_1^*)$ on binary alphabet $\mathcal{X} = \{1, 2\}$ and empirical distribution $(\hat{p}_1, 1 - \hat{p}_1)$ on $n = 30$ samples (plotting the first component). Circles correspond to the possible histograms ($\hat{p}_1 \in \{0/30, 1/30, \dots, 30/30\}$). Solid black lines correspond to the solution for arbitrary $\hat{p}_1 \in [0, 1]$, obtained numerically. The PML distribution p^* is uniform when $|\hat{p}_1 - \hat{p}_2| \leq 1/\sqrt{n}$. Outside this middle region there are two PML branches, since any permutation of the PML distribution is another PML distribution. Diagonal dashed lines correspond to the lines (\hat{p}_1, \hat{p}_1) and $(\hat{p}_1, 1 - \hat{p}_1)$.

increasing order and denote the result by p_{EM}^* . Details of the EM algorithm can be found in Appendix D. We summarize our findings in Figure 4.

4.2.3. EXACT SOLUTION FOR SIZE 2 ALPHABET WITH D DISTRIBUTIONS

Suppose we have D distributions $((p_1^{(d)}, 1 - p_1^{(d)}))_{d=1}^D$ on the same alphabet $\mathcal{X} = \{1, 2\}$ and draw $(n_d)_{d=1}^D$ samples from each, obtaining empirical distributions $((\hat{p}_1^{(d)}, 1 - \hat{p}_1^{(d)}))_{d=1}^D$. Then the D -dimensional PML distribution $(p^{*(d)})_{d=1}^D$ (defined for $D = 2$ in Section 3.2 and for general D in Appendix H) is stated in Theorem 8.

Theorem 8 *The D -tuple of PML distributions $(p^{*(d)})_{d=1}^D$ on a binary alphabet satisfies:*

$$\sum_{d=1}^D 4n_d \left(\hat{p}_1^{(d)} - \frac{1}{2} \right)^2 > 1 \quad \Rightarrow \quad (p^{*(d)})_{d=1}^D \neq \left(\left(\frac{1}{2}, \frac{1}{2} \right) \right)_{d=1}^D \quad (28)$$

See Appendix E for a proof. We conjecture that the converse of statement (28) holds.

Theorem 8 extends the results of (Orlitsky et al., 2004b) and (Fernandes and Kashyap, 2013) on the uniformity of the D -dimensional distribution beyond $D = 1$. The left side of (28) describes the interior of an ellipsoid in the space of distributions centered on $(\frac{1}{2})_{d=1}^D$, shown in Figure 5 for the case $D = 2$. Note that (28) is not consistent with D copies of condition (26) for $D \geq 2$; this reflects the fact that the probability of a D -dimensional fingerprint does not decompose into a product of D terms (see (19) for the case $D = 2$ and Appendix H for the general case).

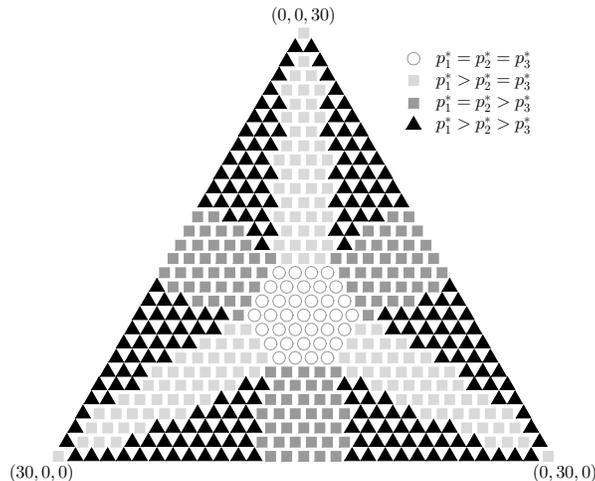


Figure 4: “Phase diagram” for the PML distribution approximation p_{EM}^* for alphabet size $|\mathcal{X}| = 3$ bins and $n = 30$ samples (analogous to Figure 3), computed by an expectation-maximization (EM) algorithm as discussed in Section 4.2.2 and Appendix D. We plot a marker for each possible empirical distribution – a type – on the simplex. Each type has three coordinates, which we then project onto the 2-dimensional image shown (so, e.g., the uniform type $n\hat{p} = (10, 10, 10)$ is in the center). The shape of the marker at position $n\hat{p}$ corresponds to the level set decomposition of the distribution p_{EM}^* , with components sorted in non-increasing order.

4.2.4. EMPIRICAL OBSERVATIONS OF THE EXACT SOLUTIONS

Figures 3, 4, and 5 suggest that the PML distribution tends to “cluster” together similar entries of the empirical histogram \hat{p} , rather than smooth them out. If $n|\hat{p}_x - \hat{p}_{x'}|$ is smaller than approximately $\sqrt{n}\hat{p}_x$, then the relative ranking of the true distribution’s x and x' components is not well resolved statistically, and the permanent of Q (14) tends to be maximized with $p_x^* = p_{x'}^*$. On the other hand if $n|\hat{p}_x - \hat{p}_{x'}|$ is larger than approximately $\sqrt{n}\hat{p}_x$, then it tends to happen that $p_x^* \neq p_{x'}^*$. The similar phenomenon of a phase transition between a uniform and nonuniform PML distribution (for a different approximation of the PML (Vontobel, 2012)) has been reported by (Fernandes and Kashyap, 2013; Chan et al., 2015).

These observations lead to the intuition that permutations $\sigma \in \mathcal{S}_{\mathcal{X}}$ that exchange sufficiently different components of \hat{p} do not contribute much to the permanent $\text{perm}(Q)$. This intuition motivates our idea of simplifying the computation of the permanent by summing over only those permutations that contribute “a lot” to the sum. Section 4.3 develops this idea.

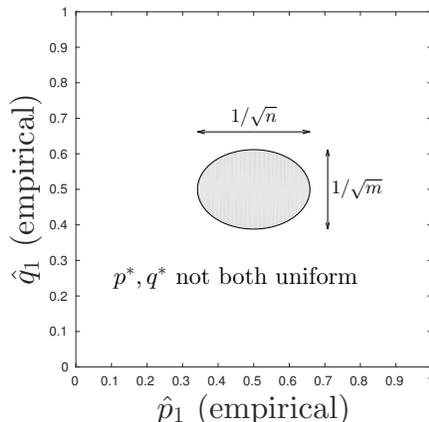


Figure 5: “Phase diagram” for the 2-D PML pair of distributions p^*, q^* (20) on binary alphabet $\mathcal{X} = \{1, 2\}$ and empirical distributions $(\hat{p}_1, 1 - \hat{p}_1), (\hat{q}_1, 1 - \hat{q}_1)$ on $n = 10, m = 20$ samples. p^*, q^* are not both uniform whenever the empirical distribution components \hat{p}_1, \hat{q}_1 both lie outside the shaded ellipse $4n \left(\hat{p}_1 - \frac{1}{2}\right)^2 + 4m \left(\hat{q}_1 - \frac{1}{2}\right)^2 \leq 1$.

4.3. Approximate PML: the case of a single distribution

In the sequel, denote by $\mathcal{A}(p) = \{\alpha\}$ the partition of \mathcal{X} into the level sets of p :

$$\mathcal{A}(p) \triangleq \{\{x \in \mathcal{X} : p_x = u\} : u \in \{p_x : x \in \mathcal{X}\}\} \quad (29)$$

In this Section we assume that the underlying distribution p is supported on finite set \mathcal{X} of known cardinality. The case of unknown or infinite support set size is treated in Section 4.6. If the support set size is known, then so is the number of “unseen” symbols \mathcal{F}_0 , so from (16) we conclude that the PML distribution p^* is a maximizer of the function:

$$V(p) \triangleq \log(\text{perm}(Q)) \quad (30)$$

We lower bound the log permanent $V(p)$ by summing over a subset of the terms in the summation – a subgroup of the symmetric group $\mathcal{S}_{\mathcal{X}}$. Denote the lower bound by $\bar{V}(p)$:

$$\bar{V}(p) \triangleq \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X}, p}} \prod_{x \in \mathcal{X}} p_x^{n \hat{p}_{\sigma(x)}} \right) \leq \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \prod_{x \in \mathcal{X}} p_x^{n \hat{p}_{\sigma(x)}} \right) = V(p) \quad (31)$$

where $\mathcal{S}_{\mathcal{X}, p}$ is a subgroup of $\mathcal{S}_{\mathcal{X}}$ consisting of all permutations that exchange only those alphabet symbols $x, x' \in \mathcal{X}$ that are within the same level set of p . Equivalently, $\mathcal{S}_{\mathcal{X}, p}$ is the stabilizer of distribution p under the action of relabeling its components:

$$\mathcal{S}_{\mathcal{X}, p} \triangleq \{\sigma \in \mathcal{S}_{\mathcal{X}} : \sigma p = p\} \cong \prod_{\alpha \in \mathcal{A}(p)} \mathcal{S}_{\alpha} \quad (32)$$

where $(\sigma p)_x \triangleq p_{\sigma(x)}$, \cong denotes group isomorphism, \times denotes the direct product of groups, and \mathcal{S}_{α} is the symmetric group acting on level set $\alpha \subset \mathcal{X}$. Then $|\mathcal{S}_{\mathcal{X}, p}| = \prod_{\alpha \in \mathcal{A}(p)} |\alpha|!$.

Our approximate profile maximum likelihood distribution \bar{p}^* maximizes the lower bound $\bar{V}(p)$. In other words, for the case of every distribution in collection \mathcal{P} having the same support set size,

$$\bar{p}^* \triangleq \operatorname{argmax}_{p \in \mathcal{P}} \bar{V}(p) \quad (33)$$

The case of unknown support set size is treated in Section 4.6.

Some intuition for the computational difficulty of the proposed approximation: the product structure (32) of $\mathcal{S}_{\mathcal{X},p}$ allows us to lower bound the log matrix permanent $V(p) = \log(\operatorname{perm}(Q)) \geq \log(\operatorname{perm}(\bar{Q})) = \bar{V}(p)$, where \bar{Q} is the block-diagonal matrix $\bar{Q}_{x,x'} = Q_{x,x'} \mathbb{1}(\hat{p}_x = \hat{p}_{x'})$ whose permanent is the product of the permanents in each block. Since p is constant within each block, the permanent of each block is easy to evaluate, so evaluation of the lower bound $\bar{V}(p)$ is dramatically simpler than evaluation of $V(p)$, taking time $O(|\mathcal{X}|)$. It moreover turns out to be possible to optimize $\bar{V}(p)$ computationally efficiently to find the approximate PML distribution \bar{p}^* ; this is done below.

The restriction to summing over $\mathcal{S}_{\mathcal{X},p}$ rather than $\mathcal{S}_{\mathcal{X}}$ seems large, but the hope is that \bar{p}^* clusters together only comparably frequent symbols (that is, $\bar{p}_x^* = \bar{p}_{x'}^*$ implies $n|\hat{p}_x - \hat{p}_{x'}|$ is less than about $\sqrt{n\hat{p}_x}$), so we hope that $\bar{V}(\bar{p}^*) \approx V(\bar{p}^*)$ and that \bar{p}^* is close to p^* , the exact profile maximum likelihood (16). Figure 1 shows that the approximate PML distribution \bar{p}^* clustering together similar enough bins of \hat{p} . We thus expect our approximate PML scheme to perform best when data is drawn from a distribution with a few well-separated level sets (like a uniform distribution or mixture of uniforms) and less well for more smoothly-varying distributions (like a Zipf distribution); this intuition is borne by the numerical experiments of Section 5.

We can rewrite our log permanent lower bound $\bar{V}(p)$ (31) as (see proof in Appendix F):

$$\bar{V}(p) = \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X},p}} \prod_{x \in \mathcal{X}} p_x^{n\hat{p}_{\sigma(x)}} \right) \quad (34)$$

$$= -n(D(\hat{p}||p) + H(\hat{p})) + \sum_{\alpha \in \mathcal{A}(p)} \log(|\alpha|!) \quad (35)$$

where $\mathcal{A}(p)$ (29) is the partition of \mathcal{X} into level sets of p , $D(\hat{p}||p)$ is the Kullback-Leibler divergence and $H(\hat{p})$ is the entropy of the empirical distribution.

Expression (35) yields some intuition about the approximate PML distribution \bar{p}^* – a maximizer of $\bar{V}(p)$. A distribution p that “clumps” many symbols together (that is, has a few large level sets α) boosts the second term (the summation over $\alpha \in \mathcal{A}(p)$) in (35), but is very different from \hat{p} , thus lowering the first term. On the other hand, by setting $p = \hat{p}$, we maximize the first term, but reduce the second term. As the sample size $n \rightarrow \infty$, the contribution of the second term in (35) vanishes relative to the first term, and we have $\bar{p}^* \rightarrow \hat{p}$, the ML distribution, in this limit.

Section 4.4 establishes some properties of the approximate PML solution \bar{p}^* (33) that enable us to compute it efficiently via a dynamic programming algorithm in Section 4.5.

4.4. Properties of the approximate PML distribution

Let $\alpha \subset \mathcal{X}$ and denote by \hat{p}_α the average value of the empirical distribution \hat{p} over α :

$$\hat{p}_\alpha \triangleq \frac{1}{|\alpha|} \sum_{x \in \alpha} \hat{p}_x \quad (36)$$

We can show that the PML distribution approximation \bar{p}^* (33) satisfies for all $x \in \mathcal{X}$

$$\bar{p}_x^* = \hat{p}_{\alpha(x, \bar{p}^*)} \quad (37)$$

where $\alpha(x, \bar{p}^*) = \{x' : \bar{p}_{x'}^* = \bar{p}_x^*\}$ is the level set of \bar{p}^* containing x . That is, \bar{p}_x^* (37) is equal to the average value of the empirical histogram \hat{p} over the symbols clumped with x into the same level set of \bar{p}^* . This follows from the fact that $D(\hat{p}||p)$ is a Bregman divergence (Jiao et al., 2017, Lemma 4).

The *averaging* property (37) determines our approximate PML distribution \bar{p}^* in terms of its partition of \mathcal{X} into level sets. Therefore, instead of maximizing $\bar{V}(p)$, we maximize $\bar{V}(\mathcal{A})$, defined:

$$\bar{V}(\mathcal{A}) \triangleq \sum_{\alpha \in \mathcal{A}} \bar{V}(\alpha) \quad (38)$$

where for $\alpha \subset \mathcal{X}$

$$\bar{V}(\alpha) \triangleq \log(|\alpha|!) + n|\alpha|\hat{p}_\alpha \log(\hat{p}_\alpha) \quad (39)$$

We can check that $\bar{V}(\mathcal{A}) = \bar{V}(\hat{p}_{\mathcal{A}})$, where $\hat{p}_{\mathcal{A}}$ is the distribution obtained by averaging the empirical distribution \hat{p} within each partition element $\alpha \in \mathcal{A}$. Let $\bar{\mathcal{A}}^*$ denote the optimal partition:

$$\bar{\mathcal{A}}^* \triangleq \underset{\mathcal{A}: \text{partition of } \mathcal{X}}{\operatorname{argmax}} \bar{V}(\mathcal{A}) \quad (40)$$

Now optimizing the lower bound to the permanent $\bar{V}(p)$ (31) is equivalent to optimizing $\bar{V}(\mathcal{A})$ (38) over partitions of \mathcal{X} , since $\bar{V}(\bar{p}^*) = \bar{V}(\bar{\mathcal{A}}^*)$ and $\bar{p}^* = \hat{p}_{\bar{\mathcal{A}}^*}$.

We make two observations for the solution of the approximate PML problem (40) that allow us to restrict the set of partitions over which we optimize. Let \mathcal{A} be a partition of \mathcal{X} . We say \mathcal{A} has the *iso-clumping* or *convexity* properties defined below:

1. *Iso-clumping* property. Symbols with the same empirical probabilities are clumped together: For all $x, x' \in \mathcal{X}$:

$$\hat{p}_x = \hat{p}_{x'} \Rightarrow \alpha(x) = \alpha(x') \quad (41)$$

Where $\alpha(x) \in \mathcal{A}$ is the partition element containing x .

2. *Convexity* property. For all $x, x', x'' \in \mathcal{X}$:

$$\left(\hat{p}_x < \hat{p}_{x'} < \hat{p}_{x''} \text{ and } \alpha(x) = \alpha(x'') \right) \Rightarrow \alpha(x) = \alpha(x') = \alpha(x'') \quad (42)$$

Theorem 9 Let $\bar{\mathcal{A}}^*$ (40) be the partition of \mathcal{X} into level sets of the approximate PML distribution \bar{p}^* (33) and let \hat{p} be the empirical distribution. Then $\bar{\mathcal{A}}^*$ has the *iso-clumping* and *convexity* properties (41) and (42).

See Appendix G for a proof.

4.5. A dynamic programming computation of the approximate PML distribution

Theorem 9 lets us efficiently maximize $\bar{V}(\mathcal{A})$ (38) by restricting to only those partitions of \mathcal{X} that have the iso-clumping and convexity properties. Once we find $\bar{\mathcal{A}}^*$ (40), the averaging property (37) lets us compute the approximate PML distribution \bar{p}^* .

Let $\mathcal{F}_+ \triangleq (\mathcal{F}_i)_{i \geq 1}$. Let

$$0 = m_0 < m_1 < m_2 < \dots < m_{F_+} \leq n \quad (43)$$

where

$$\text{Supp}(\mathcal{F}_+) = \{m \geq 1 : \mathcal{F}_m > 0\} = \{m_i\}_{i=1}^{F_+} \quad (44)$$

Then any distribution whose level set partition satisfies the iso-clumping (41) and convexity (42) properties has all level sets of the form $\mathcal{X}_{i:j}$ with $0 \leq i \leq j \leq F_+$:

$$\mathcal{X}_{i:j} \triangleq \{x \in \mathcal{X} : m_i \leq n\hat{p}_x \leq m_j\}. \quad (45)$$

Thus we optimize $\bar{V}(\mathcal{A})$ (38) over partitions of \mathcal{X} into level sets of the form (45). Note that if $\mathcal{F}_0 > 0$, then the set $\mathcal{X}_{0:0} = \{x \in \mathcal{X} : \hat{p}_x = 0\}$ of unseen symbols can not appear in the optimal partition $\bar{\mathcal{A}}^*$ because if it does, then its probability mass under $p_{\bar{\mathcal{A}}^*}$ is 0, violating our assumption that p has support \mathcal{X} . If $\mathcal{F}_0 = 0$, then $\mathcal{X}_{0:0} = \emptyset$.

This optimization can be done by a dynamic programming algorithm. For integers $i \leq j$, let $[i, j] \triangleq \{k \in \mathbb{N} : i \leq k \leq j\}$. Let \bar{V}_i denote for $0 \leq i \leq F_+$:

$$\bar{V}_i \triangleq \max_{j \in [i \vee 1, F_+]} (\bar{V}(\mathcal{X}_{i:j}) + \bar{V}_{j+1}) \quad (46)$$

$$\stackrel{(a)}{=} \max_{\mathcal{A}: \text{partition of } \mathcal{X}_{i:F_+}} \bar{V}(\mathcal{A}) \quad (47)$$

with boundary condition $\bar{V}_{F_++1} \triangleq 0$, where $i \vee j$ denotes the greater of i and j , and where $\bar{V}(\mathcal{X}_{i:j})$ is as in (39). (a) follows by induction on i downwards from F_+ . We let $j \in [i \vee 1, F_+]$ in (46) (rather than $j \in [i, F_+]$) in order to optimize over only those partitions of \mathcal{X} that do not contain the set $\mathcal{X}_{0:0}$ of unseen symbols; this restriction forces the approximate PML distribution to have support set size $|\mathcal{X}|$, as we assumed in the beginning of Section 4.3.

Then the PML distribution approximation $\bar{p}^* = \operatorname{argmax}_p \bar{V}(p)$ satisfies (setting $i = 0$ and using $\mathcal{X}_{0:F_+} = \mathcal{X}$ in (47)):

$$\bar{V}(\bar{p}^*) = \max_{\mathcal{A}: \text{partition of } \mathcal{X}} \bar{V}(\mathcal{A}) = \bar{V}_0 \quad (48)$$

We can compute the term $\bar{V}(\mathcal{X}_{i:j})$ in (46), corresponding to clustering all symbols of $\mathcal{X}_{i:j}$ into a level set of p , using (39):

$$\bar{V}(\mathcal{X}_{i:j}) = \log(|\mathcal{X}_{i:j}|!) + n|\mathcal{X}_{i:j}|\hat{p}_{\mathcal{X}_{i:j}} \log(\hat{p}_{\mathcal{X}_{i:j}}) \quad (49)$$

$$= \log\left(\left(\sum_{k=i}^j \mathcal{F}_{m_k}\right)!\right) + \left(\sum_{k=i}^j m_k \mathcal{F}_{m_k}\right) \log\left(\frac{\sum_{k=i}^j m_k \mathcal{F}_{m_k}}{n \sum_{k=i}^j \mathcal{F}_{m_k}}\right) \quad (50)$$

where we used $|\mathcal{X}_{i:j}| = \sum_{k=i}^j \mathcal{F}_{m_k}$ and $\hat{p}_{\mathcal{X}_{i:j}} = \frac{1}{n|\mathcal{X}_{i:j}|} \sum_{k=i}^j m_k \mathcal{F}_{m_k}$.

Relations (46) and (50) allow us to compute the level set decomposition \bar{p}^* by filling out a $(F_+ + 1) \times F_+$ array $(\bar{V}(\mathcal{X}_{i:j}))_{i \in [0, F_+], j \in [1, F_+]}$ and keeping track of the maximizing index each time we compute \bar{V}_i . Once we have the optimal level set decomposition $\mathcal{A}^* = \{\mathcal{X}_{i:j}\}$ of \bar{p}^* , we set for all $x \in \mathcal{X}_{i:j}$, for all i, j

$$\bar{p}_x^* = \hat{p}_{\mathcal{X}_{i:j}} \tag{51}$$

For example, in Figure 1, we have $(m_i)_{i=0}^{F_+=12} = (0, 1, 2, 3, 4, 5, 7, 9, 12, 13, 15, 16, 20)$ and the approximate PML distribution \bar{p}^* (red line) has level sets $\{\mathcal{X}_{0:6}, \mathcal{X}_{7:12}\}$.

The running time of this dynamic programming algorithm is $O(|\text{Supp}(\mathcal{F})|^2)$, where $|\text{Supp}(\mathcal{F})| \leq F_+ + 1$. In terms of the sample size n , any empirical distribution \hat{p} satisfies $|\text{Supp}(\mathcal{F}(\hat{p}))| \leq \frac{1}{2}(1 + \sqrt{8n + 1})$, with equality achieved by the empirical distribution $\hat{p}_i = \frac{2^i}{|\mathcal{X}|(|\mathcal{X}|+1)}$ for $i \in [0, |\mathcal{X}|]$ with $|\mathcal{X}| = |\text{Supp}(\mathcal{F})|$ and $n = \frac{1}{2}|\mathcal{X}|(|\mathcal{X}| + 1)$. Thus the worst-case run time for the dynamic programming algorithm is $O(\sqrt{n}^2) = O(n)$. We “usually” have $|\text{Supp}(\mathcal{F})|$ much smaller than \sqrt{n} , so this is a pessimistic estimate for a typical case. The run time of our approximate PML scheme is usually dominated by computation of \hat{p} . An implementation optimization is to pre-compute the sums $\sum_{k=i}^j \mathcal{F}_{m_k}$ and $\sum_{k=i}^j m_k \mathcal{F}_{m_k}$ for all $i \leq j$ in computing $\bar{V}(\mathcal{X}_{i:j})$.

4.6. Unknown or infinite support set size

If the support set size $K = |\mathcal{X}|$ is unknown, then we can attempt to infer it. Our estimator \bar{K}^* for the support set size is the same as (17), but replaces the permanent with our lower bound (31), $e^{\bar{V}(p)} \leq \text{perm}(Q)$:

$$\bar{K}^* \triangleq \underset{K}{\text{argmax}} \left(\frac{1}{(K - \hat{K})!} \max_{p \in \mathcal{P}_K} e^{\bar{V}(p)} \right) \tag{52}$$

whenever the max over K exists, where $\mathcal{P}_K \triangleq \{p \in \mathcal{P} : |\text{Supp}(p)| = K\}$, where $\hat{K} = |\hat{\mathcal{X}}|$ is the support set size of the empirical distribution \hat{p} . Then the PML distribution is \bar{p}^* with support set \bar{K}^* . The max over p is obtained via our approximate PML algorithm in Section 4.5. The maximizer \bar{K}^* usually exists because a larger K boosts the value of the log permanent lower bound $\bar{V}(p)$ (since there are more permutations to sum over), but reduces the first factor in (52). It may happen that \bar{K}^* does not exist, in which case we say our approximate PML distribution has a continuous part in the terminology of (Orlitsky et al., 2004b). In general the function in the argmax in (52) is multimodal in K .

It turns out we can efficiently compute \bar{K}^* if it exists and compute the corresponding approximate PML distribution; if \bar{K}^* does not exist, then we can efficiently detect this case and compute the corresponding approximate PML distribution, which in this case has a continuous part.

We state and derive our observations in Appendix J. The idea is to optimize over all possible ways to clump the unseen symbols $\mathcal{X}_{0:0}$ with the other symbols. Since the optimal clumping satisfies the iso-clumping and convexity properties, we show that there are only $|\text{Supp}(\mathcal{F}_+)| \leq \sqrt{2n} + 1$ possibilities to check and that the value of each possibility can be checked quickly, reusing the work done in the dynamic programming approach described in the previous Section, so this can be done efficiently.

4.7. Approximate PML: the case of multiple distributions (see Appendix H)

We generalize the PML distribution and approximate PML distribution to the case of $(n_d)_{d=1}^D$ samples drawn from distributions $(p^{(d)})_{d=1}^D$ on the same alphabet (see Section 3.2 for the case $D = 2$, used in estimating functionals of pairs of distributions). This requires some more notation that we delegate to Appendix H.

For $D \geq 2$, we are unable to give an efficient algorithm to compute the D -dimensional approximate PML. The difficulty is due to the lack of a natural ordering on \mathbb{N}^D for $D \geq 2$, so the dynamic programming approach we use for $D = 1$ does not work for $D \geq 2$. We settle for approximately computing our the approximate PML distribution via a greedy heuristic: we iteratively enlarge level sets of a candidate solution by merging pairs of distinct level sets until the objective function \bar{V} (suitably generalized to D distributions) stops increasing. See Appendix H for details.

4.8. Comparison to the (exact) PML distribution for size 2 alphabet

Recall the discussion of Section 4.2.1 on the PML distribution for alphabet size $|\mathcal{X}| = 2$. A computation establishes the condition for uniformity of the APML distribution:

$$\bar{p}^* = \begin{cases} (\frac{1}{2}, \frac{1}{2}) & : H(\hat{p}) \geq \log(2) (1 - \frac{1}{n}) \\ (\hat{p}_1, 1 - \hat{p}_1) & : \text{otherwise} \end{cases} \quad (53)$$

where $H(\hat{p})$ is the entropy of the empirical distribution and n is the sample size.

Let's consider the large n limit. Let \hat{p}_1^{unif} be the critical value at which the APML distribution switches between uniform and nonuniform (the least value of \hat{p}_1 satisfying the uniformity condition in (53)). Then a computation shows that

$$\hat{p}_1^{\text{unif}} = \frac{1}{2} \left(1 - \sqrt{\frac{\log(4)}{n}} + O\left(\frac{1}{n}\right) \right) \quad (54)$$

Comparing this to condition (26) for the (exact) PML, we see that in the APML case, \hat{p}_1^{unif} carries an extra factor of $\sqrt{\log(4)} \approx 1.18$ in the $1/\sqrt{n}$ term relative to the uniform/nonuniform critical value of \hat{p}_1 for the PML.

Figure 6 compares the APML and PML distributions for a size-2 alphabet. Both the APML and PML are uniform when $|\hat{p}_1 - \hat{p}_2|$ is small enough, though at different points, as remarked above. Whenever the APML solution is not uniform, it exactly matches (up to relabeling of components) the empirical distribution, unlike the PML.

5. Estimating symmetric functionals of a single discrete distribution

In the standard PML setting discussed in Section 3.1, we argued the sufficiency of the fingerprint \mathcal{F} (11) in estimating the sorted probability vector of a discrete distribution, which implies the sufficiency of the fingerprint for estimating any functional F of the sorted probability vector of the form

$$F(p) = G\left(\sum_{x \in \mathcal{X}} f(p_x)\right), \quad (55)$$

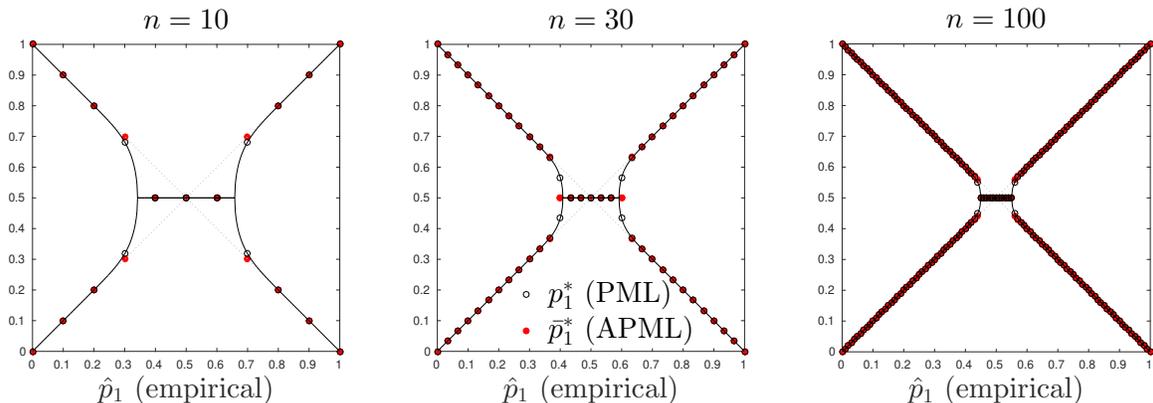


Figure 6: “Phase diagram” for the exact PML (hollow circles) and APML distributions (red circles) on binary alphabet with sample size n indicated above each subplot (compare to Figure 3). Solid black lines correspond to the solution for arbitrary $\hat{p}_1 \in [0, 1]$, obtained numerically. The PML distribution p^* is uniform when $|\hat{p}_1 - \hat{p}_2| \leq 1/\sqrt{n}$, while the APML distribution is uniform when $|\hat{p}_1 - \hat{p}_2| \leq \sqrt{\log(4)}/\sqrt{n} + O(1/n)$. Outside this middle region there are two PML (APML) branches, since any permutation of the PML (APML) distribution is another PML (APML) distribution. The APML distribution exactly matches the empirical distribution outside the uniform region. Diagonal dashed lines correspond to the lines (\hat{p}_1, \hat{p}_1) and $(\hat{p}_1, 1 - \hat{p}_1)$.

where we constrain $f(0) = 0$ to accommodate the unknown alphabet setting. Note that the fingerprint is also partially sufficient for functionals of type $\sum_{x,y \in \mathcal{X}} f(p_x, p_y)$, etc.

In order to apply Theorem 3, one also needs to show the cardinality of the fingerprint \mathcal{F} is small. The fingerprint $\mathcal{F} = (\mathcal{F}_i)_{i \geq 0}$ satisfies

$$\sum_{i \geq 0} i \mathcal{F}_i = n, \quad (56)$$

and each fingerprint corresponds to an unordered integer partition of the integer n . We now recall the Hardy-Ramanujan result on integer partitions.

Theorem 10 (*Hardy and Ramanujan, 1918*)(*Bárány and Vershik, 1992, Theorem 2*) *The cardinality of the set of fingerprints (11) on n samples is given by*

$$e^{\pi \sqrt{\frac{2n}{3}(1-o(1))}} \leq |\{(\mathcal{F}_i)_{i \geq 1}\}| \leq e^{\pi \sqrt{\frac{2n}{3}}}. \quad (57)$$

The key observation is that the cardinality of the fingerprint is sub-exponential. Then, combining with Theorem 3 and the fact that there exist estimators for various symmetric functionals with near exponential measure concentration, (Acharya et al., 2017) showed that plugging in the PML achieves the optimal sample complexity in estimating the Shannon entropy, the support size, the support coverage, and the L_1 distance to uniformity.

In this section, we extensively test the performance of symmetric functional estimation via plugging in our approximate PML (APML) distribution into the functionals. Recall

that the PML distribution p^* is defined as in (16) for empirical distribution \hat{p} on n samples with fingerprint $\mathcal{F}(\hat{p})$:

$$p^* = \operatorname{argmax}_{p \in \mathcal{P}} \mathbb{P}_p(\mathcal{F}) = \operatorname{argmax}_{p \in \mathcal{P}} \frac{\operatorname{perm} \left(\left(p_x^{n\hat{p}_{x'}} \right)_{x, x' \in \mathcal{X}} \right)}{(K - \hat{K})!}. \quad (58)$$

where the optimization is over distributions of possibly different support set size, \mathcal{X} denotes the support set of distribution p , $K = |\mathcal{X}|$, and \hat{K} denotes the support set size of the empirical distribution \hat{p} . p^* is difficult to compute exactly, so we compute the APML distribution \bar{p}^* (33), (52):

$$\bar{p}^* = \operatorname{argmax}_{p \in \mathcal{P}} \frac{e^{\bar{V}(p)}}{(K - \hat{K})!} \quad (59)$$

which maximizes a lower bound to the probability $\mathbb{P}_p(\mathcal{F})$ since $\bar{V}(p)$ (31) lower bounds the log permanent in (58). Then our estimator for function F of the form (55) is the plugin $F(\bar{p}^*)$.

Our approximate PML estimator performs comparably well to the competition (Valiant and Valiant, 2011a) (Valiant and Valiant, 2013) (Jiao et al., 2015), (Wu and Yang, 2016), (Wu and Yang, 2015) across different functions F and distributions, and performs significantly better when the true distribution is uniform. This good performance for the uniform case makes some intuitive sense: the approximate PML distribution \bar{p}^* maximizes a lower bound (22) to a matrix permanent obtained by summing over only those permutations that mix symbols within level sets of \bar{p}^* . If there is only one level set – that is, if \bar{p}^* is uniform – then we have exactly optimized the matrix permanent (or at least found a local maximum), rather than a lower bound to the permanent.

5.1. Sorted L_1 loss

Given n samples from a distribution p , the usual L_1 loss in measuring the accuracy of estimating p is given by

$$\sum_{x \in \mathcal{X}} |p_x - q_x|. \quad (60)$$

The sorted L_1 loss measures the *shape* difference between the reconstruction distribution q and the true distribution p . In other words, we first sort the distributions p, q into non-increasing probability vectors $(p_{(1)}, p_{(2)}, \dots, p_{(K)}), (q_{(1)}, q_{(2)}, \dots, q_{(K)})$, and then measure the loss by

$$\sum_{1 \leq i \leq K} |p_{(i)} - q_{(i)}|. \quad (61)$$

The sorted L_1 loss measures the error in reconstructing the sorted probability vector, for which the fingerprints are partially sufficient as shown by Theorem 6. The sorted L_1 loss can also be viewed as a Wasserstein distance (Vallender, 1974), i.e.,

$$\sum_{1 \leq i \leq K} |p_{(i)} - q_{(i)}| = K \sup_{f \in \operatorname{Lip}_1} \int f(x) \left(\sum_x \frac{1}{K} \delta_{p_x} - \sum_x \frac{1}{K} \delta_{q_x} \right), \quad (62)$$

where the supremum is over all Lipschitz functions with Lipschitz constant one.

We estimate the sorted L_1 loss using the plug-in $q = \bar{p}^*$, the APML distribution (59), computed over collection of distributions $\mathcal{P} = \Delta$, where

$$\Delta \triangleq \{p : |\text{Supp}(p)| < \infty\} \tag{63}$$

denotes the set of all discrete distributions with finite support set size, and collection of distributions $\mathcal{P} = \Delta_K$, where

$$\Delta_K \triangleq \{p : |\text{Supp}(p)| = K\} \tag{64}$$

denotes the set of all discrete distributions with support set size K .

Figure 7 shows the sorted probability vector inferred by our approximate PML distribution in Δ (63) – that is, the case of unknown support set size – along with the sorted ML (empirical) distribution and the distribution of (Valiant and Valiant, 2013). We see that both the approximate PML distribution and the distribution of (Valiant and Valiant, 2013) perform much better than the sorted ML distribution, and that the PML distribution “prefers” to have fewer and larger level sets than the distribution of (Valiant and Valiant, 2013).

Figure 8 shows the performance of our approximate PML distribution in Δ_K (64) – that is, the case of known support set size – as a plugin estimator for the sorted L_1 loss, along with the estimator of (Valiant and Valiant, 2011a) and the empirical distribution plugin estimator (MLE) (see caption for performance test parameters). We see that the approximate PML distribution performs well for the uniform and mixture-of-two-uniforms distributions, but more poorly for the Zipf distributions, consistent with the remarks in the previous Section. Good performance on near-uniform distributions is observed for other performance tests as well (see below).

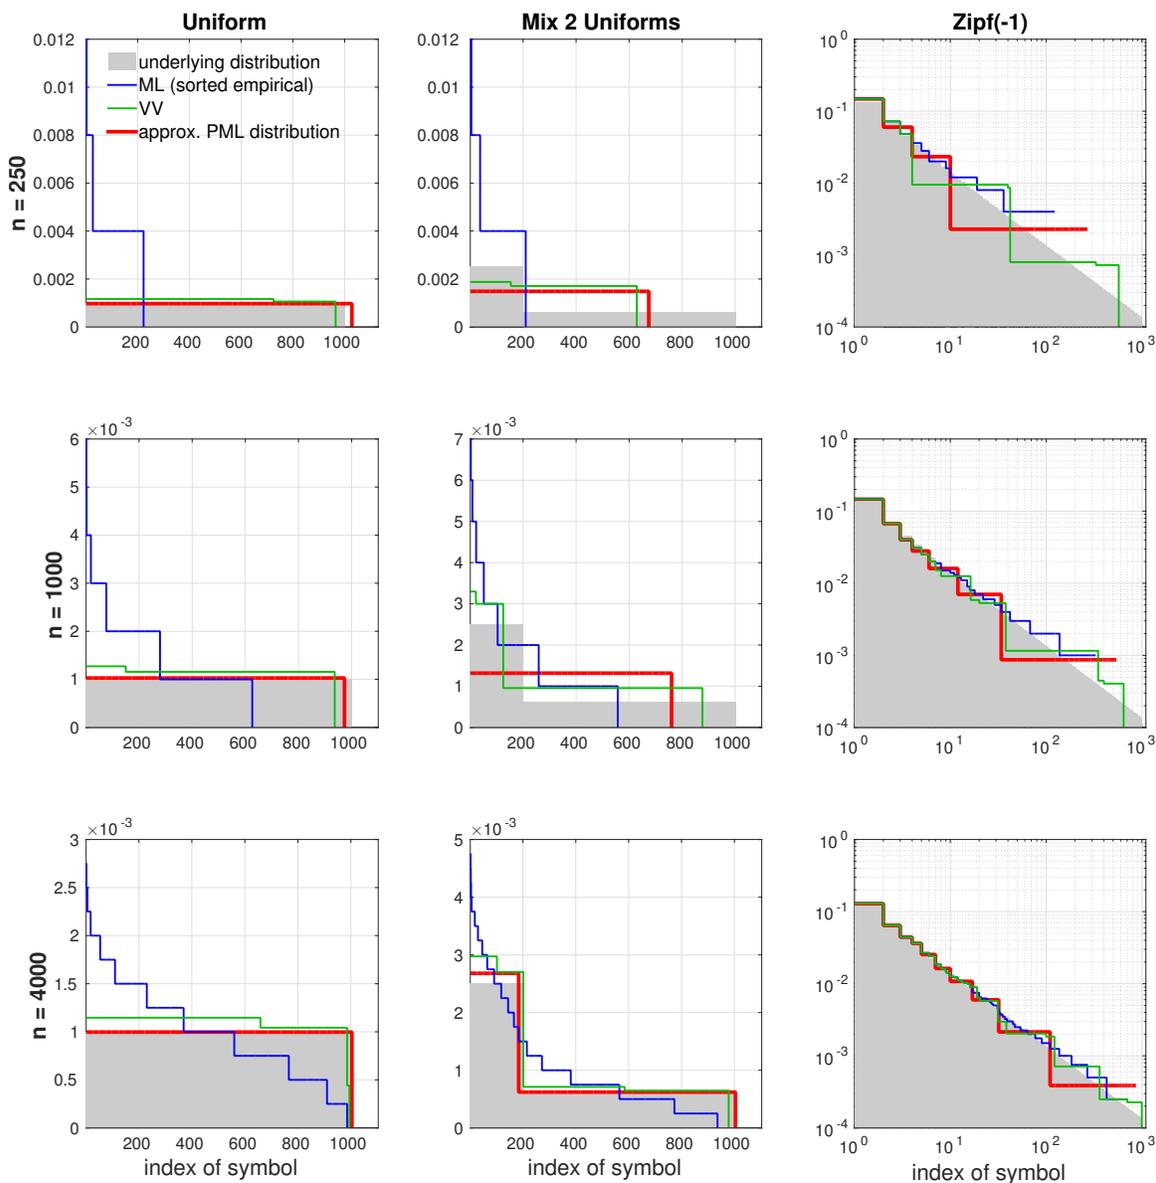


Figure 7: Inferring the sorted probability vector given samples from an underlying distribution (plotted in gray). The distribution is held constant in each column and the sample size n is held constant in each row of the figure. In all cases $K = |\mathcal{X}| = 1000$. “Uniform” is uniform on \mathcal{X} , “Mix 2 Uniforms” is a mixture of two uniform distributions, with half the probability mass on the first $K/5$ symbols, and the other half on the remaining symbols, and Zipf(α) $\sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$. ML denotes the sorted empirical distribution. VV is (Valiant and Valiant, 2013).

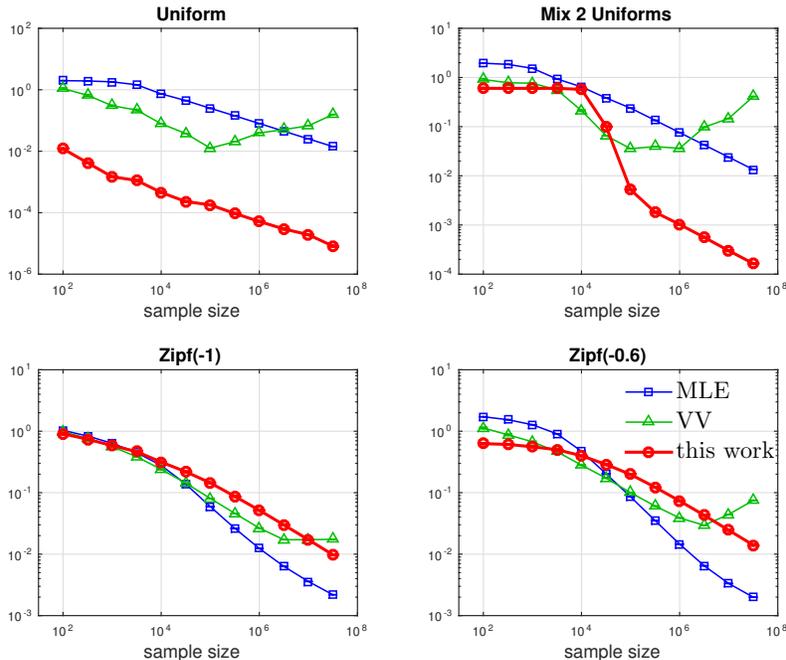


Figure 8: Sorted L_1 loss in estimating an unknown distribution p with known support set size $K = |\mathcal{X}|$. In all cases $K = 10^4$. “Uniform” is uniform on \mathcal{X} , “Mix 2 Uniforms” is a mixture of two uniform distributions, with half the probability mass on the first $K/5$ symbols, and the other half on the remaining symbols, and Zipf(α) $\sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$. MLE denotes the ML plugin (“naive”) approach of using the sorted empirical distribution in (61). VV is (Valiant and Valiant, 2013). Each data point represents 100 random trials, with 2 standard error bars smaller than the plot marker for most points.

5.2. Entropy and Rényi entropy estimation

For entropy $H(p) = -\sum_{x \in \mathcal{X}} p_x \log(p_x)$ and Rényi entropy $H_\alpha(p) = \frac{1}{1-\alpha} \log(\sum_{x \in \mathcal{X}} p_x^\alpha)$ of distribution p , the corresponding approximate PML estimator is defined as the plugin estimator $H(\bar{p}^*)$, $H_\alpha(\bar{p}^*)$, where \bar{p}^* is as in (59), optimized over the collection of distributions $\mathcal{P} = \Delta$ (63) for both the entropy and Rényi entropy. Additionally, to enable direct comparison with the entropy estimator of (Wu and Yang, 2016), which requires a support set size as input, we also find the approximate PML distribution optimized over $\mathcal{P} = \Delta_K$ (64).

Figure 9 shows the performance of our approximate PML scheme for estimating the entropy and Rényi entropy with $\alpha \in \{2, 1.5, 0.8\}$. Overall, our approximation looks competitive with the other approaches, and is particularly strong for the uniform distribution. Adding in knowledge of the true support set size in estimating the entropy did not make much of a difference, except for improvement in the case of the uniform distribution.

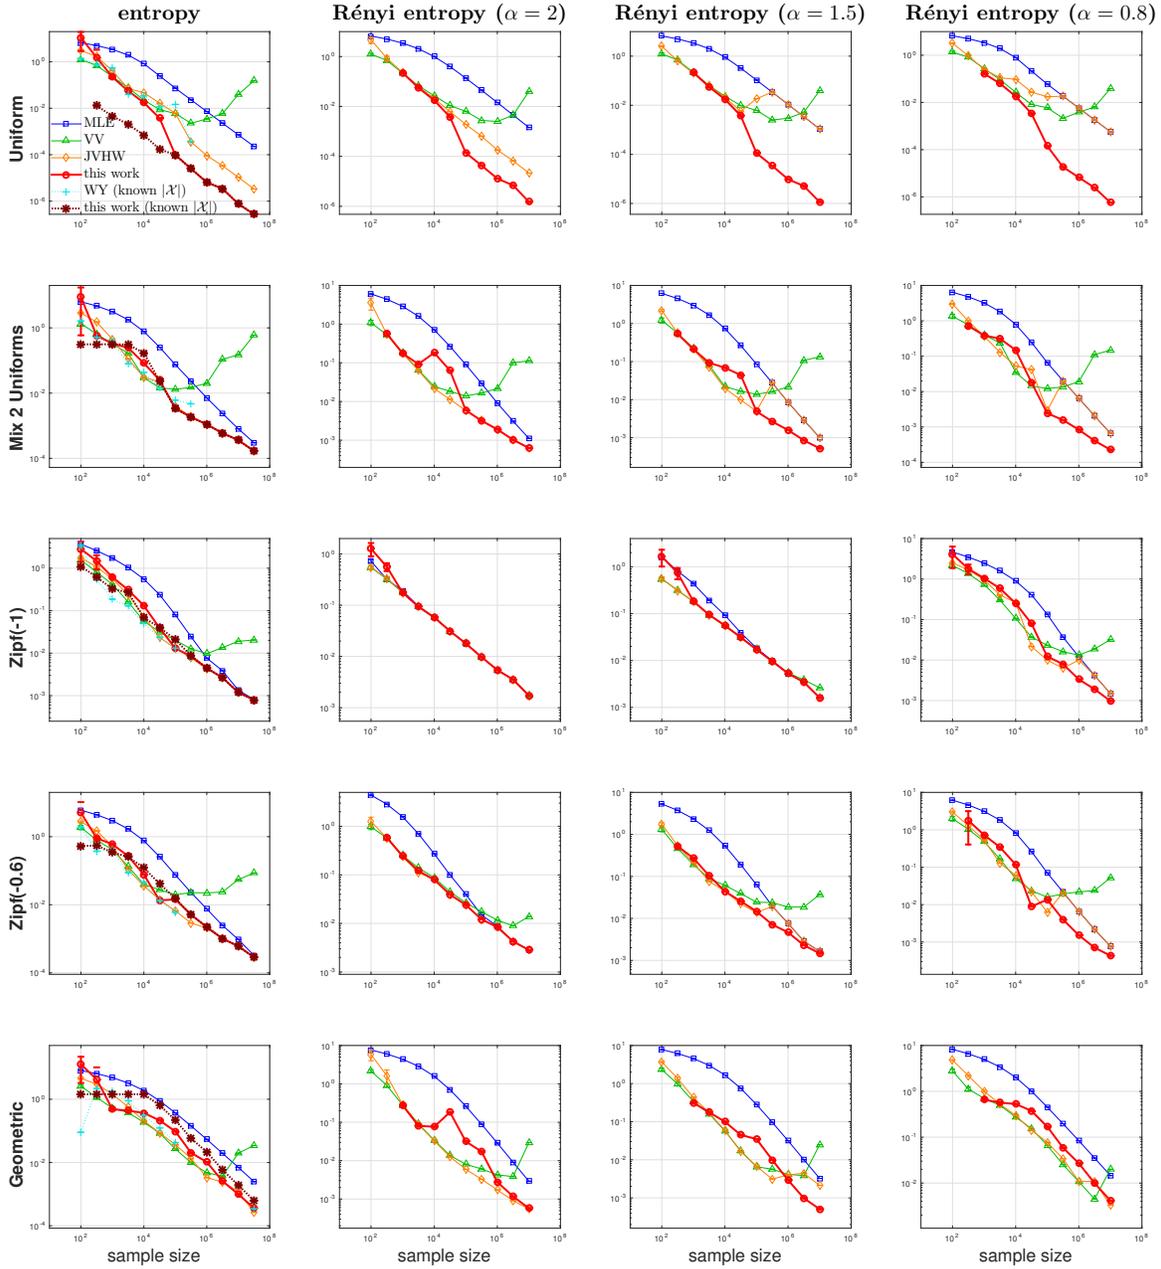


Figure 9: Root mean squared error of several estimators of entropy (first column) and Rényi entropy (last three columns for $\alpha \in \{2, 1.5, 0.8\}$). In all cases except “Geometric” we set the alphabet size $|\mathcal{X}| = 10^4$. “Uniform” is uniform on \mathcal{X} , $\text{Zipf}(\alpha) \sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$, “Geometric” is the geometric distribution with infinite support and mean K . MLE denotes the ML plugin (“naive”) approach of computing $H(\hat{p})$. VV is (Valiant and Valiant, 2013). JVHW is (Jiao et al., 2015). WY is (Wu and Yang, 2016) – this estimator requires the support set size $|\mathcal{X}|$ as an input. Our estimator optionally accepts the support set size $|\mathcal{X}|$ (shown in dark red). Each data point represents 100 random trials, with 2 standard error bars smaller than the plot marker for most points. Log base 2.

5.3. L_1 distance to uniformity

For estimating the L_1 distance to uniformity, which is $\sum_{x \in \mathcal{X}} |p_x - \frac{1}{K}|$, the corresponding PML estimator is (Acharya et al., 2017) the plugin of p^* (58) optimized over the collection of distributions $\mathcal{P} = \Delta_K$ (64) (the case of known support set size). The APML estimator is the plugin of \bar{p}^* (59) optimized over $\mathcal{P} = \Delta_K$.

Figure 10 shows the performance of our approximate PML scheme for estimating the L_1 distance to a uniform distribution with known support set size. Our approach looks competitive, and is by far the best if the true distribution is uniform.

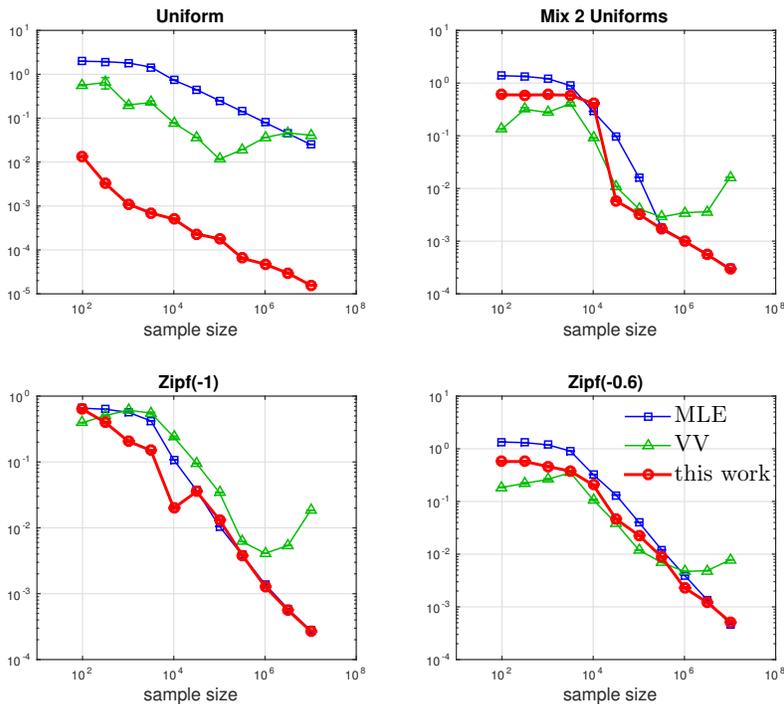


Figure 10: Root mean squared error for estimation of L_1 distance to uniformity with known support set size $K = |\mathcal{X}|$. In all cases $K = 10^4$. “Uniform” is uniform on \mathcal{X} , “Mix 2 Uniforms” is a mixture of two uniform distributions, with half the probability mass on the first $K/5$ symbols, and the other half on the remaining symbols, and $\text{Zipf}(\alpha) \sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$. MLE denotes the ML plugin (“naive”) approach of computing $\sum_{x \in \mathcal{X}} |\hat{p}_x - \frac{1}{K}|$. VV is (Valiant and Valiant, 2013). Each data point represents 100 random trials, with 2 standard error bars smaller than the plot marker for most points.

5.4. Support set size estimation

For support size estimation, the corresponding PML estimator is (Acharya et al., 2017) the plugin of p^* (58) optimized over the collection of distributions $\mathcal{P} = \Delta_{\geq \frac{1}{K}}$, where

$$\Delta_{\geq \frac{1}{K}} \triangleq \{p : p_x \geq \frac{1}{K} \forall x \in \mathcal{X}\} \quad (65)$$

denotes the set of all discrete distributions whose minimum nonzero probability for each symbol is at least $\frac{1}{K}$. The approximate PML estimator is the plugin of \bar{p}^* (59) optimized over $\mathcal{P} = \Delta_{\geq \frac{1}{K}}$.

Figure 11 shows the performance of our approximate PML scheme for estimating the support set size. Here we plot the mean and standard error of the support set size inferred by the different estimation schemes rather than a root mean squared error. Overall, our approach is comparable to the others, performing worst on the Zipf distributions, and best on the uniform distribution.

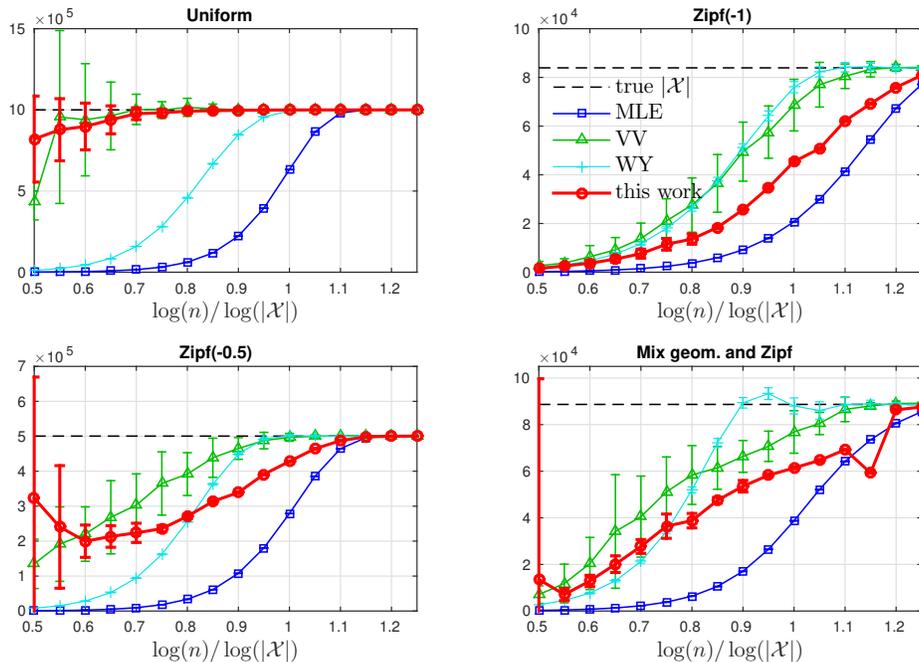


Figure 11: Support set size estimation for the same cases as in (Wu and Yang, 2015), plotting the mean and 2 standard deviations of the distribution of inferred support set sizes in 100 trials vs. sample size n . In all cases the support set size $K = |\mathcal{X}|$ is chosen such that $\min_{i \in \{1, \dots, K\}} p_i \approx 10^{-6}$. “Uniform” is uniform on $\{1, \dots, K\}$, Zipf(α) $\sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$, and “Mix geom. and Zipf” is a mixture of $p_i \sim 1/i$ for $i \in \{1, \dots, K/2\}$ and $p_i \sim (1 - 2/K)^{i-K/2-1}$ for $i \in \{K/2 + 1, \dots, K\}$ normalized so that half the mass is on the first $K/2$ symbols. MLE denotes the ML plugin (“naive”) approach of using the empirical support set size. VV is (Valiant and Valiant, 2013). WY is (Wu and Yang, 2015).

6. Estimating the divergence function between two distributions

In the 2-D PML setting discussed in Section 3.2, we argued the sufficiency of the fingerprint (18) in estimating functions of two discrete distributions of the form

$$F(p, q) = \sum_{x \in \mathcal{X}} f(p_x, q_x). \quad (66)$$

In order to apply Theorem 3, one needs to show the cardinality of the sufficient statistic is not too big. Below we present an argument that applies to the D -dimensional PML. Suppose we have D discrete distributions on the same alphabet \mathcal{X} , denoted as $p^{(1)}, p^{(2)}, \dots, p^{(D)}$. Suppose we obtain n_d samples from each distribution $p^{(d)}$ with empirical distribution $\hat{p}^{(d)}$. Let the joint fingerprint be defined for every $(i_1, i_2, \dots, i_D) \geq 0$,

$$\mathcal{F}_{i_1, i_2, \dots, i_D} = |\{x \in \mathcal{X} : (n_d \hat{p}_x^{(d)})_{d=1}^D = (i_d)_{d=1}^D\}|. \quad (67)$$

We now argue that the joint fingerprint $(\mathcal{F}_{i_1, i_2, \dots, i_D})_{(i_1, i_2, \dots, i_D) \geq 0}$ corresponds to the partition of multipartite numbers. Indeed, the joint fingerprint satisfies the following equation:

$$\sum_{(i_1, i_2, \dots, i_D) \geq 0} \mathcal{F}_{i_1, i_2, \dots, i_D} \begin{bmatrix} i_1 \\ i_2 \\ \vdots \\ i_D \end{bmatrix} = \begin{bmatrix} n_1 \\ n_2 \\ \vdots \\ n_D \end{bmatrix}. \quad (68)$$

Theorem 11 quantifies the size of the sufficient statistic.

Theorem 11 (*Auluck, 1953; Bárány and Vershik, 1992; Acharya et al., 2011*) *Suppose $D = 2$. If $n_1 = n_2 = n$, then the cardinality of the set of 2-D fingerprints on n samples is given by*

$$|\{(\mathcal{F}_{ij})_{(i,j) \neq (0,0)}\}| = e^{3(\zeta(3))^{1/3} n^{2/3} (1+o(1))}. \quad (69)$$

For general $D > 2$, if $n_d \geq 2^{D+1}$, $1 \leq d \leq D$, we have

$$|\{(\mathcal{F}_{i_1, i_2, \dots, i_D})_{(i_1, i_2, \dots, i_D) \neq \mathbf{0}}\}| \leq \exp\left(2\left(1 + \frac{1}{D}\right) \sum_{d=1}^D n_d^{\frac{D}{D+1}}\right). \quad (70)$$

where $\mathbf{0} = (0, \dots, 0)$. Moreover, if $n_d = n$ for all $d \in \{1, \dots, D\}$, we have

$$|\{(\mathcal{F}_{i_1, i_2, \dots, i_D})_{(i_1, i_2, \dots, i_D) \neq \mathbf{0}}\}| = \exp\left((D+1)(\zeta(D+1))^{1/(D+1)} n^{\frac{D}{D+1}} (1 - o(1))\right). \quad (71)$$

Here $\zeta(D+1) = \sum_{k=1}^{\infty} k^{-(D+1)}$ is the Riemann zeta function.

Fortunately, the cardinality of the D -dimensional fingerprint is still sub-exponential in the sample size for all D . Then, combining with Theorem 3 and the fact that there exist estimators for the aforementioned divergence functions with near exponential measure concentration (which can be shown in a straightforward manner as in (Acharya et al., 2017;

Han et al., 2016a,b; Bu et al., 2016; Valiant and Valiant, 2011a; Jiao et al., 2018)), we know that plugging in the 2-D PML achieves the optimal sample complexity in estimating the Kullback-Leibler divergence, L_1 distance, the squared Hellinger distance, and the χ^2 divergence. The proof of these results is to be reported elsewhere.

In this section, we extensively test the performance of divergence estimation via plugging in our approximate 2-D PML. We mainly compete with the released code of KL divergence estimator in (Han et al., 2016a,b). The concrete divergence functional estimation algorithm is as follows. Suppose we observe n samples from distribution p with empirical distribution \hat{p} and m samples from distribution q with empirical distribution \hat{q} . Let the 2-D PML estimator be as (20):

$$(p^*, q^*) = \operatorname{argmax}_{(p,q) \in \mathcal{P}} \mathbb{P}_{p,q}(\mathcal{F}), \quad (72)$$

where \mathcal{F} is the 2-D fingerprint (18) and \mathcal{P} is a collection of pairs of distributions on the same alphabet \mathcal{X} . The APML distributions maximize a lower bound to $\mathbb{P}_{p,q}(\mathcal{F})$:

$$(\bar{p}^*, \bar{q}^*) = \operatorname{argmax}_{(p,q) \in \mathcal{P}} \frac{e^{\bar{V}(p,q)}}{(K - \hat{K})!} \quad (73)$$

where \bar{V} is defined for the D -dimensional case in Appendix H, $\mathcal{X} = \operatorname{Supp}(p) \cup \operatorname{Supp}(q)$, $K = |\mathcal{X}|$, and $\hat{K} = |\operatorname{Supp}(\hat{p}) \cup \operatorname{Supp}(\hat{q})|$.

As discussed in Section 4.7, there is no natural ordering on the bins of the 2-D fingerprint, so we do not give a dynamic programming algorithm like the one in section 4.5 to find the APML distributions. Instead we use a greedy heuristic presented in Appendix H to approximately maximize $\bar{V}(p, q)$ and approximate the already-approximate PML distributions:

$$(\bar{\bar{p}}^*, \bar{\bar{q}}^*) \approx (\bar{p}^*, \bar{q}^*) \quad (74)$$

Then our estimator for functional F of the form (66) is $F(\bar{\bar{p}}^*, \bar{\bar{q}}^*)$.

Our approximate PML estimator performs overall best relative to the competition, which for the KL divergence consists of only (Han et al., 2016a,b) and the MLE plugin estimator, and for the L_1 distance consists of only the MLE plugin estimator. (Valiant and Valiant, 2013), (Valiant, 2008) generalize their approach to the 2-D fingerprint setting, but do not release code to repeat their experiments.

6.1. KL divergence estimation

For the KL divergence $D(p||q) = \sum_{x \in \mathcal{X}} p_x \log \left(\frac{p_x}{q_x} \right)$, the corresponding approximate PML estimator is the plugin $D(\bar{\bar{p}}^*||\bar{\bar{q}}^*)$, where $\bar{\bar{p}}^*, \bar{\bar{q}}^*$ are as in (74), optimized over collection of distributions $\mathcal{P} = \Delta_\rho$, where

$$\Delta_\rho = \{(p, q) : \sup_x (p_x/q_x) \leq \rho\} \quad (75)$$

denotes the set of all pairs of discrete distributions with finite support and maximum ratio ρ .

Figure 12 shows the performance of our approximate PML scheme for estimating the KL divergence. The approximate PML estimator looks mostly better than (Han et al., 2016a,b)

and the MLE. All three estimators show non-monotonicity of their root mean squared performance in the sample size. Qualitatively, the approximate PML estimator’s performance looks best-behaved among the three in terms of being roughly monotone decaying in the sample size.

6.2. L_1 distance

For the L_1 distance $\|p - q\|_1 = \sum_{x \in \mathcal{X}} |p_x - q_x|$, the corresponding approximate PML estimator is the plugin $\|\bar{p}^* - \bar{q}^*\|_1$, where \bar{p}^*, \bar{q}^* are as in (74), optimized over collection of distributions $\mathcal{P} = \Delta$, where Δ denotes the set of all pairs of discrete distributions with finite support set size (possibly different support set sizes within the pair).

Figure 13 shows the performance of our approximate PML scheme for estimating the L_1 distance between two distributions. The approximate PML estimator looks overall stronger than the MLE plugin estimator for all distributions and sample sizes, and seems to perform best when the two unknown distributions p and q are the same (corresponding to the diagonal of the matrix of plots in the figure). Although (Valiant and Valiant, 2013) did not release their code for estimating the L_1 distance, their Figure 2, leftmost pane corresponds exactly in its parameters to the top-leftmost pane of our Figure 13; from roughly checking a few sample sizes visually, the approximate PML seems to perform significantly better.

Acknowledgments

We are grateful to Huacheng Yu for providing the proof of Theorem 12. We would like to thank Jayadev Acharya, Nick Loehr, Greta Panova, Richard Stanley, and Arthur Yang for helpful discussions on the size of the sufficient statistic.

Appendix A. Proof of expression (14) for the probability of a fingerprint

Let $\mathcal{S}_{\mathcal{X}}$ denote the symmetric group on the discrete, possibly infinite set \mathcal{X} and let $\mathcal{S}_{\mathcal{X}}\hat{p}$ denote the orbit of empirical distribution \hat{p} under the action of $\mathcal{S}_{\mathcal{X}}$ by relabeling the components of \hat{p} . That is, if $\sigma \in \mathcal{S}_{\mathcal{X}}$, then $(\sigma\hat{p})_x \triangleq \hat{p}_{\sigma(x)}$, and $\mathcal{S}_{\mathcal{X}}\hat{p} = \{\sigma\hat{p}\}_{\sigma \in \mathcal{S}_{\mathcal{X}}}$. Each distribution in the orbit $\mathcal{S}_{\mathcal{X}}\hat{p}$ has the same fingerprint and distinct orbits have distinct fingerprints, so the set of distinct fingerprints \mathcal{F} is in bijection with the set of distinct orbits under the action of $\mathcal{S}_{\mathcal{X}}$, so the probability of a specific fingerprint \mathcal{F} is given by

$$\mathbb{P}_p(\mathcal{F}) = \mathbb{P}_p(\mathcal{S}_{\mathcal{X}}\hat{p}) = \sum_{\hat{p}' \in \mathcal{S}_{\mathcal{X}}\hat{p}} \mathbb{P}_p(\hat{p}') = \binom{n}{n\hat{p}} \sum_{\hat{p}' \in \mathcal{S}_{\mathcal{X}}\hat{p}} \prod_{x \in \mathcal{X}} p_x^{n\hat{p}'_x} \quad (76)$$

where the last equality follows from (10).

When the alphabet \mathcal{X} is finite, then we can replace the summation over the orbit $\mathcal{S}_{\mathcal{X}}\hat{p}$ in (76) with a summation over the finite symmetric group $\mathcal{S}_{\mathcal{X}}$, taking care to avoid double-counting, since $|\mathcal{S}_{\mathcal{X}}\hat{p}| \leq |\mathcal{S}_{\mathcal{X}}|$ with strict inequality if there exist distinct $x, x' \in \mathcal{X}$ such that $\hat{p}_x = \hat{p}_{x'}$ (see Appendix B for an example of this). The fingerprint notation lets us

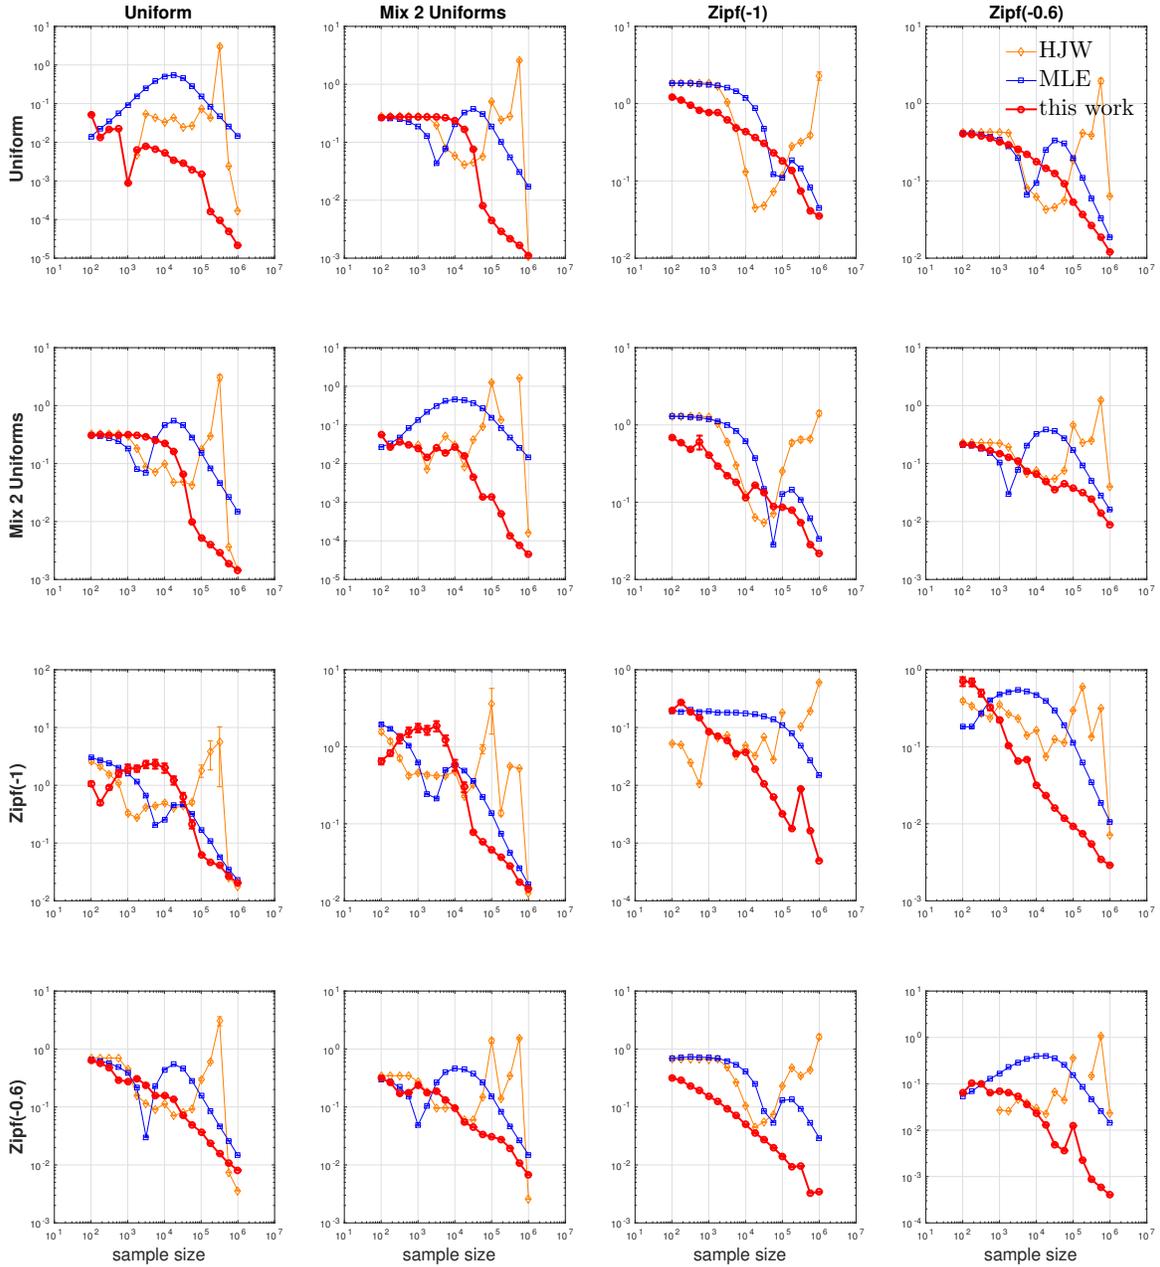


Figure 12: Root mean squared error for estimation of KL divergence between distributions indicated in row and column names (that is, $D(p_{\text{row}}||p_{\text{column}})$). In all cases $K = |\mathcal{X}| = 10^4$. “Uniform” is uniform on \mathcal{X} , “Mix 2 Uniforms” is a mixture of two uniform distributions, with half the probability mass on the first $K/5$ symbols, and the other half on the remaining symbols, and $\text{Zipf}(\alpha) \sim 1/i^\alpha$ with $i \in \{1, \dots, K\}$. MLE denotes the ML plugin (“naive”) approach of computing $D(\hat{p}||\hat{q})$. HJW is (Han et al., 2016a,b). Each data point represents 100 random trials, with 2 standard error bars smaller than the plot marker for most points. Log base 2.

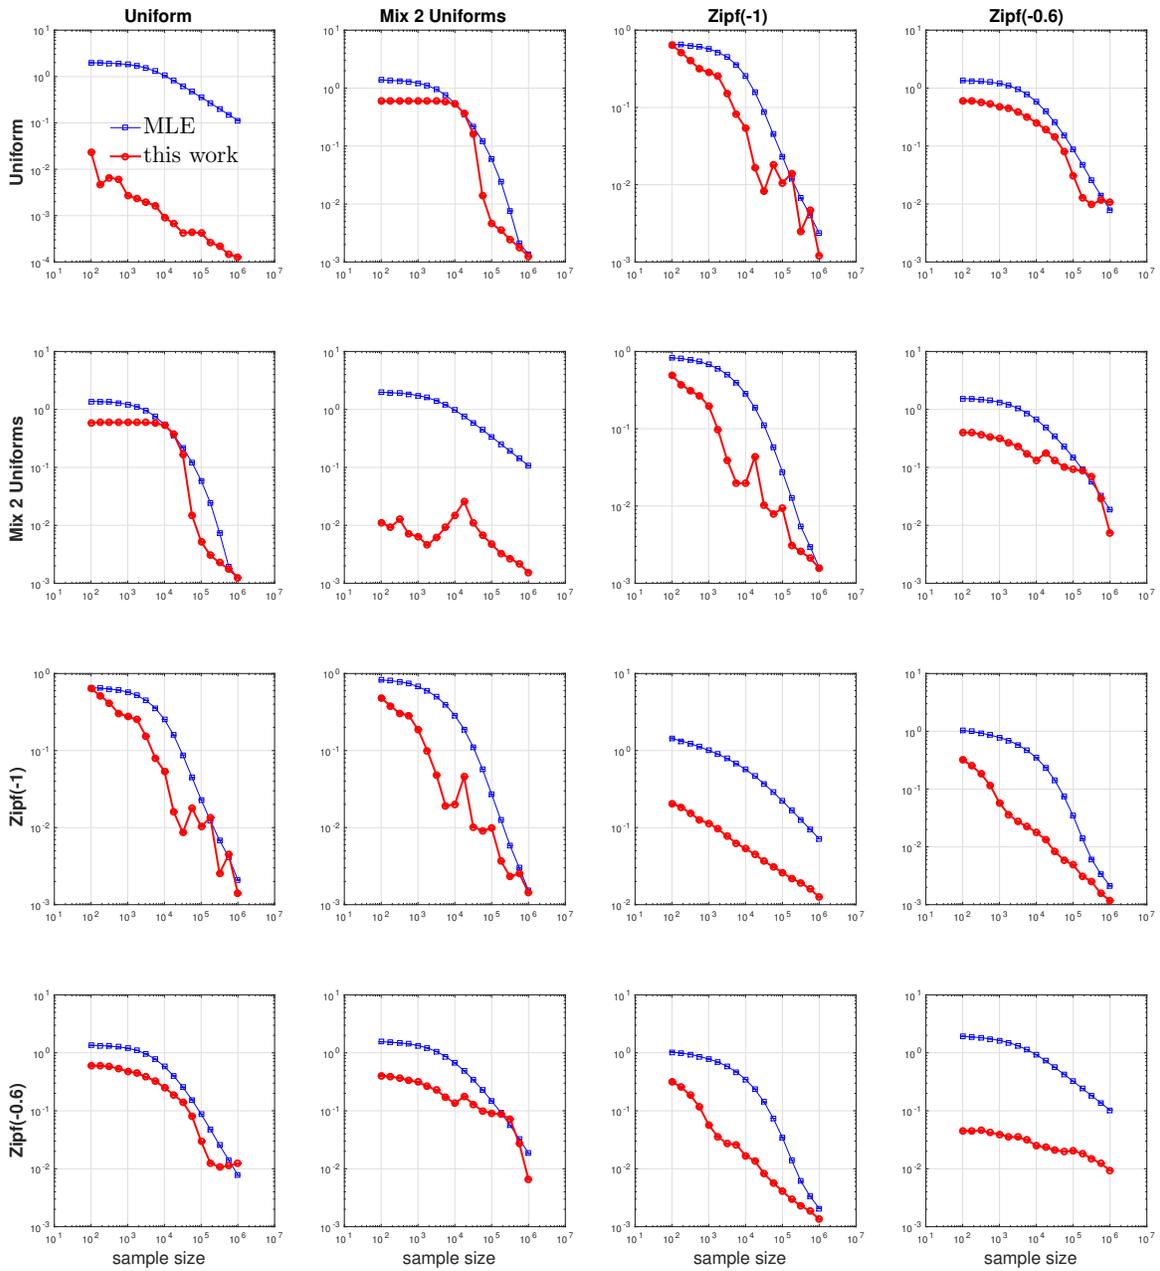


Figure 13: Root mean squared error for estimation of L_1 distance between distributions indicated in row and column names. Same parameters as in Figure 12. Since the L_1 distance is symmetric in its arguments, the performance matrix is symmetric.

conveniently express $|\mathcal{S}_{\mathcal{X}}\hat{p}|$:

$$|\mathcal{S}_{\mathcal{X}}\hat{p}| = \binom{|\mathcal{X}|}{\mathcal{F}} = |\mathcal{X}|! \prod_{i \geq 0} \frac{1}{\mathcal{F}_i!}. \quad (77)$$

The probability of a specific fingerprint \mathcal{F} under distribution p with finite support \mathcal{X} is given by

$$\mathbb{P}_p(\mathcal{F}) = \mathbb{P}_p(\mathcal{S}_{\mathcal{X}}\hat{p}) = \sum_{\hat{p}' \in \mathcal{S}_{\mathcal{X}}\hat{p}} \mathbb{P}_p(\hat{p}') \quad (78)$$

$$\stackrel{(a)}{=} \frac{|\mathcal{S}_{\mathcal{X}}\hat{p}|}{|\mathcal{S}_{\mathcal{X}}|} \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \mathbb{P}_p(\sigma\hat{p}) \quad (79)$$

$$\stackrel{(b)}{=} \frac{1}{|\mathcal{X}|!} \binom{|\mathcal{X}|}{\mathcal{F}} \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \mathbb{P}_p(\sigma\hat{p}) \quad (80)$$

$$\stackrel{(c)}{=} \left(\prod_{i \geq 0} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \prod_{x \in \mathcal{X}} p_x^{n\hat{p}_{\sigma(x)}} \quad (81)$$

$$= \left(\prod_{i \geq 0} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \text{perm} \left(\underbrace{\left(p_x^{n\hat{p}_{x'}} \right)_{x, x' \in \mathcal{X}}}_Q \right) \quad (82)$$

$$\stackrel{(d)}{=} \frac{1}{|\mathcal{X} \setminus \text{Supp}(\hat{p})|!} \left(\prod_{i \geq 1} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \text{perm}(Q) \quad (83)$$

where in (a) the prefactor corrects for multiple-counting points in the orbit $\mathcal{S}_{\mathcal{X}}\hat{p}$ when we sum over $\sigma \in \mathcal{S}_{\mathcal{X}}$, (b) follows from (77), (c) follows from (10), and (d) follows since $\mathcal{F}_0 = |\mathcal{X} \setminus \text{Supp}(\hat{p})|$.

When the alphabet \mathcal{X} is countably infinite, then we group terms in the sum (76) by the support set $\text{Supp}(\hat{p}')$. Let $\hat{K} = |\text{Supp}(\hat{p})|$ (13) denote the empirical support set size and let $\Pi_{\hat{K}}(\mathcal{X})$ denote the set of all subsets of \mathcal{X} of size \hat{K} . Then

$$\mathbb{P}_p(\mathcal{F}) = \sum_{\hat{p}' \in \mathcal{S}_{\mathcal{X}}\hat{p}} \mathbb{P}_p(\hat{p}') \quad (84)$$

$$= \sum_{\mathcal{X}' \in \Pi_{\hat{K}}(\mathcal{X})} \sum_{\hat{p}' \in \mathcal{S}_{\mathcal{X}}\hat{p}: \text{Supp}(\hat{p}') = \mathcal{X}'} \mathbb{P}_p(\hat{p}') \quad (85)$$

$$\stackrel{(a)}{=} \sum_{\mathcal{X}' \in \Pi_{\hat{K}}(\mathcal{X})} \frac{|\mathcal{S}_{\mathcal{X}'}\hat{p}|}{|\mathcal{S}_{\mathcal{X}'}|} \sum_{\sigma \in \mathcal{S}_{\mathcal{X}'}} \mathbb{P}_p(\sigma\hat{p}') \quad (86)$$

$$\stackrel{(b)}{=} \left(\prod_{i \geq 1} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \sum_{\mathcal{X}' \in \Pi_{\hat{K}}(\mathcal{X})} \text{perm} \left(\left(p_x^{n\hat{p}_{x'}} \right)_{x, x' \in \mathcal{X}'} \right) \quad (87)$$

where in (a) the prefactor corrects for multiple-counting points in the orbit $\mathcal{S}_{\mathcal{X}'}\hat{p}'$ when we sum over $\sigma \in \mathcal{S}_{\mathcal{X}'}$ and (b) follows from (10) and the fact that \hat{p}' is supported on \mathcal{X}' (so there is no \mathcal{F}_0 term, unlike the case when \mathcal{X} is finite).

such assignments. Therefore, the number of equivalent classes Z is at least

$$Z \geq (\alpha^2 n)^n / (\alpha n)^{2\alpha n} = n^{(1-2\alpha)n} \cdot \alpha^{2n(1-\alpha)}.$$

The above inequality holds for any $\alpha > 0$. In particular, for any constant α , we have $Z \geq n^{(1-2\alpha-o(1))n}$, which implies the claimed bound.

Now we prove the upper bound. Consider the following procedure for generating an assignment to a $n \times n$ array.

1. Choose an integer $1 \leq m \leq n$ arbitrarily;
2. Choose m entries from the n^2 entries arbitrarily;
3. For each of the m chosen entries from last step, fill in a positive integer such that the m entries sum up to n .

It is easy to verify that for every equivalent class, some assignment in the class can be generated by the above procedure. On the other hand, the number of different assignments that can be generated is at most

$$\sum_{m=1}^n \binom{n^2}{m} \cdot \binom{n}{m} \leq 2^n \cdot n \cdot \max_{1 \leq m \leq n} \binom{n^2}{m} \leq 2^n \cdot n \cdot (e \cdot n)^n = n^{(1+o(1))n}.$$

This proves the upper bound.

Appendix D. The EM algorithm used to numerically solve the PML for small alphabets in Section 4

D.1. 1-D case

We treat p as the unknown parameter to be inferred and assume that p is sorted. Recall that for a fixed alphabet size, the PML distribution is:

$$p^* = \operatorname{argmax}_p \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \prod_{x \in \mathcal{X}} p_x^{n\hat{p}_{\sigma(x)}} \quad (92)$$

We treat permutation σ as the hidden data drawn uniformly on $\mathcal{S}_{\mathcal{X}}$, and the observed data \hat{p} drawn from the permuted distribution σp . Then if $p_{(t)}$ is our estimate for p after t iterations of the EM algorithm, the next value is:

$$p_{(t+1)} \sim \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} (\sigma \hat{p}) \prod_{x \in \mathcal{X}} p_{(t),x}^{n\hat{p}_{\sigma(x)}} \quad (93)$$

where $\sigma \hat{p}$ denotes the permuted empirical histogram and \sim denotes equality up to normalization. For the starting point we choose $p_{(0)} = \hat{p}$.

D.2. 2-D case

Analogously to the 1D case above, let $p_{(t)}$, $q_{(t)}$ be our estimates for p , q after t iterations of the EM algorithm. Then the next values are:

$$p_{(t+1)} \sim \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} (\sigma \hat{p}) w_{\sigma} \quad (94)$$

$$q_{(t+1)} \sim \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} (\sigma \hat{q}) w_{\sigma} \quad (95)$$

where

$$w_{\sigma} \triangleq \prod_{x \in \mathcal{X}} p_{(t),x}^{n \hat{p}_{\sigma(x)}} q_{(t),x}^{m \hat{q}_{\sigma(x)}} \quad (96)$$

and \sim denotes equality up to normalization. For the starting point we choose $p_{(0)} = \hat{p}$, $q_{(0)} = \hat{q}$.

Appendix E. Proof of Theorem 8 (D -dimensional PML distribution on a binary alphabet)

The D -dimensional PML distributions $\mathbf{p}^* \triangleq (p^{(d)*})_{d=1}^D$ are (see Appendix H for a derivation):

$$\mathbf{p}^* \triangleq \underset{\mathbf{p} \in \mathcal{P}}{\operatorname{argmax}} \operatorname{perm} \left(\left(\prod_{d=1}^D (p_x^{(d)})^{n_d \hat{p}_{x'}} \right)_{x, x' \in \{1, 2\}} \right) \quad (97)$$

$$= \underset{\mathbf{p} \in \mathcal{P}}{\operatorname{argmax}} \left(\prod_{d=1}^D (p_1^{(d)})^{n_d \hat{p}_1^{(d)}} (1 - p_1^{(d)})^{n_d (1 - \hat{p}_1^{(d)})} + \prod_{d=1}^D (p_1^{(d)})^{n_d (1 - \hat{p}_1^{(d)})} (1 - p_1^{(d)})^{n_d \hat{p}_1^{(d)}} \right) \quad (98)$$

where \mathcal{P} is a collection of D -tuples of distributions on the binary alphabet $\mathcal{X} = \{1, 2\}$. To lighten notation, let

$$r_d \triangleq p_1^{(d)} \quad (99)$$

and

$$\Delta_d \triangleq n_d \left(\hat{p}_1^{(d)} - \frac{1}{2} \right) \quad (100)$$

Then $\mathbf{p}^* = ((r_d^*, 1 - r_d^*))_{d=1}^D$, where

$$(r_d^*)_{d=1}^D = \underset{(r_d)_{d=1}^D \in [0, 1]^D}{\operatorname{argmax}} \underbrace{\left(\prod_{d=1}^D r_d^{n_d/2 + \Delta_d} (1 - r_d)^{n_d/2 - \Delta_d} + \prod_{d=1}^D r_d^{n_d/2 - \Delta_d} (1 - r_d)^{n_d/2 + \Delta_d} \right)}_{-U(r_1, \dots, r_d)} \quad (101)$$

$$= \underset{(r_d)_{d=1}^D \in [0, 1]^D}{\operatorname{argmin}} U(r_1, \dots, r_d) \quad (102)$$

If $r_d^* = 1/2$ for all d , then $U(r_1, \dots, r_d)$ must have positive definite Hessian H_U at $(r_d)_{d=1}^D = (1/2)_{d=1}^D$. Let's compute the Hessian at $r_d = 1/2$ for all d , letting $[D] \triangleq \{1, \dots, D\}$:

$$H_U^* \triangleq H_U|_{(r_d)_{d=1}^D = (1/2)_{d=1}^D} = \left(\frac{\partial^2 U}{\partial r_d \partial r_{d'}} \right)_{d, d' \in [D]} \Big|_{(r_d)_{d=1}^D = (1/2)_{d=1}^D} \quad (103)$$

$$= c \left((\Delta_d \Delta_{d'})_{d, d' \in [D]} - \text{diag} \left(\left(\frac{n_d}{4} \right)_{d \in [D]} \right) \right) \quad (104)$$

where $c = 2^{5 - \sum_{d=1}^D n_d}$ and in the second term in (104) $\text{diag}(\cdot)$ denotes the diagonal matrix with entry n_d in the d -th row and d -th column.

The matrix H_U^* is positive definite when $\Delta_d = 0$ for all d with eigenvalues $(-cn_d/4)_{d=1}^D$. Since the eigenvalues of H_U^* vary continuously in $(\Delta_d)_{d=1}^D$, we conclude that H_U^* is positive definite when $(\Delta_d)_{d=1}^D$ is in the neighborhood of $(0)_{d=1}^D$ determined by:

$$\det(H) = c' \left(-1 + \sum_{d=1}^D \frac{4\Delta_d^2}{n_d} \right) > 0 \quad (105)$$

$$\begin{aligned} & \Downarrow \\ & \sum_{d=1}^D 4n_d \left(\hat{p}_1^{(d)} - \frac{1}{2} \right)^2 > 1 \end{aligned} \quad (106)$$

where $c' = c 2^{2D} \prod_{d=1}^D n_d$. Thus H_U^* is not positive definite at $(r_d)_{d=1}^D = (1/2)_{d=1}^D$ when the empirical distribution components $(\hat{p}_1^{(d)})_{d=1}^D$ lie outside the ellipse determined by (106), so the D PML distributions $(p^{(d*)})_{d=1}^D$ are not uniform in this case.

Appendix F. Proof of permanent lower bound identity (35)

Given distribution p supported on set \mathcal{X} , denote by $\mathcal{X}_u(p)$ a level set of p :

$$\mathcal{X}_u(p) \triangleq \{x \in \mathcal{X} : p_x = u\} \quad (107)$$

Denote by $\mathcal{U}(p)$ the set of unique entries of p :

$$\mathcal{U}(p) \triangleq \{p_x : x \in \mathcal{X}\} \quad (108)$$

and denote by $\mathcal{A}(p)$ the partition of \mathcal{X} into level sets of p (equivalent to the definition (29)):

$$\mathcal{A}(p) \triangleq \{\mathcal{X}_u(p) : u \in \mathcal{U}(p)\} \quad (109)$$

It is convenient to express the group $\mathcal{S}_{\mathcal{X}, p}$ (32) as isomorphic to:

$$\mathcal{S}_{\mathcal{X}, p} \cong \prod_{\alpha \in \mathcal{A}(p)} \mathcal{S}_\alpha = \prod_{u \in \mathcal{U}(p)} \mathcal{S}_{\mathcal{X}_u} \quad (110)$$

We write

$$V(p) = \log(\text{perm}(Q)) \geq \bar{V}(p) \quad (111)$$

$$= \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X}, p}} \prod_{x \in \mathcal{X}} p_x^{n \hat{p}_{\sigma(x)}} \right) \quad (112)$$

$$\stackrel{(a)}{=} \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X}, p}} \prod_{x \in \mathcal{X}} p_{\sigma(x)}^{n \hat{p}_x} \right) \quad (113)$$

$$\stackrel{(b)}{=} \log \left(\sum_{\sigma \in \times_{u \in \mathcal{U}(p)} \mathcal{S}_{\mathcal{X}_u}} \prod_{x \in \mathcal{X}} p_{\sigma(x)}^{n \hat{p}_x} \right) \quad (114)$$

$$\stackrel{(c)}{=} \log \left(\prod_{u \in \mathcal{U}(p)} \sum_{\sigma \in \mathcal{S}_{\mathcal{X}_u}} \prod_{x \in \mathcal{X}_u} u^{n \hat{p}_x} \right) \quad (115)$$

$$\stackrel{(d)}{=} \log \left(\prod_{u \in \mathcal{U}(p)} |\mathcal{S}_{\mathcal{X}_u}| u^{n \sum_{x \in \mathcal{X}_u} \hat{p}_x} \right) \quad (116)$$

$$= \sum_{u \in \mathcal{U}(p)} \left(\log(|\mathcal{S}_{\mathcal{X}_u}|) + n \left(\sum_{x \in \mathcal{X}_u} \hat{p}_x \right) \log(u) \right) \quad (117)$$

$$= \sum_{u \in \mathcal{U}(p)} \log(|\mathcal{S}_{\mathcal{X}_u}|) + n \sum_{u \in \mathcal{U}(p)} \sum_{x \in \mathcal{X}_u} \hat{p}_x \log(u) \quad (118)$$

$$\stackrel{(e)}{=} \sum_{u \in \mathcal{U}(p)} \log(|\mathcal{X}_u|!) + n \sum_{x \in \mathcal{X}} \hat{p}_x \log(p_x) \quad (119)$$

$$= \sum_{u \in \mathcal{U}(p)} \log(|\mathcal{X}_u|!) + n \sum_{x \in \mathcal{X}} \hat{p}_x \left(\log \left(\frac{p_x}{\hat{p}_x} \right) + \log(\hat{p}_x) \right) \quad (120)$$

$$\stackrel{(f)}{=} \sum_{\alpha \in \mathcal{A}(p)} \log(|\alpha|!) - n(D(\hat{p}||p) + H(\hat{p})) \quad (121)$$

where in (a) we changed order of summation over the group $\mathcal{S}_{\mathcal{X}}$, in (b) we summed over the isomorphic group (110), in (c) we used the product structure (110) of $\mathcal{S}_{\mathcal{X}, p}$ to rewrite the sum as a product of sums and used the fact that $p_x = u$ for all $x \in \mathcal{X}_u$ (107). In (d) we used the fact that the summands are independent of σ and moved the sum into the exponent. In (e) we used $|\mathcal{S}_{\alpha}| = |\alpha|!$ and used the fact that $u = p_x$ for all $x \in \mathcal{X}_u$. In (f) we used $\alpha \in \mathcal{A}(p) \Leftrightarrow \exists u \in \mathcal{U}(p) : \mathcal{X}_u = \alpha$.

Appendix G. Proof of Theorem 9 (properties of the approximate PML distribution's level set partition $\bar{\mathcal{A}}^*$)

G.1. Proof of iso-clumping property (41)

Suppose $\bar{\mathcal{A}}^*$ does not satisfy property (41). Then there exist $x, y \in \mathcal{X}$ such that $\hat{p}_x = \hat{p}_y$, but $\alpha(x) \neq \alpha(y)$. We show that the partition obtained by either “reassigning” y to the

same level set as x or reassigning x to the same level set as y achieves a higher value of the log permanent lower bound $\bar{V}(\cdot)$ (38), contradicting the optimality of $\bar{\mathcal{A}}^*$ (40).

Let $m \triangleq n\hat{p}_x = n\hat{p}_y$, let $N_{\alpha(x)} \triangleq \sum_{x \in \alpha(x)} n\hat{p}_x$. Let V_δ denote for $\delta \in [-1, 1]$:

$$\begin{aligned} V_\delta \triangleq & \log((|\alpha(x)| + \delta)!) + (N_{\alpha(x)} + m\delta) \log\left(\frac{N_{\alpha(x)} + m\delta}{n(|\alpha(x)| + \delta)}\right) \\ & + \log((|\alpha(y)| - \delta)!) + (N_{\alpha(y)} - m\delta) \log\left(\frac{N_{\alpha(y)} - m\delta}{n(|\alpha(y)| - \delta)}\right) \end{aligned} \quad (122)$$

Then $V_0 = \bar{V}(\alpha(x)) + \bar{V}(\alpha(y))$ (39) corresponds to two terms in $\bar{V}(\bar{\mathcal{A}}^*)$ (38), and V_1 (resp. V_{-1}) corresponds to two terms in the value of the partition of \mathcal{X} obtained by reassigning symbol y to the same level set as x (resp. reassigning symbol x to the same level set as y). We can check that V_δ is strictly convex in $\delta \in [-1, 1]$. A strictly convex function attains its maximum on the boundaries of a convex set, so either $V_{-1} > V_0$ or $V_1 > V_0$, contradicting the optimality of partition $\bar{\mathcal{A}}^*$, thus establishing the iso-clumping property.

G.2. Proof of convexity property (42)

Suppose $\bar{\mathcal{A}}^*$ does not satisfy property (42). Then there exist $x_1, x_2, x_3 \in \mathcal{X}$ such that $\hat{p}_{x_1} < \hat{p}_{x_2} < \hat{p}_{x_3}$, but $\alpha(x_1) = \alpha(x_3) \neq \alpha(x_2)$. For all $x \in \mathcal{X}$, let $N_{\alpha(x)} \triangleq \sum_{x \in \alpha(x)} n\hat{p}_x$. If $\frac{N_{\alpha(x_1)}}{|\alpha(x_1)|} \geq \frac{N_{\alpha(x_2)}}{|\alpha(x_2)|}$, then let $x = x_2$ and $y = x_1$; otherwise let $x = x_3$, $y = x_2$. Let $\Delta \triangleq n(\hat{p}_x - \hat{p}_y)$. Then $\Delta > 0$ and

$$\frac{N_{\alpha(y)}}{|\alpha(y)|} \geq \frac{N_{\alpha(x)}}{|\alpha(x)|}. \quad (123)$$

Let \mathcal{A}' denote the partition of \mathcal{X} obtained by swapping the set memberships of x and y in $\bar{\mathcal{A}}^*$. We show that $\bar{V}(\mathcal{A}') > \bar{V}(\bar{\mathcal{A}}^*)$, contradicting the optimality of $\bar{\mathcal{A}}^*$ (\mathcal{A}' might not satisfy the iso-clumping property (41)).

Since \mathcal{A}' is obtained from $\bar{\mathcal{A}}^*$ by swapping the level set memberships of x and y , the elements of \mathcal{A}' and $\bar{\mathcal{A}}^*$ have the same sizes, so $\sum_{\alpha \in \mathcal{A}'} \log(|\alpha|!) = \sum_{\alpha \in \bar{\mathcal{A}}^*} \log(|\alpha|!)$, so using (38) and (39) we write:

$$\begin{aligned} \bar{V}(\mathcal{A}') - \bar{V}(\bar{\mathcal{A}}^*) = & (N_{\alpha(x)} - \Delta) \log\left(\frac{N_{\alpha(x)} - \Delta}{n|\alpha(x)|}\right) + (N_{\alpha(y)} + \Delta) \log\left(\frac{N_{\alpha(y)} + \Delta}{n|\alpha(x)|}\right) \\ & - N_{\alpha(x)} \log\left(\frac{N_{\alpha(x)}}{n|\alpha(x)|}\right) - N_{\alpha(y)} \log\left(\frac{N_{\alpha(y)}}{n|\alpha(y)|}\right) \end{aligned} \quad (124)$$

$$= \Delta \log\left(\frac{N_{\alpha(y)} |\alpha(x)|}{|\alpha(y)| N_{\alpha(x)}}\right) + f(N_{\alpha(x)}, \Delta) + f(N_{\alpha(y)}, -\Delta) \quad (125)$$

$$\stackrel{(a)}{>} 0 \quad (126)$$

where $f(N, \Delta) \triangleq \Delta + (N - \Delta) \log\left(\frac{N - \Delta}{N}\right)$ with $f(N, N) \triangleq \lim_{\Delta \rightarrow N} f(N, \Delta) = N$. (a) follows because we can show that $f(N, \Delta) \geq 0$ for $\Delta \leq N$ with strict inequality iff $\Delta \neq 0$, so the last two terms in (125) are positive since $\Delta > 0$. The first term in (125) is non-negative since $\Delta > 0$ and (123). Thus $\bar{V}(\mathcal{A}') > \bar{V}(\bar{\mathcal{A}}^*)$, contradicting the optimality of partition $\bar{\mathcal{A}}^*$, thus establishing the convexity property.

Appendix H. Our approximate PML algorithm: the case of multiple distributions

We generalize the discussion of Section 3.2 to D -dimensional fingerprints. In this case we are not able to optimize a D -dimensional analogue of our objective function \bar{V} (31) and will instead settle for an approximation of our already-approximate PML distribution. We use a greedy heuristic to approximate the optimal partition $\bar{\mathcal{A}}^*$ (40) by iteratively, greedily merging partition elements (enlarging the second term in a D -dimensional analogue of (35) until the first term becomes too large) (see Section H.4). Even this doubly-approximate algorithm turns out to be too slow to run on examples of practical interest, so we introduce a further heuristic to speed things up (only trying to merge partition elements that are close in Euclidean distance), which comes at a cost of adding some tunable knobs to the algorithm.

H.1. Some notation

Suppose we have D distributions $(p^{(d)})_{d=1}^D$ on the same alphabet \mathcal{X} and for each d draw n_d samples i.i.d. from $p^{(d)}$ with empirical distribution $\hat{p}^{(d)}$. Let bold symbols denote D -component collections:

$$\mathbf{p} \triangleq (p^{(d)})_{d=1}^D \quad \mathbf{n} \triangleq (n_d)_{d=1}^D \quad \hat{\mathbf{p}} \triangleq (\hat{p}^{(d)})_{d=1}^D \quad \mathbf{n}\hat{\mathbf{p}} \triangleq (n_d \hat{p}^{(d)})_{d=1}^D \quad (127)$$

We assume that the D distributions \mathbf{p} satisfy:

$$\sum_{d=1}^D p_x^{(d)} > 0 \quad \forall x \in \mathcal{X} \quad (128)$$

This condition ensures that there are no symbols x such that $p_x^{(d)} = 0$ for all d , which simplifies the expressions below. Denote by $\hat{\mathcal{X}}$ the empirical support set:

$$\hat{\mathcal{X}} \triangleq \bigcup_{d=1}^D \text{Supp}(\hat{p}^{(d)}) = \{x \in \mathcal{X} : \sum_{d=1}^D \hat{p}_x^{(d)} > 0\} \quad (129)$$

Finally let $K \triangleq |\mathcal{X}|$ and $\hat{K} \triangleq |\hat{\mathcal{X}}|$.

The probability to draw empirical distributions $\hat{\mathbf{p}}$ under distributions \mathbf{p} is (using the independence of the D empirical distributions and (10)):

$$\mathbb{P}_{\mathbf{p}}(\hat{\mathbf{p}}) = \left(\prod_{d=1}^D \binom{n_d}{n_d \hat{p}^{(d)}} \right) \prod_{x \in \mathcal{X}} \prod_{d=1}^D (p_x^{(d)})^{n_d \hat{p}_x^{(d)}} \quad (130)$$

where the first product is over multinomial coefficients in the notation of (10).

Hereon denote by $\mathcal{F} = \mathcal{F}(\hat{p}^{(1)}, \dots, \hat{p}^{(D)})$ the D -dimensional fingerprint (Raghunathan et al., 2017) indexed by vector $\mathbf{i} = (i_d)_{d=1}^D \in \mathbb{N}^D$ (which we write as satisfying $\mathbf{i} \geq \mathbf{0} \Leftrightarrow i_d \geq 0 \forall d$, where $\mathbf{0} \triangleq (0)_{d=1}^D$):

$$\mathcal{F} = \mathcal{F}(\hat{\mathbf{p}}) = (\mathcal{F}_{\mathbf{i}})_{\mathbf{i} \geq \mathbf{0}} \triangleq \left(|\{x \in \mathcal{X} : (n_d \hat{p}_x^{(d)})_{d=1}^D = \mathbf{i}\}| \right)_{\mathbf{i} \geq \mathbf{0}} \quad (131)$$

Note that $\mathcal{F}_0 = |\mathcal{X} \setminus \hat{\mathcal{X}}| = K - \hat{K}$ counts the number of symbols “unseen” by any of the D samples, and is thus unknown if the support set size K is unknown. Then the probability to draw the D -dimensional fingerprint \mathcal{F} under \mathbf{p} is (obtained analogously to the $D = 1$ case in Appendix A):

$$\mathbb{P}_{\mathbf{p}}(\mathcal{F}) = \left(\prod_{\mathbf{i} \geq \mathbf{0}} \frac{1}{\mathcal{F}_{\mathbf{i}}!} \right) \left(\prod_{d=1}^D \binom{n_d}{n_d \hat{p}^{(d)}} \right) \sum_{\sigma \in \mathcal{S}_{\mathcal{X}}} \prod_{x \in \mathcal{X}} \prod_{d=1}^D (p_x^{(d)})^{n_d \hat{p}_{\sigma(x)}^{(d)}} \quad (132)$$

$$= \frac{1}{(K - \hat{K})!} \left(\prod_{\mathbf{i} \geq \mathbf{0}, \mathbf{i} \neq \mathbf{0}} \frac{1}{\mathcal{F}_{\mathbf{i}}!} \right) \left(\prod_{d=1}^D \binom{n_d}{n_d \hat{p}^{(d)}} \right) \text{perm} \left(\underbrace{\left(\prod_{d=1}^D (p_x^{(d)})^{n_d \hat{p}_{x'}^{(d)}} \right)}_{Q} \right)_{x, x' \in \mathcal{X}} \quad (133)$$

where in (133) the first product is over all D -dimensional indices \mathbf{i} such that $\min_d i_d \geq 0$ and $\max_d i_d > 0$.

H.2. The PML distributions in D dimensions

For a given collection \mathcal{P} of D -tuples of distributions, the D -dimensional PML distributions \mathbf{p}^* are:

$$\mathbf{p}^* \triangleq (p^{(d)*})_{d=1}^D \triangleq \underset{\mathbf{p} \in \mathcal{P}}{\text{argmax}} \mathbb{P}_{\mathbf{p}}(\mathcal{F}(\hat{\mathbf{p}})) \quad (134)$$

$$= \underset{\mathbf{p} \in \mathcal{P}}{\text{argmax}} \frac{\text{perm}(Q)}{(K - \hat{K})!} \quad (135)$$

where in (135) we discarded all \mathbf{p} -independent factors of $\mathbb{P}_{\mathbf{p}}(\mathcal{F})$ (133). $\mathcal{F}_0 = K - \hat{K} = |\mathcal{X} \setminus \hat{\mathcal{X}}|$ depends on p through its support set size. Note that $\mathbb{P}_{\mathbf{p}}(\mathcal{F})$ is invariant under relabeling of the components of \mathbf{p} , so we can choose \mathbf{p}^* to be non-increasing in the same ordering as we choose for the support set \mathcal{X} . Note that the set \mathcal{P} is not necessarily the same as the set of all discrete distributions.

Note that finding the D -dimensional PML distributions (135) is not equivalent to finding D 1-dimensional PML distributions since $\mathbb{P}_{\mathbf{p}}(\mathcal{F})$ (133) does not factor into D pieces.

If the collection of D -tuples of distributions \mathcal{P} includes distributions with different support set sizes, then we can estimate the support set size by breaking up the optimization in (135) into two steps:

$$K^* \triangleq \underset{K}{\text{argmax}} \left(\frac{1}{(K - \hat{K})!} \max_{\mathbf{p} \in \mathcal{P}_K} \text{perm}(Q) \right) \quad (136)$$

whenever the max over K exists, where $\mathcal{P}_K \triangleq \{\mathbf{p} \in \mathcal{P} : |\bigcup_{d=1}^D \text{Supp}(p^{(d)})| = K\}$.

The case $D \geq 2$ is more complicated than the case $D = 1$ because the D support sets $(\text{Supp}(p^{(d)}))_{d=1}^D$ can overlap partially. If \mathcal{X} is the joint support set – that is, $\mathcal{X} = \bigcup_{d=1}^D \text{Supp}(p^{(d)})$ – then we can associate with each symbol $x \in \mathcal{X}$ an indicator function $\phi_d(x) = \mathbb{1}(x \in \text{Supp}(p^{(d)}))$, so there are $2^D - 1$ possible values $(\phi_d(x))_{d=1}^D$ for each $x \in \mathcal{X}$ ⁹.

9. 1 less than 2^D since $\phi_d(x) = 0 \forall d$ implies $x \notin \mathcal{X}$.

We could then be more ambitious and estimate not just the joint support set size $K = |\mathcal{X}|$, but also the optimal value $(\phi_d(x))_{d=1}^D$ for all symbols. Our proposed approximate PML approximation scheme for $D \geq 2$ is less ambitious and only supports estimation of $|\mathcal{X}|$, but not of the finer structure of the joint support set \mathcal{X} in terms of D partially overlapping support sets $(\text{Supp}(p^{(d)}))_{d=1}^D$.

H.3. Approximate PML in D dimensions: the case of known support set size

In this Section we assume the support set size $K = |\mathcal{X}|$ is known. Analogously to the notation in the $D = 1$ case of Section 4.3 and the proof of the permanent lower bound for $D = 1$ in Appendix F, define the level sets $\mathcal{X}_{\mathbf{u}}(\mathbf{p})$ of \mathbf{p} indexed by D -dimensional real vector $\mathbf{u} = (u_d)_{d=1}^D$:

$$\mathcal{X}_{\mathbf{u}}(\mathbf{p}) \triangleq \{x \in \mathcal{X} : (p_x^{(d)})_{d=1}^D = \mathbf{u}\} \quad (137)$$

the unique D -tuple entries of \mathbf{p} by $\mathcal{U}(\mathbf{p})$:

$$\mathcal{U}(\mathbf{p}) \triangleq \{(p_x^{(d)})_{d=1}^D : x \in \mathcal{X}\} \quad (138)$$

and the partition $\mathcal{A}(\mathbf{p})$ of \mathcal{X} into level sets of \mathbf{p} :

$$\mathcal{A}(\mathbf{p}) \triangleq \{\mathcal{X}_{\mathbf{u}}(\mathbf{p}) : \mathbf{u} \in \mathcal{U}(\mathbf{p})\} \quad (139)$$

We assume $K = |\mathcal{X}|$ is known, so \mathbf{p}^* (135) is a maximizer of

$$V(\mathbf{p}) \triangleq \log(\text{perm}(Q)) \quad (140)$$

As for the $D = 1$ case, we do not know how to computationally efficiently maximize $V(\mathbf{p})$ over \mathbf{p} , so we settle for maximizing a lower bound $\bar{V}(\mathbf{p})$, analogous to (31) and motivated by our empirical observations for the case $D = 2$ in Section 4.2:

$$V(\mathbf{p}) \geq \bar{V}(\mathbf{p}) \triangleq \log \left(\sum_{\sigma \in \mathcal{S}_{\mathcal{X}, \mathbf{p}}} \prod_{x \in \mathcal{X}} \prod_{d=1}^D (p_x^{(d)})^{n_d \hat{p}_{\sigma(x)}^{(d)}} \right) \quad (141)$$

$$= - \sum_{d=1}^D n_d \left(D(\hat{p}^{(d)} || p^{(d)}) + H(\hat{p}^{(d)}) \right) + \sum_{\alpha \in \mathcal{A}(\mathbf{p})} \log(|\alpha|!) \quad (142)$$

where $\mathcal{S}_{\mathcal{X}, \mathbf{p}}$ is analogous to (32) – a subgroup of $\mathcal{S}_{\mathcal{X}}$ consisting of all permutations that exchange only those alphabet symbols that are in the same level set of \mathbf{p} :

$$\mathcal{S}_{\mathcal{X}, \mathbf{p}} \triangleq \{\sigma \in \mathcal{S}_{\mathcal{X}} : \sigma \mathbf{p} = (\sigma p^{(d)})_{d=1}^D = \mathbf{p}\} \quad (143)$$

(142) follows from (141) by an analogous argument to the one given in Appendix F for the case $D = 1$.

As in the $D = 1$ case, the intuition is that we approximate the permanent by summing over only those permutations that contribute “a lot” to the value, and then maximize this lower bound $\bar{V}(\mathbf{p})$ over \mathbf{p} . Our approximate PML distributions are the D -tuple:

$$\bar{\mathbf{p}}^* = (\bar{p}^{(d)*})_{d=1}^D \triangleq \underset{\mathbf{p} \in \mathcal{P}}{\text{argmax}} \bar{V}(\mathbf{p}) \quad (144)$$

As in the 1-D case, the first term in (142) encourages $\bar{\mathbf{p}}^*$ to clump many symbols together, while the second term encourages $\bar{\mathbf{p}}^*$ to be similar to $\hat{\mathbf{p}}$, dominating as $n \rightarrow \infty$.

Note that optimizing $\bar{V}(\mathbf{p})$ over \mathbf{p} is not equivalent to solving D independent optimization problems because the D distributions “interact” through the second summation in expression (142) (the summation over $\mathcal{A}(\mathbf{p})$).

As in the 1-D case, we can show that for all d the approximate PML distribution $\bar{p}^{(d)*}$ (144) satisfies the averaging property (37). Therefore as in the 1-D case, $\bar{\mathbf{p}}^*$ is determined by its partition of \mathcal{X} into level sets. Thus we seek to maximize the quantity

$$\bar{V}(\mathcal{A}) \triangleq \sum_{\alpha \in \mathcal{A}} \bar{V}(\alpha) \quad (145)$$

where for $\alpha \subset \mathcal{X}$

$$\bar{V}(\alpha) \triangleq \log(|\alpha|!) + |\alpha| \sum_{d=1}^D n_d \hat{p}_\alpha^{(d)} \log(\hat{p}_\alpha^{(d)}) \quad (146)$$

The optimal partition is

$$\bar{\mathcal{A}}^* \triangleq \operatorname{argmax}_{\mathcal{A}: \text{partition of } \mathcal{X}} \bar{V}(\mathcal{A}) \quad (147)$$

Finally $\bar{\mathbf{p}}^*$ is obtained by averaging the empirical distributions $\hat{\mathbf{p}}^*$ within each partition element $\alpha \in \bar{\mathcal{A}}^*$.

As in the 1-dimensional case, the approximate PML partition $\bar{\mathcal{A}}^*$ has the iso-clumping property (generalizing (41) to D dimensions):

$$(\hat{p}_x^{(d)})_{d=1}^D = (\hat{p}_{x'}^{(d)})_{d=1}^D \Rightarrow \alpha(x) = \alpha(x') \quad (148)$$

for all $x, x' \in \mathcal{X}$, where $\alpha(x), \alpha(x') \in \bar{\mathcal{A}}^*$ are the partition elements containing x, x' , respectively. This can be shown by adapting the proof for the 1-D case in Appendix G.1.

We do not know if $\bar{\mathcal{A}}^*$ satisfies a convexity property analogous to the 1-D case (42).

H.4. A greedy heuristic to approximately maximize the permanent for $D \geq 2$ distributions

Unlike the $D = 1$ case, we do not solve the optimization problem (144) exactly for $D \geq 2$. The difficulty is that there is no natural ordering on \mathbb{N}^D for $D \geq 2$, so we can no longer define a structured collection of subsets of \mathcal{X} from which to build the approximate PML partition $\bar{\mathcal{A}}^*$ (147) \mathcal{X} , so we do not offer a dynamic programming algorithm like the one described in Section 4.3. Our search for an optimal partition of \mathcal{X} is similar in flavor to a clustering problem, so the added difficulty in $D \geq 2$ dimensions makes some sense. For example, k-means clustering in 1 dimension is exactly solvable by a dynamic programming algorithm (Wang and Song, 2011) similar to the one we present, but no such algorithm is available for dimension greater than 1 (indeed k-means is NP-hard for dimension greater than 1).

We settle for a greedy heuristic to approximately maximize $\bar{V}(\mathcal{A})$, and call the resulting approximate optimizer $\bar{\bar{\mathcal{A}}}^*$:

$$\bar{\bar{\mathcal{A}}}^* \approx \operatorname{argmax}_{\mathcal{A}: \text{partition of } \mathcal{X}} \bar{V}(\mathcal{A}) \quad (149)$$

the double bar signaling the two levels of approximation (an approximation to the approximate PML partition $\bar{\mathcal{A}}^*$ (144)). The doubly-approximate PML distribution $\bar{\mathbf{p}}^*$ is obtained by averaging the empirical distributions $\hat{\mathbf{p}}^*$ within each partition element $\alpha \in \bar{\mathcal{A}}^*$.

The greedy heuristic used to compute $\bar{\mathcal{A}}^*$ is given in algorithm (1). In words, we start with initial partition $\mathcal{A}(\hat{\mathbf{p}})$ of \mathcal{X} into the level sets of the empirical distributions $\hat{\mathbf{p}}$ (we clump equally empirically frequent symbols into the same level sets since we know that $\bar{\mathcal{A}}^*$ satisfies the iso-clumping property; see Section H.3) and then iteratively merge partition elements until it is no longer profitable to do so in terms of the objective function \bar{V} increasing. At each step we merge two partition elements that most boost the value of the objective function.

Algorithm 1 Approximation to level set partition of the approximate D -dimensional PML distribution

- 1: **function** $\bar{\mathcal{A}}^*(\hat{\mathbf{p}}, \mathbf{n}, |\mathcal{X}|)$ ▷ Input: D empirical distributions, D sample sizes, and assumed support set size $|\mathcal{X}|$
 - 2: $\mathcal{A} \leftarrow \mathcal{A}(\hat{\mathbf{p}})$ ▷ Initial partition into level sets of empirical distribution $\hat{\mathbf{p}}$
 - 3: **while** $\max_{\alpha_1, \alpha_2 \in \mathcal{A}: \alpha_1 \neq \alpha_2} (\bar{V}(\alpha_1 \cup \alpha_2) - (\bar{V}(\alpha_1) + \bar{V}(\alpha_2))) > 0$ **do**
 - 4: $(\alpha_1^*, \alpha_2^*) \leftarrow \operatorname{argmax}_{\alpha_1, \alpha_2 \in \mathcal{A}: \alpha_1 \neq \alpha_2} (\bar{V}(\alpha_1 \cup \alpha_2) - (\bar{V}(\alpha_1) + \bar{V}(\alpha_2)))$
 - 5: $\mathcal{A} \leftarrow (\mathcal{A} \cup \{\alpha_1^* \cup \alpha_2^*\}) \setminus \{\alpha_1^*, \alpha_2^*\}$ ▷ Merge two level sets into one
 - 6: **end while**
 - 7: **return** \mathcal{A} ▷ Output: $\bar{\mathcal{A}}^*(\hat{\mathbf{p}})$
 - 8: **end function**
-

The inputs $\hat{\mathbf{p}}$ and \mathbf{n} in algorithm (1) are used when evaluating $\bar{V}(\alpha)$ (146) for $\alpha \subset \mathcal{X}$. The input of the support set size $|\mathcal{X}|$ is used in line (2) to initially assign the $|\mathcal{X}| - |\mathcal{X}(\hat{\mathbf{p}})|$ unobserved symbols to a partition element. The loop in algorithm (1) always terminates since we only merge distinct partition elements of the finite support set \mathcal{X} . We can implement algorithm (1) efficiently by observing that for $\alpha_1, \alpha_2 \subset \mathcal{X}$ $\alpha_1 \neq \alpha_2$:

$$\hat{p}_{\alpha_1 \cup \alpha_2} = \frac{|\alpha_1| \hat{p}_{\alpha_1} + |\alpha_2| \hat{p}_{\alpha_2}}{|\alpha_1| + |\alpha_2|} \quad (150)$$

Relation (150) lets us evaluate $\bar{V}(\alpha_1 \cup \alpha_2)$ without summing over all points in $\alpha_1 \cup \alpha_2$, so we can reuse the work we did when we previously computed $\bar{V}(\alpha_1)$, $\bar{V}(\alpha_2)$.

We obtain the doubly-approximate set of D distributions $\bar{\mathbf{p}}^*$ by setting for all $x \in \mathcal{X}$, $d \in \{1, \dots, D\}$

$$\bar{p}_x^{(d)*} = \hat{p}_{\alpha(x)}^{(d)} \quad (151)$$

where $\alpha(x) \in \bar{\mathcal{A}}^*$ is the partition element containing x , and $\hat{p}_\alpha^{(d)}$ (36) is the average of $\hat{p}^{(d)}$ over $\alpha \subset \mathcal{X}$.

We know that the partition $\bar{\mathcal{A}}^*$ computed by algorithm (1) is suboptimal (that is, does not maximize \bar{V} (147)) because it only tries to merge pairs of existing partition elements rather than triplets or larger collections; it is possible to construct an example where the optimal partition consists of a single set (all of \mathcal{X}) but no pair of partition elements of $\mathcal{A}(\hat{\mathbf{p}})$ is “worth” merging in the sense of satisfying the loop condition in line (3).

The running time of the heuristic (1) is $O(|\text{Supp}(\mathcal{F})|^2) = O(n)$, where $n = \sum_{d=1}^D n_d$. In line (2), we have $|\mathcal{A}(\hat{\mathbf{p}})| = |\text{Supp}(\mathcal{F})| = O(\sqrt{n})$. To evaluate the max over $\alpha_1, \alpha_2 \in \mathcal{A}$ for the first time in line (3) we consider $O(|\mathcal{A}(\hat{\mathbf{p}})|^2) = O(n)$ pairs of partition elements. On every subsequent iteration of the loop, we add one new partition element ($\alpha_1^* \cup \alpha_2^*$) to \mathcal{A} and remove two partition elements (α_1^* and α_2^*) from \mathcal{A} , so there are $O(|\mathcal{A}|) = O(\sqrt{n})$ new pairs of distinct partition elements to check in computing the max in line (3). The size of $|\mathcal{A}|$ decrements by one each time the loop is run, so the loop is run at most $O(|\mathcal{A}(\hat{\mathbf{p}})|) = O(\sqrt{n})$ times, so the overall complexity of algorithm 1 is $O(n) + O(\sqrt{n})O(\sqrt{n}) = O(n)$.

If the support set size $|\mathcal{X}|$ is unknown, then we vary the assumed support set size $|\mathcal{X}|$ to optimize $\bar{V}(\bar{\mathcal{A}}^*(\hat{\mathbf{p}}, \mathbf{n}, |\mathcal{X}|))$ using a bisection search on the range $[|\mathcal{X}(\hat{\mathbf{p}})|, \max_{d \in [1, D]} n_d^2]$, increasing the overall running time of the algorithm (optimizing over both the level set partition \mathcal{A} and the support set size) to $O(n \log(n))$. Numerical experiments show that $\bar{V}(\bar{\mathcal{A}}^*(\hat{\mathbf{p}}, \mathbf{n}, |\mathcal{X}|))$ is not unimodal in $|\mathcal{X}|$, so this heuristic might choose a suboptimal support set size.

The $O(n)$ worst-case running time of algorithm (1) turns out to be slow for examples of practical interest (in making the performance plots of Section 6), so we introduce another heuristic to speed things up at the cost of adding some tunable knobs to the algorithm. The heuristic is to avoid checking all pairs of partition elements $\alpha_1, \alpha_2 \in \mathcal{A}$, instead checking only the “promising” pairs. A pair $\alpha_1, \alpha_2 \in \mathcal{A}$ is called *promising* if the point $(\hat{p}_{\alpha_1}^{(d)})_{d=1}^D \in \mathbb{R}^D$ is one of the k nearest neighbors of the point $(\hat{p}_{\alpha_2}^{(d)})_{d=1}^D \in \mathbb{R}^D$ in Euclidean distance¹⁰. This heuristic is motivated by the observation that the level sets that get merged in the loop of algorithm (1) tend to correspond to nearby points in Euclidean distance. Then we modify algorithm (1) to first compute the set of k nearest neighbors (doable in time $O(|\text{Supp}(\mathcal{F})| \log(|\text{Supp}(\mathcal{F})|)) = O(\sqrt{n} \log(\sqrt{n})) = O(\sqrt{n} \log(n))$) and then to update the set of nearest neighbors after each loop execution. This updating can be done heuristically by taking the k nearest neighbors of $\hat{p}_{\alpha_1 \cup \alpha_2}^{(d)}$ to be the k nearest points chosen from among only the union of the k nearest neighbors of $(\hat{p}_{\alpha_1}^{(d)})_{d=1}^D$ and the k nearest neighbors of $(\hat{p}_{\alpha_2}^{(d)})_{d=1}^D$ (rather than the k nearest points chosen from among $(\hat{p}_{\alpha}^{(d)})_{d=1}^D$ for all $\alpha \in \mathcal{A}$). The loop is executed $O(\sqrt{n})$ times, so the overall running time of the algorithm is $O(\sqrt{n} \log(n) + k\sqrt{n})$. The time to compute the empirical histograms $\hat{\mathbf{p}}$ is $O(n)$, asymptotically in n dominating the runtime of our algorithm for large n , but tends to take less time in our use cases than running algorithm 1 modified with our k nearest neighbor heuristic. For the performance plots of Section 6 we used $k = 5$.

Figures 2 and 14 show the doubly-approximate PML level set partition $\bar{\mathcal{A}}^*$ (149) for $D = 2$ and $D = 3$, respectively.

Appendix I. Mutual and Lautum information

For the case of estimating the mutual information $I(X; Y)$, defined as

$$I(X; Y) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_{x,y} \log \left(\frac{p_{x,y}}{p_x p_y} \right), \quad (152)$$

10. The nearest neighbor relation is not symmetric, so α_1, α_2 being promising does not imply that α_2, α_1 is a promising pair.

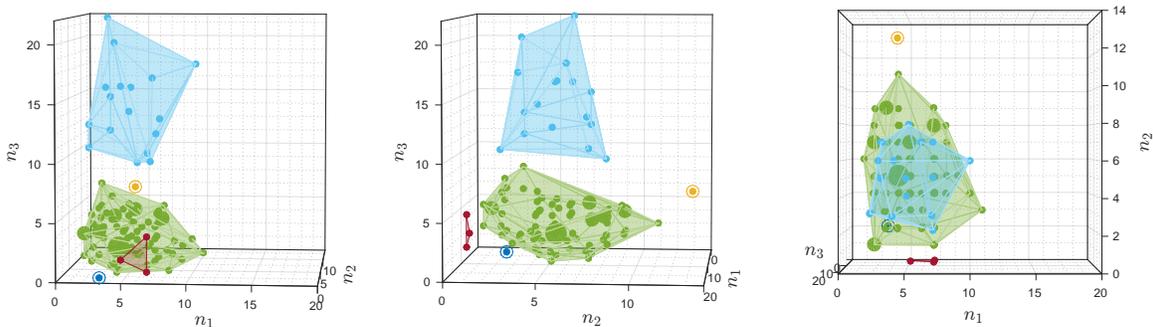


Figure 14: Computing the doubly-approximate PML distribution level set partition $\bar{\mathcal{A}}^*$ (149) by clumping entries of a 3-D fingerprint $\mathcal{F}_{i,j,k}$ using algorithm 1 modified with the 5-nearest neighbors heuristic as described in Section H.4, shown in three projections. The underlying distributions are $p^{(1)} = p^{(2)}$ uniform on $\{1, \dots, 100\}$ and $p^{(3)}$ is a mixture of two uniforms, with half its mass uniform on $\{1, \dots, 20\}$ and the other half uniform on $\{21, \dots, 100\}$. The sample size is 500 for all three empirical distributions $(\hat{p}^{(d)})_{d=1}^3$. Colored convex hulls correspond to symbols assigned to the same clump (partition element). The marker size of point (i, j, k) is roughly proportional to $\mathcal{F}_{i,j,k}$. We can see two large clumps, corresponding to symbols $\{1, \dots, 20\}$ and $\{21, \dots, 100\}$.

or lautum information (Palomar and Verdú, 2008), defined as

$$L(X; Y) = \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_x p_y \log \left(\frac{p_x p_y}{p_{x,y}} \right), \quad (153)$$

we have n samples from a joint distribution p on $\mathcal{X} \times \mathcal{Y}$ with joint histogram \hat{p} . Here $p_x = \sum_y p_{x,y}, p_y = \sum_x p_{x,y}$ are the marginal distributions of X and Y , respectively.

I.1. Product of permutation groups: one candidate for PML

It is clear that the mutual information and the lautum information are invariant to the product of two permutation groups $\mathcal{S}_{\mathcal{X}} \times \mathcal{S}_{\mathcal{Y}}$, acting on p as:

$$((\sigma_{\mathcal{X}}, \sigma_{\mathcal{Y}})p)_{x,y} \triangleq p_{\sigma_{\mathcal{X}}(x), \sigma_{\mathcal{Y}}(y)} \quad (154)$$

for all $(\sigma_{\mathcal{X}}, \sigma_{\mathcal{Y}}) \in \mathcal{S}_{\mathcal{X}} \times \mathcal{S}_{\mathcal{Y}}$. Let $(\mathcal{S}_{\mathcal{X}} \times \mathcal{S}_{\mathcal{Y}})p$ denote the orbit of p under the group action. Following the general profile maximum likelihood methodology for the product of these two permutation groups, the PML distribution p^* maximizes the following quantity over a set

of distribution \mathcal{P} on $\mathcal{X} \times \mathcal{Y}$:

$$\mathbb{P}_p((\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y})\hat{p}) = \sum_{\hat{p}' \in (\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y})\hat{p}} \mathbb{P}_p(\hat{p}') \quad (155)$$

$$\stackrel{(a)}{=} \frac{|(\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y})\hat{p}|}{|\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y}|} \sum_{(\sigma_\mathcal{X}, \sigma_\mathcal{Y}) \in \mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y}} \mathbb{P}_p((\sigma_\mathcal{X}, \sigma_\mathcal{Y})\hat{p}) \quad (156)$$

$$= \left(\prod_{i \geq 0} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \sum_{(\sigma_\mathcal{X}, \sigma_\mathcal{Y}) \in \mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y}} \prod_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p_{x,y}^{n\hat{p}_{\sigma_\mathcal{X}(\mathcal{X}^{(x)}, \sigma_\mathcal{Y}^{(y)})}} \quad (157)$$

$$= \left(\prod_{i \geq 0} \frac{1}{\mathcal{F}_i!} \right) \binom{n}{n\hat{p}} \text{perm} \left(\left(p_{x,y}^{n\hat{p}_{x',y'}} \right)_{(x,y),(x',y') \in \mathcal{X} \times \mathcal{Y}} \right), \quad (158)$$

where in (a) the prefactor corrects for multiple-counting points of the orbit $(\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y})\hat{p}$ when we sum over $(\sigma_\mathcal{X}, \sigma_\mathcal{Y}) \in (\mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y})$, and where $(\mathcal{F}_i)_{i \geq 0} = (|\{(x,y) \in \mathcal{X} \times \mathcal{Y} : n\hat{p}_{x,y} = i\}|)_{i \geq 0}$ is the fingerprint of the distribution when one treats $\mathcal{X} \times \mathcal{Y}$ as a single alphabet. Note that in this setting $(\mathcal{F}_i)_{i \geq 0}$ is not partially sufficient for estimating the mutual (lautum) information: the mutual (lautum) information cannot be written as a functional of the sorted probabilities of $p_{x,y}$.

In this setting, it follows from Theorem 6 that the partially sufficient statistic could be defined via the following equivalence class representation: for any empirical distribution $\hat{p}^{(1)} = (\hat{p}_{x,y}^{(1)})$, $\hat{p}^{(2)} = (\hat{p}_{x,y}^{(2)})$, we say $\hat{p}^{(1)}$ is equivalent to $\hat{p}^{(2)}$ if there exists an element $g \in \mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y}$ such that $g\hat{p}^{(1)} = \hat{p}^{(2)}$. The partially sufficient statistic is the corresponding equivalent classes.

In this PML formulation, we analyze the cardinality of the corresponding partially sufficient statistic below. This number appeared in (OEIS sequence A007716), and its precise asymptotics is obtained in the following theorem.

Theorem 12 *For any two empirical distributions $\hat{p}^{(1)} = (\hat{p}_{x,y}^{(1)})$, $\hat{p}^{(2)} = (\hat{p}_{x,y}^{(2)})$ with identical sample size n , we say $\hat{p}^{(1)}$ is equivalent to $\hat{p}^{(2)}$ if there exists an element $(\sigma_\mathcal{X}, \sigma_\mathcal{Y}) \in \mathcal{S}_\mathcal{X} \times \mathcal{S}_\mathcal{Y}$ such that $(\sigma_\mathcal{X}, \sigma_\mathcal{Y})\hat{p}^{(1)} = \hat{p}^{(2)}$, where $\mathcal{S}_\mathcal{X}, \mathcal{S}_\mathcal{Y}$ are the permutation groups on \mathcal{X}, \mathcal{Y} , respectively. Then, the number of equivalence classes is given by $e^{(1+o(1))n \log n}$.*

Theorem 12 shows that the cardinality of the partially sufficient statistic is no longer sub-exponential. This negative result immediately renders Theorem 3 not applicable, which leaves us with no guarantee on the performance of this type of PML on estimating mutual information or lautum information.

I.2. Three permutation groups: the other candidate for the PML

The mutual information $I(X; Y)$ can be rewritten as

$$I(X; Y) = H(P_X) + H(P_Y) - H(P_{XY}) \quad (159)$$

$$= \sum_{x \in \mathcal{X}} p_x \log \left(\frac{1}{p_x} \right) + \sum_{y \in \mathcal{Y}} p_y \log \left(\frac{1}{p_y} \right) - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p_{x,y} \log \left(\frac{1}{p_{x,y}} \right), \quad (160)$$

where $H(P_X) = \sum_x p_x \log \frac{1}{p_x}$ denotes the Shannon entropy of discrete distribution P_X .

Following the discussion in Section 3.1, we introduce

$$(\mathcal{F}_i^{XY})_{i \geq 0} = (|\{(x, y) \in \mathcal{X} \times \mathcal{Y} : n\hat{p}_{x,y} = i\}|)_{i \geq 0} \quad (161)$$

$$(\mathcal{F}_i^X)_{i \geq 0} = (|\{x \in \mathcal{X} : n\hat{p}_x = i\}|)_{i \geq 0} \quad (162)$$

$$(\mathcal{F}_i^Y)_{i \geq 0} = (|\{y \in \mathcal{Y} : n\hat{p}_y = i\}|)_{i \geq 0}, \quad (163)$$

and it is clear that $(\mathcal{F}_i^{XY})_{i \geq 0}$ is partially sufficient for $H(P_{XY})$, $(\mathcal{F}_i^X)_{i \geq 0}$ is partially sufficient for $H(P_X)$, and $(\mathcal{F}_i^Y)_{i \geq 0}$ is partially sufficient for $H(P_Y)$. Since $I(X; Y) = H(P_X) + H(P_Y) - H(P_{XY})$, it motivates the following PML definition:

$$p^* \triangleq \operatorname{argmax}_{p \in \mathcal{P}} \mathbb{P}_p(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^Y). \quad (164)$$

In other words, the PML distribution $(p_{x,y}^*)$ on $\mathcal{X} \times \mathcal{Y}$ maximizes the probability of observing the three joint fingerprints $(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^Y)$. This PML formulation, which only applies to mutual information but not lautum information, does not suffer from the super-exponential cardinality problem. Applying Theorem 10, it follows that the cardinality of the statistic is at most

$$\left(e^{\pi \sqrt{\frac{2n}{3}}} \right)^3 = e^{\pi \sqrt{6n}}, \quad (165)$$

which is still sub-exponential. However, it seems to be a non-trivial task to even approximately solve the PML formulation (164). Moreover, this PML formulation does not apply to lautum information.

Note that there may not exist a group G (a subgroup of the symmetric group $\mathcal{S}_{\mathcal{X} \times \mathcal{Y}}$) acting on the joint empirical distributions \hat{p} such that the set of distinct triplets $(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^Y)$ is in bijection with the set of distinct orbits under the action of G . For example, let \hat{p} , \hat{p}' , and \hat{p}'' be three empirical distributions on $\mathcal{X} = \{x_1, x_2\}$, $\mathcal{Y} = \{y_1, y_2, y_3\}$ with sample size $n = 1110$:

$$n\hat{p} = \begin{pmatrix} 1 & 4 & 100 \\ 2 & 3 & 1000 \end{pmatrix} \quad n\hat{p}' = \begin{pmatrix} 2 & 3 & 100 \\ 1 & 4 & 1000 \end{pmatrix} \quad n\hat{p}'' = \begin{pmatrix} 2 & 100 & 3 \\ 1 & 1000 & 4 \end{pmatrix} \quad (166)$$

where rows (columns) correspond to elements of \mathcal{X} (\mathcal{Y}). \hat{p} , \hat{p}' , and \hat{p}'' have the same \mathcal{X} - and \mathcal{Y} -marginal distributions ((105, 1005)/1110 and (3, 7, 1100)/1110, respectively) and the same set of entries, so their fingerprint triplets $(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^Y)$ are the same.

Suppose there exists a group $G \subset \mathcal{S}_{\mathcal{X} \times \mathcal{Y}}$ such that the orbit $G\hat{p}$ contains exactly those distributions with the same fingerprint triplet $(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^Y)$. Then $\hat{p}', \hat{p}'' \in G\hat{p}$. Let $\sigma \in G$, where $\hat{p}' = \sigma\hat{p}$. Then σ switches x_1 and x_2 iff $y \neq y_3$:

$$(\sigma p)_{x,y} = \begin{cases} p_{x_2,y} & : x = x_1, y \neq y_3 \\ p_{x_1,y} & : x = x_2, y \neq y_3 \\ p_{x,y} & : \text{otherwise} \end{cases} \quad (167)$$

On the other hand,

$$n(\sigma\hat{p}'') = \begin{pmatrix} 1 & 1000 & 3 \\ 2 & 100 & 4 \end{pmatrix} \quad (168)$$

has \mathcal{X} -marginal $(1004, 106)/1110$, and thus a different \mathcal{F}^X , so the action of σ does not preserve the triplet $(\mathcal{F}^{XY}, \mathcal{F}^X, \mathcal{F}^T)$, so group G does not exist.

Appendix J. Estimating the support set size of the approximate PML distribution

We first state our results in the notation introduced in Sections 4.5 and 4.6.

Theorem 13 *Let \bar{p}^* be the approximate PML distribution. Then $\mathcal{X}_{0:i^*} \in \mathcal{A}(\bar{p}^*)$ is a level set of \bar{p}^* , where*

$$i^* \triangleq \operatorname{argmax}_{i \in [1, F_+]} \left(\sup_{\mathcal{F}_0 \in \mathbb{N}} (-\log(\mathcal{F}_0!) + \bar{V}(\mathcal{X}_{0:i})) + \bar{V}_{i+1} \right) \quad (169)$$

For each value of $i \in [1, F_+]$ the supremum over \mathcal{F}_0 in (169) is finite and we compute it below.

Let $K_i \triangleq |\mathcal{X}_{1:i}|$ and $N_i \triangleq \sum_{x \in \mathcal{X}_{1:i}} n\hat{p}_x$.

If $i = 1$ and $\mathcal{F}_1 \geq 1$, then $N_1 = K_1 = \mathcal{F}_1 \geq 1$ is the number of symbols seen exactly once in the sample, and the supremum in (169) is equal to $N_1 \log(\frac{N_1}{n})$. If $\mathcal{F}_1 > 1$, then the supremum is approached as $\mathcal{F}_0 \rightarrow \infty$ and is not achieved by any $\mathcal{F}_0 \in \mathbb{N}$. If $\mathcal{F}_1 = 1$, then the supremum is achieved by all $\mathcal{F}_0 \in \mathbb{N}$, so we arbitrarily choose $\mathcal{F}_0^{(1)} = 0$ as an achiever.

If $i > 1$ or $\mathcal{F}_1 = 0$, then $N_i > K_i$ and the supremum in (169) is achieved by $\mathcal{F}_0^{(i)} \in \mathbb{N}$ where

$$\mathcal{F}_0^{(i)} \triangleq \operatorname{argmax}_{\mathcal{F}_0 \in \mathbb{N}} \left(\log \left(\frac{(K_i + \mathcal{F}_0)!}{\mathcal{F}_0!} \right) - N_i \log(K_i + \mathcal{F}_0) \right) \begin{cases} = 0 & : N_i \geq K_i H_{K_i} \\ \in [0, \frac{K_i^2 - N_i}{N_i - K_i} + 1] & : K_i H_{K_i} > N_i > K_i \end{cases} \quad (170)$$

where $H_k \triangleq \sum_{j=1}^k \frac{1}{j}$ is the k -th harmonic number.

If the supremum in (169) is achieved by finite $\mathcal{F}_0 = \mathcal{F}_0^{(i^*)}$ (this is the case unless $i^* = 1$ and $\mathcal{F}_1 > 1$) then the approximate PML distribution \bar{p}^* has finite support set size $\bar{K}^* = \mathcal{F}_0^{(i^*)} + \hat{K}$, where \hat{K} is the support set size of the empirical distribution \hat{p} . Then we compute \bar{p}^* as in Section 4.5 with known support set size \bar{K}^* . If the supremum is not achieved by finite \mathcal{F}_0 (this is the case if $i^* = 1$ and $\mathcal{F}_1 > 1$), then we say \bar{p}^* has a continuous part of mass N_1/n and discrete part \bar{p}_d^* of mass $1 - N_1/n$ supported on $\mathcal{X}_{2:F_+} = \mathcal{X} \setminus \mathcal{X}_{0:1}$. Then we compute \bar{p}_d^* as in Section 4.5 with known support set size $|\mathcal{X}_{2:F_+}|$.

Since the function in the argmax in (170) is unimodal in \mathcal{F}_0 and $0 \leq \mathcal{F}_0^{(i)} \leq \frac{K_i^2 - N_i}{N_i - K_i} + 1 \leq N_i^2 + 2$ (where the last inequality follows from $N_i \geq K_i + 1$), we can compute the optimizing integer $\mathcal{F}_0^{(i)}$ in $O(\log(N_i))$ iterations of a bisection search. We must do this bisection search

$F_+ = |\text{Supp}(\mathcal{F}_+)|$ times. As remarked earlier, since $|\text{Supp}(\mathcal{F}_+)| \leq \sqrt{2n} + 1$ and $N_i \leq n$ for all i , then the total run time of our support set size optimization scheme is $O(\sqrt{n} \log(n))$.

Thus to estimate the support set size of the approximate PML distribution we solve F_+ optimization problems, corresponding to estimation of the support set size of a uniform distribution (since all symbols in $\mathcal{X}_{0:i}$ are assigned to a single level set) given a sample of size N_i with K_i distinct symbols.

Proof

Optimizing over the support set size K , the approximate PML distribution \bar{p}^* satisfies:

$$\log(\mathbb{P}_{\bar{p}^*}(\mathcal{F})) = \sup_K \left(-\log((K - \hat{K})!) + \max_{p \in \mathcal{P}_K} (\bar{V}(p)) \right) \quad (171)$$

$$= \sup_{\mathcal{F}_0} \left(-\log(\mathcal{F}_0!) + \max_{p \in \mathcal{P}_{\hat{K} + \mathcal{F}_0}} (\bar{V}(p)) \right) \quad (172)$$

$$\stackrel{(a)}{=} \sup_{\mathcal{F}_0} \left(-\log(\mathcal{F}_0!) + \bar{V}_0 \right) \quad (173)$$

$$\stackrel{(b)}{=} \sup_{\mathcal{F}_0} \left(-\log(\mathcal{F}_0!) + \max_{i \in [1, F_+]} (\bar{V}(\mathcal{X}_{0,i}) + \bar{V}_{i+1}) \right) \quad (174)$$

$$\stackrel{(c)}{=} \max_{i \in [1, F_+]} \left(\bar{V}_{i+1} + \sup_{\mathcal{F}_0} (\bar{V}(\mathcal{X}_{0,i}) - \log(\mathcal{F}_0!)) \right) \quad (175)$$

$$\stackrel{(d)}{=} \max_{i \in [1, F_+]} \left(\bar{V}_{i+1} + N_i \log \left(\frac{N_i}{n} \right) + \underbrace{\sup_{\mathcal{F}_0} \left(\log \left(\frac{(K_i + \mathcal{F}_0)!}{\mathcal{F}_0!} \right) - N_i \log(K_i + \mathcal{F}_0) \right)}_{f(\mathcal{F}_0, K_i, N_i)} \right) \quad (176)$$

where (a) follows from (48), (b) follows from (46), (c) follows since \bar{V}_i is independent of \mathcal{F}_0 (and hence K) for $i \geq 1$, and (d) from (39) and (50). We use $\sup_{\mathcal{F}_0}$ (rather than $\max_{\mathcal{F}_0}$) since a maximizing value of \mathcal{F}_0 might not exist. We later show that the supremum in (176) is finite, so we interchange the order of the max and supremum. The function $f(\mathcal{F}_0, K, N)$ is defined for non-negative arguments as in (176). Note that we do not allow $i = 0$ in the optimization (176) in order to exclude $\mathcal{X}_{0,0}$ (the set of unseen symbols) as a level set of p ; otherwise, due to the averaging-over-clumps property (37), we would have p vanish on $\mathcal{X}_{0,0}$, violating the assumption that the support set size of p is K . It is possible to have $\bar{K}^* = \hat{K} = |\text{Supp}(\hat{p})|$, corresponding to the case $\mathcal{F}_0 = 0$, so that $\mathcal{X}_{0,i} = \mathcal{X}_{1,i}$ for all i .

Let's compute the supremization in (176). We can check that $f(\mathcal{F}_0, N, N) < 0$ for all $N > 1$ and $\mathcal{F}_0 \geq 0$, and

$$\lim_{\mathcal{F}_0 \rightarrow \infty} f(\mathcal{F}_0, N, N) = 0 \quad (177)$$

so if $i = 1$ and $\mathcal{F}_1 > 1$, then $N_1 = K_1 = \mathcal{F}_1 > 1$ and the supremum in (176) is equal to 0, is approached as $\mathcal{F}_0 \rightarrow \infty$, and is not achieved by any $\mathcal{F}_0 \in \mathbb{N}$.

We can check that $f(\mathcal{F}_0, 1, 1) = 0$ for all \mathcal{F}_0 , so if $i = 1$ and $\mathcal{F}_1 = 1$, then $N_1 = K_1 = \mathcal{F}_1 = 1$ and the supremum in (176) is equal to 0 and is achieved by all $\mathcal{F}_0 \in \mathbb{N}$, so we arbitrarily choose $\mathcal{F}_0^{(1)} = 0$ as an achiever.

If $i > 1$ or $\mathcal{F}_1 = 0$, then $N_i > K_i$. When $N > K$ the function $f(\mathcal{F}_0, K, N)$ is continuous in \mathcal{F}_0 ¹¹ and unimodal with a global maximum. Solving

$$\frac{\partial}{\partial \mathcal{F}_0} f(\mathcal{F}_0, K, N^*) = 0 \quad (178)$$

for N^* we find $N^* = KH_K$, where H_K is the K -th harmonic number. Thus if $N_i \geq K_i H_{K_i}$, then the supremum in (176) is achieved by $\mathcal{F}_0 = 0$. If $N_i < K_i H_{K_i}$, then we can upper bound the supremizing value \mathcal{F}_0 by finding the unique inflection point of f in \mathcal{F}_0 beyond the global maximum. Setting the “discrete” second derivative to 0 (we take the discrete derivative because the function $\frac{\partial f}{\partial \mathcal{F}_0}$ is easier to work with for integer \mathcal{F}_0):

$$\frac{\partial}{\partial \mathcal{F}_0} (f(\mathcal{F}_0 + 1, K, N) - f(\mathcal{F}_0, K, N)) \Big|_{\mathcal{F}_0 = \mathcal{F}_0^*} = 0 \quad (179)$$

we find $\mathcal{F}_0^* = \frac{K^2 - N}{N - K}$. If $K < N < KH_K$ then $K > 1$ (since $1 = H_1$) and $N < K^2$, so $\mathcal{F}_0^* > 0$. Thus we can upper bound the inflection point of f by $\frac{K^2 - N}{N - K} + 1$, so if $K_i < N_i < K_i H_{K_i}$, then $\mathcal{F}_0^{(i)} \in [0, \frac{K_i^2 - N_i}{N_i - K_i} + 1]$. ■

References

- Jayadev Acharya, Alon Orlitsky, and Shengjun Pan. The maximum likelihood probability of unique-singleton, ternary, and length-7 patterns. In *Information Theory (ISIT), 2009 IEEE International Symposium on*, pages 1135–1139. IEEE, 2009.
- 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, 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.
- Dragi Anevski, Richard D Gill, and Stefan Zohren. Estimating a probability mass function with unknown labels. *The Annals of Statistics*, 45(6):2708–2735, 2017.
- FC Auluck. On partitions of bipartite numbers. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 49, pages 72–83. Cambridge Univ Press, 1953.
- I Bárány and AM Vershik. On the number of convex lattice polytopes. *Geometric and Functional Analysis*, 2(4):381–393, 1992.
- D Basu. On partial sufficiency: A review. In *Selected Works of Debabrata Basu*, pages 291–303. Springer, 2011.

11. Replacing $x!$ with $\Gamma(x + 1)$, where Γ is the gamma function.

- Tugkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing that distributions are close. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 259–269. IEEE, 2000.
- Allan Birnbaum. On the foundations of statistical inference. *Journal of the American Statistical Association*, 57(298):269–306, 1962.
- David A Blackwell and Meyer A Girshick. *Theory of games and statistical decisions*. Courier Corporation, 1979.
- Yuheng Bu, Shaofeng Zou, Yingbin Liang, and Venugopal V Veeravalli. Estimation of kl divergence: Optimal minimax rate. *arXiv preprint arXiv:1607.02653*, 2016.
- Chun Lam Chan, Winston Fernandes, Navin Kashyap, and Majunath Krishnapur. Phase transitions for the uniform distribution in the pattern maximum likelihood problem and its Bethe approximation. *SIAM J. Discrete Math*, 31:597–631, 2015.
- Bradley Efron. Maximum likelihood and decision theory. *The Annals of Statistics*, pages 340–356, 1982.
- Winston Fernandes and Navin Kashyap. A phase transition for the uniform distribution in the pattern maximum likelihood problem. In *ITW*, 2013.
- Jaroslav Hájek. On basic concepts of statistics. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probabilities*, volume 1, pages 139–162, 1967.
- YanJun Han, Jiantao Jiao, and Tsachy Weissman. Minimax rate-optimal estimation of divergences between discrete distributions. *arXiv preprint arXiv:1605.09124*, 2016a.
- YanJun Han, Jiantao Jiao, and Tsachy Weissman. Minimax rate-optimal estimation of KL divergence between discrete distributions. In *Information Theory and Its Applications (ISITA), 2016 International Symposium on*, pages 256–260. IEEE, 2016b.
- Godfrey H Hardy and Srinivasa Ramanujan. Asymptotic formulæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.
- David A Harville. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*, 72(358):320–338, 1977.
- Jiantao Jiao, Kartik Venkat, YanJun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *Information Theory, IEEE Transactions on*, 61(5): 2835–2885, 2015.
- Jiantao Jiao, Kartik Venkat, and Tsachy Weissman. Relations between information and estimation in discrete-time Lévy channels. *IEEE Transactions on Information Theory*, 63(6):3579–3594, 2017.
- Jiantao Jiao, YanJun Han, and Tsachy Weissman. Minimax estimation of the L_1 distance. *IEEE Transactions on Information Theory*, 64(10):6672–6706, 2018.

- A Kolmogoroff. Sur l'estimation statistique des paramètres de la loi de Gauss. *Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya*, 6(1):3–32, 1942.
- L Le. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, pages 1419–1455, 1964.
- Lucien Le Cam. Maximum likelihood: an introduction. *Statistics Branch, Department of Mathematics, University of Maryland*, 1979.
- Erich Leo Lehmann and George Casella. *Theory of point estimation*, volume 31. Springer, 1998.
- Susan A Murphy and Aad W Van der Vaart. On profile likelihood. *Journal of the American Statistical Association*, 95(450):449–465, 2000.
- OEIS Foundation Inc. The online encyclopedia of integer sequences (OEIS)–A007716. <https://oeis.org/A007716>. Accessed: 2017-05-16.
- Alon Orlitsky, Sajama, Narayana Santhanam, Krishnamurthy Viswanathan, and Junan Zhang. Algorithms for modeling distributions over large alphabets. In *Information Theory (ISIT), 2004 IEEE International Symposium on*, page 304. IEEE, 2004a.
- 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, 2004b.
- Daniel P Palomar and Sergio Verdú. Lautum information. *IEEE transactions on information theory*, 54(3):964–975, 2008.
- Liam Paninski. Estimation of entropy and mutual information. *Neural Computation*, 15(6):1191–1253, 2003.
- H Desmond Patterson and Robin Thompson. Recovery of inter-block information when block sizes are unequal. *Biometrika*, 58(3):545–554, 1971.
- Dmitri S. Pavlichin, Jiantao Jiao, and Tsachy Weissman. Approximate profile maximum likelihood (software), 2017. URL <https://doi.org/10.5281/zenodo.1043617>. URL <https://github.com/dmitrip/PML>.
- Aditi Raghunathan, Greg Valiant, and James Zou. Estimating the unseen from multiple populations. *arXiv preprint arXiv:1707.03854*, 2017.
- Gregory Valiant and Paul Valiant. The power of linear estimators. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 403–412. IEEE, 2011a.
- 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 *the 43rd ACM Symposium on Theory of Computing (STOC)*, 2011b.

- Gregory Valiant and Paul Valiant. Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems*, pages 2157–2165, 2013.
- L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.
- Paul Valiant. *Testing Symmetric Properties of Distributions*. PhD thesis, Massachusetts Institute of Technology, 2008.
- SS Vallender. Calculation of the wasserstein distance between probability distributions on the line. *Theory of Probability & Its Applications*, 18(4):784–786, 1974.
- Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Pascal O. Vontobel. The Bethe approximation of the pattern maximum likelihood distribution. In *ISIT*, 2012.
- Abraham Wald. *Statistical decision functions*. Wiley, 1950.
- Haizhou Wang and Mingzhou Song. Ckmeans.1d.dp: Optimal k-means clustering in one dimension by dynamic programming. *The R Journal*, 3(2):29–33, 2011.
- Yihong Wu and Pengkun Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *arXiv preprint arXiv:1504.01227*, 2015.
- Yihong Wu and Pengkun Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, 2016.