

© 2019 I Chien

REGULARIZED WEIGHTED CHEBYSHEV APPROXIMATIONS FOR
SUPPORT ESTIMATION

BY

I CHIEN

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2019

Urbana, Illinois

Adviser:

Professor Olgica Milenkovic

ABSTRACT

We introduce a new method for estimating the support size of an unknown distribution which provably matches the performance bounds of the state-of-the-art techniques in the area and outperforms them in practice. In particular, we present both theoretical and computer simulation results that illustrate the utility and performance improvements of our method. The theoretical analysis relies on introducing a new weighted Chebyshev polynomial approximation method, jointly optimizing the bias and variance components of the risk, and combining the weighted minmax polynomial approximation method with discretized semi-infinite programming solvers. Such a setting allows for casting the estimation problem as a linear program (LP) with a small number of variables and constraints that may be solved as efficiently as the original Chebyshev approximation problem. Our technique is tested on synthetic data and textual data (Shakespeare's plays), and is used to address an important problem in computational biology - estimating the number of bacterial genera in the human gut. On synthetic datasets, for practically relevant sample sizes, we observe significant improvements in the value of the worst-case risk compared to existing methods. The same is true of the text data. For the bioinformatics application, using metagenomic data from the NIH Human Gut and the American Gut Microbiome Projects, we generate a list of frequencies of bacterial taxa that allows us to estimate the number of bacterial genera to ~ 2300 .

To my parents, for their love and support.

ACKNOWLEDGMENTS

My thanks go to my wonderful adviser, Professor Olgica Milenkovic, for her intelligent suggestions and valuable discussions. She also gives me a lot of freedom in the direction of research. Meanwhile, her great personality also makes the ambiance in our group quite friendly and interactive. I really enjoy doing research in the group of Professor Milenkovic. I look forward to working as a PhD student in the next few years with her on the topic of clustering and machine learning. I would like to thank my collaborators Pan Li, Huozhi Zhou, Chao Pan, Anuththari Gamage and Jianhao Peng for their contributions to the papers we have published during my master's program. Also thanks to my summer internship mentors, Antonia Maria Tulino and Jaime Llorca, from Nokia Bell Labs. I learned a lot from them, and doing research with them is indeed an experience of a lifetime. I am grateful to my labmates Chuanyi Zhang, Srilakshmi Pattabiraman, Abhishek Agarwal, Vishal Rana and Seiyun Shin for many good discussions. I thank my undergrad adviser I-Hsiang Wang, who led me to the study of learning on graphs. Finally, I owe special thanks to my parents for their love and support throughout my life.

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION	1
1.1 Prior work	1
1.2 Our contributions	2
CHAPTER 2 PROBLEM FORMULATION AND NEW POLY- NOMIAL CLASS ESTIMATORS	5
CHAPTER 3 ESTIMATOR ANALYSIS	8
3.1 Solving problem (3.1)	10
CHAPTER 4 SIMULATIONS	13
CHAPTER 5 CONCLUSION	18
APPENDIX A PROOF OF THEOREM 3.1	19
APPENDIX B PROOF OF THEOREM 3.2	23
B.1 Proof of convergence	23
B.2 Proof for the convergence rate	25
APPENDIX C THEORETICAL RESULTS SUPPORTING RE- MARK 3.1	28
APPENDIX D CONSTRUCTION OF THE RWC-S ESTIMATOR	29
REFERENCES	31

CHAPTER 1

INTRODUCTION

Estimating the support size of a discrete distribution is an important theoretical and data-driven problem [1, 2]. In computer science, this task frequently arises in large-scale database mining and network monitoring where the objective is to estimate the types of database entries or IP addresses from a limited number of observations [3, 4, 5]. In machine learning, support estimation is used to bound the number of clusters in clustering problems arising in semi-supervised or active learning [6, 7, 8]. The most challenging practical support estimation issues arise in the “small sample set” regime in which one has only a limited number of observations for a large-support distribution. In such a setting, classical maximum likelihood frequency techniques are known to perform poorly [9]. It is for this sampling regime that the estimation problem has received significant attention from both the theoretical computer science and machine learning communities, as well as researchers in various computational data processing areas [10, 11, 12, 13, 5, 14, 15, 16].

1.1 Prior work

Approaches to distribution estimation in the small sample regime may be roughly grouped into two categories [17, 18, 19, 20, 21, 22, 23]. The first line of works [17, 19, 21] makes use of the maximum likelihood principle. While [19] constructs estimators based on the profile maximum likelihood (PML) [24], the work reported in [17] focuses on sequence maximum likelihood (SML) estimators [25]. The main advantage of ML-based methods is that they easily generalize to many other estimation tasks. For example, the authors of [19] showed that one and the same method may be used for entropy estimation, support coverage and distance to uniformity analysis. However, most ML-based estimators require large computational resources [21, 18].

To address the computational issue, a sophisticated approximate PML technique that reduces the computational complexity of support estimation at the expense of some performance loss was described in [21]. On the other hand, the second line of works [18, 20, 22, 23] formulates support estimation as an approximation problem. The underlying methods, which we henceforth refer to as approximation-based methods, design estimators by minimizing the worst case risk. In particular, [18] uses shifted and scaled Chebyshev polynomials of the first kind to construct efficient estimators. In contrast, the authors of [23] suggest disposing with minmax estimators and introduce a data amplification technique with analytical performance guarantees. The aforementioned estimator is based on polynomial smoothing [26] related to approximation techniques. Note that all described approximation-based estimators are computationally efficient, with the exception of [22], as reported in [23].

1.2 Our contributions

We adopt the approximation approach put forward in [18], but significantly improve it in practice by using *weighted* polynomial approximation techniques largely unknown in the machine learning community [27]. The weighted approximation approach significantly improves the performance of known approximation methods while allowing one to seamlessly combine certain ideas already explored in different contexts for the support estimation problem. In particular, the technical novelty of the proposed method includes uniquely combining the following approaches:

Using weighted Chebyshev approximations: The authors of [18] used classical Chebyshev approximations to construct their estimator, while ignoring the exponential weighting term that arises due to Poissonization. As will be shown in subsequent analyses, the exponential weights play a major role in improving the performance of our method as well as making it computationally tractable through a new take on the “localization idea”.

Using the variance as a regularizer: The idea of jointly optimizing the bias and variance of entropy estimators has been previously reported in [11], but only within the framework of ML estimation. The authors of [18] focused on optimizing the bias, but accounted for the variance through a separate

parameter tuning technique. To the best of our knowledge, *jointly optimizing the bias and variance terms in approximation-based methods* is novel.

Using semi-infinite programming (SIP) techniques with discretization: SIP techniques have also been explored in the approximation-based support estimation literature [22], but for a different objective and without rigorously establishing the convergence results of the discretization technique. We show that the solution of our discretized SIP for the regularized weighted Chebyshev approximation problem converges to the true unique optimal solution.

The theoretical and practical results presented are as follows. First, we cast the estimation problem in terms of a weighted and regularized Chebyshev approximation [28] problem for a normalized quadratic loss function, in which the normalization term is an upper bound on the support [18]. The regularized weighted Chebyshev estimator provably offers *worst case risk bounds* that match or improve those of the estimator in [18]. Although the risk bound of the aforementioned estimator is order-optimal, significant improvements in the risk exponent are possible and demonstrated in the experimental section of our work. It is worth pointing out that *adding weights* into the approximation formulation is of crucial importance in reducing the error in the bias. We solve the regularized weighted Chebyshev approximation problem via its epigraph formulation, which takes the form of a semi-infinite program (SIP) [29]. We prove that the underlying SIP can be solved consistently and efficiently through discretization, resulting in a small linear program (LP) of size *decreasing with the number of samples*. Interestingly, despite the fact that we also use an LP solver as was done in one of the best performing ML-based approaches [17], the ML-LP formulation has a number of variables and constraints that increase with the number of samples. Our experimental results also reveal that the regularized weighted Chebyshev estimator significantly outperforms the low-complexity ML estimator described in [21], the benchmark approximation algorithm [18], as well as the recently proposed sample amplification estimator [23]. The running times of all described algorithms are comparable on moderate-to-large sample set sizes ($\sim 10^6$).

Additional simulations also account for the fact that the normalized square loss causes the worst-case distribution to be close to uniform, resulting in performance issues previously described in [23]. To address this problem, we use a different normalization term in the squared loss equal to the actual support size. We show that the weighted regularized Chebyshev estimator

that minimizes the worst case risk under the modified loss outperforms all methods mentioned above and performs consistently well on many tested distributions.

The thesis is organized as follows. In Chapter 2, we introduce the relevant notation and the support estimation problem, and describe a class of estimators termed *polynomial class estimators*. The same chapter also provides a brief review of relevant results from [18]. In Chapter 3, we outline the bias analysis for polynomial class estimators and describe the additional technical challenges one needs to overcome when dealing with weighted min-max polynomial approximations. The same chapter contains our worst-case risk analysis. Chapter 4 is devoted to experimental verifications and testing, both on synthetically generated data and metagenomic data samples from the NIH Human Microbiome [30] and American Gut Microbiome project (<http://americangut.org/publications/>).

CHAPTER 2

PROBLEM FORMULATION AND NEW POLYNOMIAL CLASS ESTIMATORS

Let $P = (p_1, p_2, \dots)$ be a discrete distribution over some countable alphabet and let X_1, \dots, X_n be i.i.d. samples drawn according to the distribution P . The problem of interest is to estimate the support size, defined as $S(P) = \sum_i \mathbf{1}_{\{p_i > 0\}}$. Henceforth, we use S instead of $S(P)$ to avoid notational clutter. We make the assumption that the minimum non-zero probability of the distribution P is greater than $\frac{1}{k}$, for some $k \in \mathbb{R}^+$, i.e., $\inf\{p \in P \mid p > 0\} \geq \frac{1}{k}$. Furthermore, we let D_k denote the space of all probability distributions satisfying $\inf\{p \in P \mid p > 0\} \geq \frac{1}{k}$. Clearly, $S \leq k$, $\forall P \in D_k$. A sufficient statistic for X_1, \dots, X_n is the empirical distribution (i.e., histogram) $N = (N_1, N_2, \dots)$, where $N_i = \sum_{j=1}^n \mathbf{1}_{\{X_j=i\}}$ and $\mathbf{1}_A$ stands for the indicator function of the event A .

The focal point of our analysis is $R^*(k, n)$, the minmax risk under normalized squared loss,

$$R^*(k, n) = \inf_{\hat{S}} \sup_{P \in D_k} \mathbb{E} \left[\left(\frac{\hat{S}(N) - S}{k} \right)^2 \right]. \quad (2.1)$$

We seek a support estimator \hat{S} that minimizes

$$\begin{aligned} & \sup_{P \in D_k} \mathbb{E} \left[\left(\frac{\hat{S}(N) - S}{k} \right)^2 \right] \\ &= \sup_{P \in D_k} \left[\mathbb{E}^2 \left(\frac{\hat{S}(N) - S}{k} \right) + \text{var} \left(\frac{\hat{S}(N) - S}{k} \right) \right]. \end{aligned}$$

The first term within the supremum captures the expected bias of the estimator \hat{S} . The second term represents the variance of the estimator \hat{S} . Hence, “good” estimators are required to balance out the worst-case contributions of the bias and variance.

The Chebyshev polynomial of the first kind of degree L is defined as $T_L(x) = \cos(L \arccos(x)) = \frac{z^L + z^{-L}}{2}$, where z is the solution of the quadratic equation $z + z^{-1} = 2x$. From the definition, it is easy to see that the polynomial T_L is bounded in the interval $[-1, 1]$. Chebyshev polynomials may be scaled and shifted to lie in an interval $[l, r]$ not necessarily equal to $[-1, 1]$,

$$R_L(x) = -\frac{T_L\left(\frac{2x-r-l}{r-l}\right)}{T_L\left(\frac{-r-l}{r-l}\right)} \triangleq \sum_{j=0}^L \tilde{a}_j x^j.$$

Clearly, $R_L(0) = -1$, and $\tilde{a}_0 = -1$. The coefficients in the above expansion equal

$$\tilde{a}_j = \frac{R_L^{(j)}(0)}{j!}. \quad (2.2)$$

The estimator used in [18] takes the form $\tilde{S} = \sum_i \tilde{g}_L(N_i)$, where

$$\tilde{g}_L(j) = \begin{cases} \tilde{a}_j j! + 1, & \text{if } j \leq L, \\ 1, & \text{otherwise,} \end{cases}, \quad (2.3)$$

$$L \triangleq \lfloor c_0 \log k \rfloor, \quad [l, r] \triangleq \left[\frac{n}{k}, c_1 \log k \right].$$

The estimator \tilde{S} is order-optimal *in the exponent* under the risk (2.1). Nevertheless, the estimator \tilde{S} is designed to minimize only a bias term that ignores a multiplicative exponential and for a given pair of parameters (c_0, c_1) . The parameters (c_0, c_1) are set to $c_0 \approx 0.558$ and $c_1 = 0.5$ in order to ensure provably good balance between the bias and variance of the estimator. The performance guarantees of the estimator are stated in the theorem that follows.

Theorem 2.1 (Theorem 1 in [18]). For all $k, n \geq 2$, the optimal achievable risk (2.1) is bounded as

$$R^*(k, n) = \exp(-\Theta(\sqrt{\frac{n \log k}{k}} \vee \frac{n}{k} \vee 1)) \quad (2.4)$$

In addition, if $\frac{k}{\log k} \ll n \ll k \log k$, as $k \rightarrow \infty$,

$$\begin{aligned} \exp(-(3.844 + o(1))\sqrt{\frac{n \log k}{k}}) &\leq R^*(k, n) \\ &\leq \exp(-(1.579 + o(1))\sqrt{\frac{n \log k}{k}}). \end{aligned}$$

The estimator \tilde{S} attains the upper bound, which has a constant negative exponent equal to one half of the one corresponding to the lower bound. Consequently, \tilde{S} may be improved by accounting for the omitted exponential weights and jointly optimizing the bias and variance term in the squared loss. This is accomplished by our proposed estimator that *directly* operates on the weighted squared loss.

To this end, we define an extended class of polynomial based estimators as follows. Given a parameter $L \in \mathbb{N}$, we say that an estimator \hat{S} is a polynomial class estimator with parameter L (i.e., a $Poly(L)$ estimator) if it takes the form $\hat{S} = \sum_i g_L(N_i)$, where g_L is defined as

$$g_L(j) = \begin{cases} a_j j! + 1, & \text{if } j \leq L \\ 1, & \text{otherwise.} \end{cases} \quad (2.5)$$

Here, $a_j \in \mathbb{R}$, and $a_0 = -1$, since this choice ensures that $g_L(0) = 0$. One can associate an estimator \hat{S} with its corresponding coefficients \mathbf{a} , and define a family of estimators

$$Poly(L) = \left\{ \mathbf{a} \in \mathbb{R}^{L+1} \mid a_0 = -1 \right\}.$$

Clearly, $\tilde{S} \in Poly(L)$, with corresponding coefficients

$$\tilde{\mathbf{a}} = \arg \min_{\mathbf{a} \in Poly(L)} \sup_{\lambda \in [l=\frac{n}{k}, r=c_1 \log k]} \left| \sum_{l=0}^L a_l \lambda_i^l \right|. \quad (2.6)$$

Next, we show that the original minmax problem can be rewritten as a regularized exponentially weighted Chebyshev approximation problem [27]. Once a proper interval is identified, it is possible to efficiently determine the best coefficients \mathbf{a} in the class $Poly(L)$.

CHAPTER 3

ESTIMATOR ANALYSIS

In order to jointly optimize the bias and variance term in the squared loss, we start by analyzing $\sup_{P \in D_k} \mathbb{E} \left(\frac{S - \hat{S}}{k} \right)^2$ directly. Classical Poissonization arguments [18] lead to

$$\begin{aligned} \mathbb{E} \left(\frac{S - \hat{S}}{k} \right)^2 &= \frac{1}{k^2} \left\{ \sum_{i: \lambda_i > 0} \left(\sum_{l=0}^L e^{-\lambda_i} a_l^2 \lambda_i^l l! \right) \right. \\ &\quad \left. + \sum_{i \neq j: \lambda_i \lambda_j > 0} \left(e^{-\lambda_i} P_L(\lambda_i, \mathbf{a}) \right) \left(e^{-\lambda_j} P_L(\lambda_j, \mathbf{a}) \right) \right\}, \end{aligned}$$

where $P_L(\lambda, \mathbf{a}) \triangleq \sum_{l=0}^L a_l \lambda^l$. Taking the supremum over D_k we can bound the risk as

$$\begin{aligned} &\leq \sup_{\lambda: \lambda_i \in [\frac{n}{k}, n]} \frac{1}{k^2} \left\{ \sum_{i: \lambda_i > 0} \left(\sum_{l=0}^L e^{-\lambda_i} a_l^2 \lambda_i^l l! \right) \right. \\ &\quad \left. + \sum_{i \neq j: \lambda_i \lambda_j > 0} \left(e^{-\lambda_i} P_L(\lambda_i, \mathbf{a}) \right) \left(e^{-\lambda_j} P_L(\lambda_j, \mathbf{a}) \right) \right\} \\ &\leq \sup_{\lambda \in [\frac{n}{k}, n]} \left\{ \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\}. \end{aligned}$$

In the above inequality, we used the fact that $S \leq k$ and $\left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) > 0$, for all $\lambda > 0$. Hence, the optimization problem at hand reads as

$$\inf_{\mathbf{a} \in \text{Poly}(L)} \sup_{\lambda \in [\frac{n}{k}, n]} \left\{ \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\}. \quad (3.1)$$

Problem (3.1) represents a regularized weighted Chebyshev approximation problem. If we ignore the first term in (3.1), the optimization problem be-

comes

$$\inf_{\mathbf{a} \in \text{Poly}(L)} \sup_{\lambda \in [\frac{n}{k}, n]} \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2.$$

The term $e^{-\lambda} P_L(\lambda, \mathbf{a})$ corresponds to the bias of the estimator. It is straightforward to see that the optimal choice of \mathbf{a} for the above problem is a solution of

$$\inf_{\mathbf{a} \in \text{Poly}(L)} \sup_{\lambda \in [\frac{n}{k}, n]} \left| e^{-\lambda} P_L(\lambda, \mathbf{a}) \right|. \quad (3.2)$$

Problem (3.2) is an exponential weighted Chebyshev approximation problem [28].

The first term $\frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right)$, which corresponds to the variance, may be written as

$$\begin{aligned} \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) &= \mathbf{a}^T \mathbf{M}(\lambda) \mathbf{a} \triangleq \|\mathbf{a}\|_{\mathbf{M}(\lambda)}^2, \\ \mathbf{M}(\lambda) &\triangleq \frac{e^{-\lambda}}{k} \text{Diag}(\lambda^0 0!, \lambda^1 1!, \dots, \lambda^L L!). \end{aligned}$$

Clearly, $\|\cdot\|_{\mathbf{M}(\lambda)}$ is a valid norm, and consequently, the first term in (3.1) can be viewed as a regularizer.

Note that both problems (2.6) and (3.2) aim to minimize the bias term. However, simple algebra reveals

$$\sup_{P \in D_k} \frac{1}{k} |\mathbb{E}(S - \hat{S}(N))| \leq \sup_{\lambda \in [\frac{n}{k}, n]} |e^{-\lambda} P_L(\lambda, \mathbf{a})| \quad (3.3)$$

$$\leq e^{-\frac{n}{k}} \sup_{\lambda \in [\frac{n}{k}, n]} |P_L(\lambda, \mathbf{a})|, \quad (3.4)$$

where (3.3) is equivalent to (3.2), while (3.4) resembles problem (2.6), except for a different optimization interval used within the supremum. Optimizing (3.3) instead of (3.4) should give us a smaller bias as the exponential weight is inherent to the formulation. Note that the authors of [18] choose a shorter interval in (2.6) in order to decrease the contribution of the variance to the loss. The modified bound in (3.4) was minimized with respect to the coefficients \mathbf{a} , using the minmax property of Chebyshev polynomials [31, 28], resulting in $\tilde{\mathbf{a}}$.

Let us now turn back to the solution of the minmax problem (3.1), denoted by \mathbf{a}^* . Clearly, using \mathbf{a}^* instead of $\tilde{\mathbf{a}}$ is guaranteed to reduce the reported *upper bound* on the risk since \mathbf{a}^* *jointly* minimizes both contributing terms.

Proposition 3.1. Let \hat{S}^* be the estimator associated with the coefficients \mathbf{a}^* optimizing (3.1). Whenever $\frac{k}{\log k} \ll n \ll k \log k$, and $k \rightarrow \infty$, the estimator \hat{S}^* has a smaller worst-case risk bound than \tilde{S} .

Proof. The proof follows directly from the fact that the upper bound of Theorem 2.1 is derived by upper bounding (3.1), plugging $\tilde{\mathbf{a}}$ into the expression, and replacing $e^{-\lambda}$ by $e^{-\frac{n}{k}}$. \square

We remark that improving upper bounds does not establish improvements in the actual worst-case risk. Nevertheless, whenever faced with complicated analytical settings, comparisons of bounds may provide useful insights, and in our case, these insights are supported by strong simulation evidence.

3.1 Solving problem (3.1)

To directly solve the optimization problem of interest, we turn to weighted Chebyshev approximations [28] and semi-infinite programming. Solving for \mathbf{a}^* directly appears to be difficult, so we instead resort to numerically solving the epigraph formulation of problem (3.1) and proving that the numerical solution is asymptotically consistent.

The epigraph formulation of (3.1) is of the form ([32], Chapter 6.1)

$$\begin{aligned} & \min_{t, a_1, \dots, a_L} t \quad \text{subject to} \\ & \left\{ \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\} \leq t, \\ & \forall \lambda \in \left[\frac{n}{k}, n \right], \text{ with } a_0 = -1. \end{aligned} \quad (3.5)$$

Note that (3.5) is a semi-infinite programming problem. There are many algorithms that can be used to numerically solve (3.5), such as the discretization method, and the central cutting plane, KKT reduction and SQP reduction methods [29, 33]. For simplicity, we focus on the discretization method. For this purpose, we first form a grid of the interval $[\frac{n}{k}, n]$ involving s points,

denoted by $\text{Grid}([\frac{n}{k}, n], s)$. Problem (3.5) may be viewed as an LP with infinitely many quadratic constraints, which is not solvable. Hence, instead of addressing (3.5), we focus on solving the relaxed problem

$$\begin{aligned} & \min_{t, a_1, \dots, a_L} t \quad \text{subject to} \\ & \left\{ \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\} \leq t, \\ & \forall \lambda \in \text{Grid}([\frac{n}{k}, n], s), \text{ with } a_0 = -1. \end{aligned} \quad (3.6)$$

As will be discussed in greater detail in Section 3.1.1, the solution of the relaxed problem is asymptotically consistent with the solution of the original problem (i.e., as s goes to infinity, the optimal values of the objectives of the original and relaxed problem are equal). Problem (3.6) is an LP with a finite number of quadratic constraints that may be solved using standard optimization tools. Unfortunately, the number of constraints scales with the length of the grid interval, which in the case of interest is linear in n . This is an undesired feature of the approach, but it may be mitigated through the following theorem which demonstrates that an optimal solution of the problem may be found over an interval of length proportional to the significantly smaller value $\log k$, where $\frac{k}{\log k} \lesssim n$ is the fundamental bound for support estimation. We relegate the proof to Appendix A.

Theorem 3.1. For any $\mathbf{a} \in \text{Poly}(L)$ and $L = \lfloor c_0 \log k \rfloor$, and $c_0 = 0.558$, let

$$g(\mathbf{a}, \lambda) = \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2.$$

Then, we have

$$\sup_{\lambda \in [\frac{n}{k}, n]} g(\mathbf{a}, \lambda) = \begin{cases} \sup_{\lambda \in [\frac{n}{k}, 6.5L]} g(\mathbf{a}, \lambda) & \text{if } \frac{n}{k} \leq 6.5L \\ g(\mathbf{a}, \frac{n}{k}) & \text{if } \frac{n}{k} > 6.5L. \end{cases}$$

Remark 3.1. Using exponential weighted approximation theory [27] (Theorem C.1 in the Appendix C), one may shrink the optimization interval to $[\frac{n}{k}, \frac{\pi}{2}L + \frac{n}{k}]$ if the regularization term is omitted. However, it remains an open problem to extend the approach of [27] used in our proof to account for more general weighted sums of polynomials.

Consequently, the optimization problem of interest equals

$$\begin{aligned} & \min_{t, a_1, \dots, a_L} t \quad \text{subject to} \\ & \left\{ \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\} \leq t, \\ & \forall \lambda \in \text{Grid}\left(\left[\frac{n}{k}, 6.5L\right], s\right), \text{ with } a_0 = -1. \end{aligned} \quad (3.7)$$

Since $L = \lfloor c_0 \log k \rfloor$, the length of the optimization interval in (3.7) is proportional to $\log k$.

3.1.1 Convergence of the discretization method

It seems intuitive to assume that as s grows, the solution of the relaxed semi-infinite program approaches the optimal solution of the original problem (3.5). This intuition can be rigorously justified for the case of objective functions and constraints that are “well-behaved”, as defined in [34] and [35]. The first line of work describes the conditions needed for convergence, while the second establishes the convergence rate given that the discretized solver converges. We use these results in conjunction with a number of properties of our objective SIP to establish the claim in the following theorem. Once again, the proof is relegated to Appendix B.

Theorem 3.2. Let s be the number of uniformly placed grid points on the interval (3.7), and let $d \triangleq \frac{6.5L - \frac{n}{k}}{s-1}$ be the length of the discretization interval. As $d \rightarrow 0$, the optimal objective value t_d of the discretized SIP (3.7) converges to the optimal objective value of the original SIP t^* . Moreover, the optimal solution is unique. The convergence rate of t_d to t^* equals $O(d^2)$. If the optimal solution of the SIP is a strict minimum of order one (i.e., if $t - t^* \geq C \|\mathbf{a} - \mathbf{a}^*\|$ for some constant $C > 0$ and for all feasible neighborhoods of \mathbf{a}^*), then the solution of the discretized SIP also converges to an optimal solution with rate $O(d^2)$.

CHAPTER 4

SIMULATIONS

Next, we compare our estimator, referred to as the regularized weighted Chebyshev (RWC) method, with the Good-Turing (GT) estimator, the WY estimator of [18], the PJW estimator described in [21] and the HOSW estimator of [23]. We do not compare our method with the estimators introduced in [17, 22] due to their high computational complexity [18, 23].

In the first experiment, we evaluate the maximum risk with normalized squared loss of all listed estimators over six different distributions: The uniform distribution with $p_i = \frac{1}{k}$, the Zipf distributions with $p_i \propto i^{-\alpha}$, and α equal to 1.5, 1, 0.5 or 0.25, and the Benford distribution with $p_i \propto \log(i+1) - \log(i)$. We choose the support sizes for the Zipf and Benford distribution so that the minimum non-zero probability mass is roughly 10^{-6} . We ran the estimator 100 times to calculate the risk. For both approximation-based estimators, we fix c_0 to be 0.558. With our proposed method, we solved (3.7) on a grid with $s = 1000$ points on the proposed interval $[\frac{n}{k}, 6.5L]$. For the estimator described in [18], we set $c_1 = 0.5$ according to the recommendation made in the cited paper. The GT method used for comparison first estimates the total probability of seen symbols (e.g., sample coverage) according to $\hat{C} = 1 - \frac{h_1}{n}$, and then estimates the support size according to $\hat{S}_{GT} = \frac{\hat{S}_c}{\hat{C}}$; here, \hat{S}_c stands the (naive) counting estimator. Note that h_1 equals the number of different alphabet symbols observed only once in the n samples.

Figure 4.1(c) shows that the RWC estimator has a significantly better worst case performance compared to all other methods when tested on the above described collection of distributions, provided that $n \geq 0.2k$. Also, both RWC and WY estimators have significantly better error exponents compared to the GT, PJW and HOSW estimators. The GT and PJW estimators perform better than RWC if $n \lesssim \frac{k}{\log k}$, which confirms the results of Theorem 2.1 and of our Proposition 3.1.

In the second set of experiments, we change the normalization from $(1/k)^2$ to $(1/S)^2$ as was also done in [23]. The RWC-S estimator minimizes an upper bound on the worst case risk $\mathbb{E} \left(\frac{\hat{S}-S}{S} \right)^2$. A detailed description of this algorithm and intuitive explanation why it may outperform the RWC method is relegated to Appendix D. Figures 4.1(d) (i)-(l) illustrate that our RWC-S estimator significantly outperforms all other estimators with respect to the worst case risk. Moreover, the RWC-S estimator also outperforms all known estimators on almost all tested distributions.

Considering the exponential weighting term when designing the estimator leads to significant performance improvements in the bias of the estimator, as illustrated in Figure 4.1(a). There, we see that a classical Chebyshev approximation introduces a larger bias than the weighted one whenever the underlying distribution is close to uniform (i.e., when $\lambda \sim \frac{n}{k} = 1$). This phenomenon persists even when regularization is taken into account, as shown in Figure 4.1(m). Furthermore, regularized weighted Chebyshev estimators exhibit a different behavior than that observed for classical approximations: Figure 4.1(b) provides an example for the values of g_L used in RWC that do not alternate in sign. Nevertheless, these and all other estimators tested run for ≤ 1 s on sample sizes $\approx 10^6$.

One common approach to testing support estimators on real data is to estimate the number of distinct words in selected books. The authors of [18, 17] applied their estimators to *Hamlet* in order to estimate the number of words used in the play. Here, we repeat this experiment not only on *Hamlet* but *Othello*, *Macbeth* and *King Lear* as well. The comparative results for different estimators, following the setup in [18], are presented in Figure 4.2. In the experiments, we randomly sampled words in the text with replacement and used the obtained counts to estimate the number of distinct words. For simplicity, we set k to equal the total number of words. For example, as the total number of words in *Hamlet* equals 30,364, we set $k = 30,364$. As may be clearly seen, our method significantly outperforms all other competitive techniques both in terms of convergence rate and the accuracy of the estimated support for all experiments. These results further strengthen the case for the practicality of our estimators.

We now turn our attention to a new support estimation problem, concerned with determining the bacterial diversity of the human gut micro-

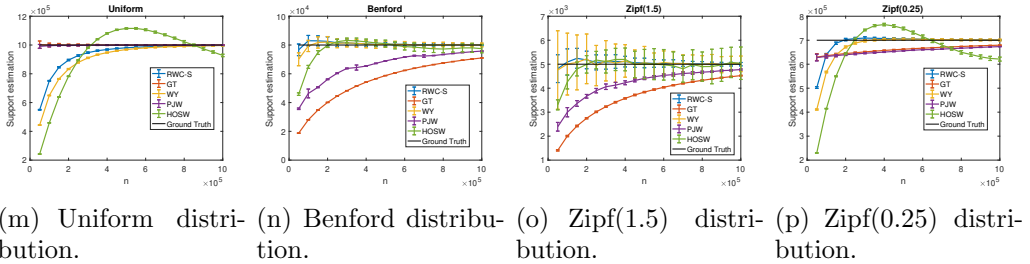
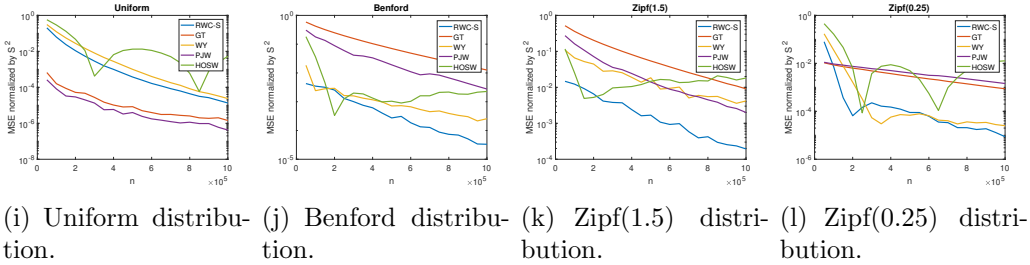
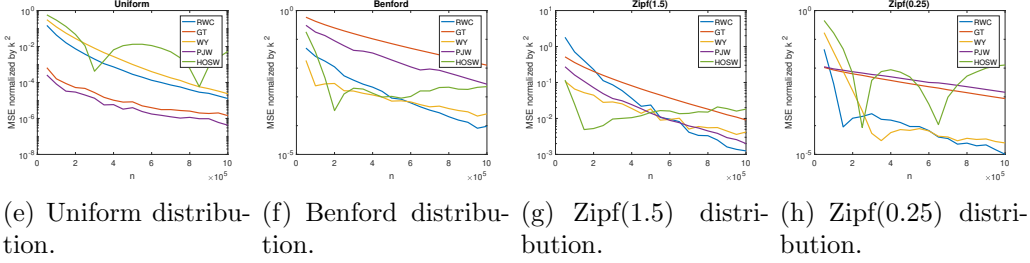
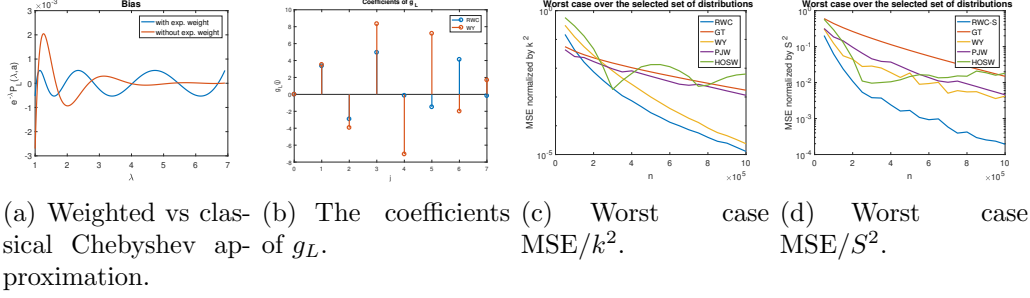


Figure 4.1: Figure (a) illustrates the importance of exponential weights for reducing the bias. The WY estimator aims to minimize the “worst case” risk which arises for small values of λ (close to uniform). As λ increases, the variance dominates the bias (i.e. Figure (o)). Figure (b) gives an example for the coefficients g_L of the RWC and WY estimators. We set $n = k = 10^6$ and $c_0 = 0.558$ for both Figure (a) and (b). The remaining figures plot the MSE of all estimators considered. The y -axis is on the log scale. Figures (m)-(p) show the mean and variance of the estimators for the tested distributions.

biome. Although it is known that the human gut hosts trillions of bacterial cells [36], very little effort has been made to rigorously estimate the actual number of different bacterial genera in the gut. To address this

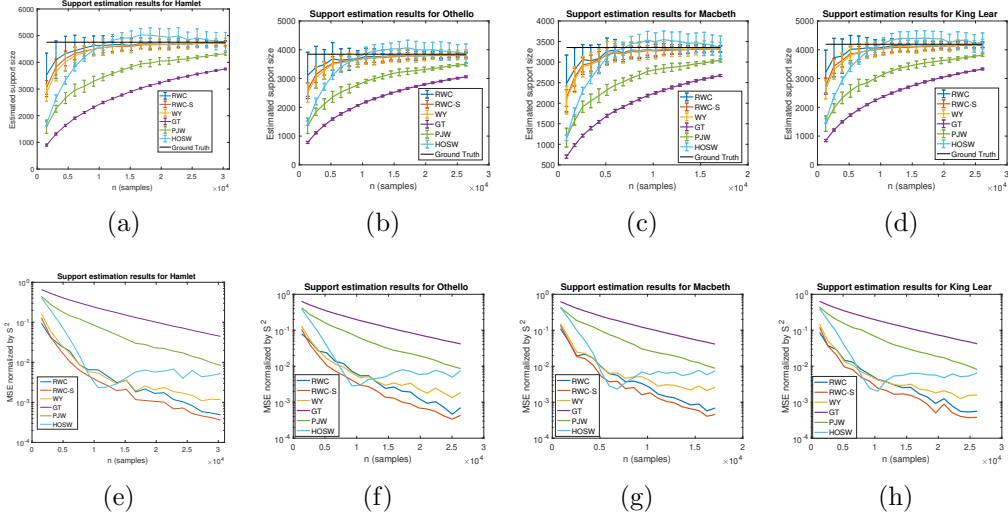


Figure 4.2: The result is obtained over 100 independent trials. The first row of figures shows the mean and standard deviation of the estimators, while the second row of figures shows the MSE normalized by S^2 .

problem, we retrieved 1374 human gut microbiome datasets from the NIH Human Gut Microbiome [30] and the America Gut Microbiome websites (<http://americangut.org>). To determine which bacterial species are presented in the samples, we ran Kraken [37], a taxonomy profiler, on each dataset, using a library of size 8 GB. This lead to $n = 7,415,847$ samples, which we used for obtaining the bacterial genera sample histograms (depicted in Figure 4.3), and to specify $k = n$. The obtained estimates are listed in Table 4.1.




Figure 4.3: The histogram of bacterial genera in the human gut. The plots (a)-(c) correspond to the entire histogram, the histogram for the 100 most frequently encountered genera, and the histogram for the 25 most frequently encountered genera, respectively.

Table 4.1: Estimates for the number of different bacterial genera in the human gut obtained using four different methods. The numbers in brackets correspond to the parameter c_0 .

RWC	RWC-S	PJW	Naive	GT	WY	HOSW
2364	2346	2250	2151	2151	2663	2673

CHAPTER 5

CONCLUSION

In this thesis, we study the problem of support estimation. We proposed an efficient estimator based on regularized weighted Chebyshev approximation. In theory, our estimator is proved to achieve a better risk bound compared to the previous state-of-the-art result. In practice, we show through extensive experiments on both synthetic data and real world data that our estimator outperforms all the previous reported estimators. We also proposed a new bioinformatics application of the support estimation problem: Using metagenomic data from the NIH Human Gut and the American Gut Microbiome Projects, we generate a list of frequencies of bacterial taxa that allows us to estimate the number of bacterial genera.

Future work will involve characterization of the tighter upper bound for our proposed estimator. One will need to analyze our new regularized weighted Chebyshev approximation problem which, to the best of our knowledge, has not been studied in the literature. Another interesting direction would be combine the result with the currently proposed polynomial interpolation techniques that have become rife in this area. Other possible directions include the noisy sample setting (i.e. sticky channel) and finding more applications such that some prior information can be used.

APPENDIX A

PROOF OF THEOREM 3.1

To prove the result, we need to show that $\forall \lambda \geq 6.5L$, $\frac{\partial}{\partial \lambda} g(\mathbf{a}, \lambda) < 0$. The derivative of the first term in g equals

$$\frac{\partial}{\partial \lambda} \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) = \frac{1}{k} \left(\sum_{l=0}^L \left(\frac{l}{\lambda} - 1 \right) e^{-\lambda} a_l^2 \lambda^l l! \right).$$

Clearly, the right-hand side in the above expression is negative for all $\lambda > L$. The second term of the derivative equals

$$\begin{aligned} & \frac{\partial}{\partial \lambda} \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \\ &= 2 \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right) \left(-e^{-\lambda} P_L(\lambda, \mathbf{a}) + e^{-\lambda} \sum_{l=0}^L \frac{l}{\lambda} a_l \lambda^l \right) \\ &= 2e^{-2\lambda} \left(P_L(\lambda, \mathbf{a}) \right) \left(\sum_{l=0}^L \left(\frac{l}{\lambda} - 1 \right) a_l \lambda^l \right). \end{aligned}$$

To analyze the two terms of the derivative, we introduce the vectors $\mathbf{y}, \mathbf{z}, \mathbf{1}$ and the diagonal matrix \mathbf{D} according to

$$\begin{aligned} \mathbf{y} &= (a_0 \lambda^0, a_1 \lambda^1, \dots, a_L \lambda^L)^T, \\ \mathbf{z} &= \left(\left(\frac{0}{\lambda} - 1 \right), \left(\frac{1}{\lambda} - 1 \right), \dots, \left(\frac{L}{\lambda} - 1 \right) \right)^T, \\ \mathbf{1} &= (1, 1, \dots, 1)^T, \\ D_{ii} &= \left(-1 + \frac{i-1}{\lambda} \right) \frac{(i-1)!}{\lambda^{(i-1)}}. \end{aligned}$$

Consequently, we have

$$\begin{aligned}\frac{\partial}{\partial \lambda} \frac{1}{k} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) &= \frac{e^{-\lambda}}{k} \mathbf{y}^T \mathbf{D} \mathbf{y}, \\ \frac{\partial}{\partial \lambda} \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 &= 2e^{-2\lambda} \mathbf{y}^T \mathbf{z} \mathbf{z}^T \mathbf{y} \\ &= e^{-2\lambda} \mathbf{y}^T (\mathbf{z} \mathbf{z}^T + \mathbf{z} \mathbf{1}^T) \mathbf{y}.\end{aligned}$$

Therefore,

$$\frac{\partial}{\partial \lambda} g(\mathbf{a}, \lambda) = e^{-2\lambda} \mathbf{y}^T \left(\frac{e^\lambda}{k} \mathbf{D} + (\mathbf{z} \mathbf{z}^T + \mathbf{z} \mathbf{1}^T) \right) \cdot \mathbf{y}$$

To show that $\frac{\partial}{\partial \lambda} g(\mathbf{a}, \lambda) < 0$ for all polynomials of degree L whenever $\lambda > CL$, we show that the matrix $\left(\frac{e^\lambda}{k} \mathbf{D} + (\mathbf{z} \mathbf{z}^T + \mathbf{z} \mathbf{1}^T) \right)$ is negative-definite whenever $\lambda > CL$, for some constant $C > 0$. It suffices to show that the sum of the maximum eigenvalues of $\frac{e^\lambda}{k} \mathbf{D}$ and $(\mathbf{z} \mathbf{z}^T + \mathbf{z} \mathbf{1}^T)$ is negative, since $\frac{e^\lambda}{k} \mathbf{D}$ is a diagonal matrix. Thus, we turn our attention to determining the maximum eigenvalues of these two matrices. For $\frac{e^\lambda}{k} \mathbf{D}$, the maximum eigenvalue satisfies

$$\frac{e^\lambda}{k} \max_{i \in \{0, 1, \dots, L\}} \left(-1 + \frac{i}{\lambda} \right) \frac{i!}{\lambda^i} \leq -\frac{e^\lambda}{2k} \min_{i \in \{0, 1, \dots, L\}} \frac{i!}{\lambda^i},$$

since for $\lambda > 2L$, one has $(-1 + \frac{i}{\lambda}) \leq -\frac{1}{2}$. When $\lambda > L$, it is clear that $\frac{i!}{\lambda^i}$ is decreasing in i , for $i \in \{0, 1, \dots, L\}$, so that

$$\min_{i \in \{0, 1, \dots, L\}} \frac{i!}{\lambda^i} = \frac{L!}{\lambda^L} \geq \left(\frac{L}{e\lambda} \right)^L.$$

The last inequality is a consequence of Stirling's formula, which asserts that $n! \geq (\frac{n}{e})^n$. Combining the above expressions, we obtain

$$\frac{e^\lambda}{k} \max_{i \in \{0, 1, \dots, L\}} \left(-1 + \frac{i}{\lambda} \right) \frac{i!}{\lambda^i} \leq -\frac{e^\lambda}{2k} \left(\frac{L}{e\lambda} \right)^L.$$

Next, we derive an upper bound on the maximum eigenvalue of the second matrix. The i, j entry of the matrix $(\mathbf{z} \mathbf{z}^T + \mathbf{z} \mathbf{1}^T)$ equals $\frac{i+j-2}{\lambda} - 2$, and all these values are negative when $\lambda > L$. Moreover, it is clear that the matrix of interest has rank equal to 2. Therefore, the matrix has exactly two non-zero

eigenvalues.

Let $\mathbf{A} = -(\mathbf{1z}^T + \mathbf{z1}^T)$. All entries of \mathbf{A} are positive whenever $\lambda > L$. By Gershgorin's theorem, we can upper bound the maximum eigenvalues of the matrix \mathbf{A} by its maximum row sum. It is obvious that the maximum row sum equals

$$2(L+1) - \frac{L(L+1)}{2\lambda}.$$

Moreover, the trace of \mathbf{A} equals

$$2(L+1) - \frac{L(L+1)}{\lambda}.$$

This implies that the minimum eigenvalue of \mathbf{A} is lower bounded by $-\frac{L(L+1)}{2\lambda}$, which directly implies that the maximum eigenvalue of $(\mathbf{1z}^T + \mathbf{z1}^T)$ is upper bounded by $\frac{L(L+1)}{2\lambda}$.

Summing up the two previously derived upper bounds gives

$$h(\lambda) \triangleq -\frac{e^\lambda}{2k} \left(\frac{L}{e\lambda}\right)^L + \frac{L(L+1)}{2\lambda},$$

whenever $\lambda > 2L$. Note that $h(\lambda) < 0$ is equivalent to

$$\begin{aligned} \frac{L(L+1)}{2\lambda} &< \frac{e^\lambda}{2k} \left(\frac{L}{e\lambda}\right)^L \\ &\Leftrightarrow \log(L) + \log(L+1) + \log(k) - L\log(L) + L \\ &< \lambda + \log(\lambda) - L\log(\lambda). \end{aligned} \tag{A.1}$$

The function $\lambda + \log(\lambda) - L\log(\lambda)$ is non-decreasing in λ whenever $\lambda > L$ since

$$\frac{d}{d\lambda}(\lambda + \log(\lambda) - L\log(\lambda)) = 1 - \frac{L-1}{\lambda}.$$

By the definition of $L = \lfloor c_0 \log(k) \rfloor$, we also have $\log(k) \leq \frac{L+1}{c_0}$. Using $\log(x+1) \leq x$, which holds $\forall x \geq 1$. Hence $\forall \lambda > CL$ where $C > 2$, the

sufficient condition for (A.1) to hold is

$$\begin{aligned} \log(L) + L + \frac{L+1}{c_0} - L \log(L) + L \\ < CL + \log(CL) - L \log(CL). \end{aligned}$$

Rearranging terms leads to

$$\left(C - \log(C) - 2 - \frac{1}{c_0} \right) L + \log(C) > \frac{1}{c_0}.$$

Sufficient conditions for the above inequality to hold are $\log(C) \geq \frac{1}{c_0}$ and $(C - \log(C) - 2 - \frac{1}{c_0}) > 0$. The first condition implies $C \geq e^{\frac{1}{c_0}} = 6.0021$, while the second condition holds with $C = 6.5$, for which the first condition is also satisfied. This completes the proof.

APPENDIX B

PROOF OF THEOREM 3.2

The proof consists of two parts. In the first part, we establish the conditions for convergence, while in the second part, we determine the convergence rate.

B.1 Proof of convergence

We start by introducing the relevant terminology. Let $\Pi \subset \mathbb{R}^{L+1}$ be a closed set of parameters, and let f be a continuous functional on Π . Assume that $B \subset \mathbb{R}$ is compact and that $g : \Pi \mapsto \mathcal{C}(B)$ is a continuous mapping from Π into $\mathcal{C}(B)$, where $\mathcal{C}(B)$ is the space of continuous functions over B equipped with the supremum norm $\|\cdot\|_\infty$. For each $D \subset B$ let

$$M(D) = \{\mathbf{c} \in \Pi \mid g(\mathbf{c}, x) \leq 0, x \in D\}$$

denote the set of feasible points of the optimization problem

$$\min f(\mathbf{c}) \text{ over } \mathbf{c} \in M(D).$$

Assuming that $M(D) \neq \emptyset$, let

$$\mu(D) = \inf\{f(\mathbf{c}) \mid \mathbf{c} \in M(D)\},$$

and define the level set

$$\text{Level}(\mathbf{c}_0, D) = \{\mathbf{c} \in \Pi \mid f(\mathbf{c}) \leq f(\mathbf{c}_0)\} \cap M(D).$$

We also make the following two assumptions:

Assumption B.1 (Fine grid). Let $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. There exists a sequence $\{B_i\}$ of compact subsets of B with $B_i \subset B_{i+1}$, $i \in \mathbb{N}_0$, for which

$\lim_{i \rightarrow \infty} h(B_i, B) = 0$, such that

$$h(B_i, B) = \sup_{x \in B} \inf_{y \in B_i} \|x - y\|.$$

Assumption B.2 (Bounded level set). $M(B)$ is nonempty, and there exists a $\mathbf{c}_0 \in M(B)$ such that the level set $\text{Level}(\mathbf{c}_0, B_0)$ is bounded and hence compact in \mathbb{R}^{L+1} .

Theorem B.1 (Convergence of the discretization method, Theorem 2.1 in [34]). Under Assumptions B.1 and B.2, the solution of the discretized problem converges to the optimal solution. More formally, we have

$$\begin{aligned} \mu(B_i) &\leq \mu(B_{i+1}) \leq \mu(B), \forall t \in \mathbb{N}_0 \\ \lim_{i \rightarrow \infty} \mu(B_i) &= \mu(B). \end{aligned}$$

If \mathbf{c}^* is the unique optimal solution of the original problem, and \mathbf{c}_i^* is the optimal solution of the discretization relaxation with grid B_i , then

$$\lim_{i \rightarrow \infty} \|\mathbf{c}^* - \mathbf{c}_i^*\|_2 = 0.$$

It is straightforward to see that our chosen grid is arbitrary fine. Hence, we only need to prove that there exists a \mathbf{c}_0 such that the level set $\text{Level}(\mathbf{c}_0, D)$ is bounded.

Let $\mathbf{c} = (\mathbf{a}; t)$ and note that in our setting, $f(\mathbf{c}) = t$. Rewrite $g(\mathbf{c}, \lambda)$ in matrix form as

$$g(\mathbf{c}, \lambda) = \mathbf{a}^T \mathbf{M}(\lambda) \mathbf{a} + \mathbf{a}^T \Lambda \Lambda^T \mathbf{a} - t,$$

where

$$\Lambda \triangleq e^{-\lambda} (\lambda^0, \lambda^1, \dots, \lambda^L)^T.$$

Note that only a_1, \dots, a_L are allowed to vary since we fixed $a_0 = -1$. Obviously, $\Lambda \Lambda^T$ is positive semi-definite and the previously introduced $\mathbf{M}(\lambda)$ is positive definite for all $\lambda > 0$. Since the constraints on g in (3.7) are positive definite with respect to a_1, \dots, a_L , g is coercive in a_1, \dots, a_L . Furthermore, for any given t , the set of feasible coefficients a_1, \dots, a_L is bounded. Therefore, given a t_0 , the level set $\text{Level}(\mathbf{c}_0, B_0)$ is bounded. This ensures that Assumption B.2 holds for our optimization problem.

Next, we prove the uniqueness of the optimal solution \mathbf{c}^* . Note that proving this result is equivalent to proving the uniqueness of \mathbf{a}^* . Hence, we once again refer to the original minmax formulation of our problem,

$$\inf_{\mathbf{a}:a_0=-1} \sup_{\lambda \in [\frac{n}{k}, 6.5L]} \mathbf{a}^T(\mathbf{M}(\lambda) + \mathbf{\Lambda}\mathbf{\Lambda}^T)\mathbf{a} \triangleq \inf_{\mathbf{a}:a_0=-1} \sup_{\lambda \in [\frac{n}{k}, 6.5L]} h_\lambda(\mathbf{a}). \quad (\text{B.1})$$

Clearly, $\forall \lambda \in [\frac{n}{k}, 6.5L]$, the function $h_\lambda(\mathbf{a})$ is strictly convex since $(\mathbf{M}(\lambda) + \mathbf{\Lambda}\mathbf{\Lambda}^T) \succ 0$, $\forall \lambda \in [\frac{n}{k}, 6.5L]$. Taking the supremum over λ preserves strict convexity since $\forall \theta \in (0, 1)$, giving

$$\begin{aligned} & \sup_{\lambda \in [\frac{n}{k}, 6.5L]} h_\lambda(\theta \mathbf{x} + (1 - \theta)\mathbf{y}) \\ & < \sup_{\lambda \in [\frac{n}{k}, 6.5L]} \theta h_\lambda(\mathbf{x}) + (1 - \theta)h_\lambda(\mathbf{y}) \\ & \leq \sup_{\lambda \in [\frac{n}{k}, 6.5L]} \theta h_\lambda(\mathbf{x}) + \sup_{\lambda' \in [\frac{n}{k}, 6.5L]} (1 - \theta)h_{\lambda'}(\mathbf{y}). \end{aligned}$$

Hence $\sup_{\lambda \in [\frac{n}{k}, 6.5L]} h_\lambda(\mathbf{a})$ is strictly convex, which consequently implies the uniqueness of \mathbf{a}^* and hence \mathbf{c}^* . This proves the convergence result of Theorem 3.2.

B.2 Proof for the convergence rate

In what follows, and for reasons of simplicity, we omit the constraint $a_0 = -1$ in the SIP formulation. The described proof only requires small modifications to accommodate the constraint $a_0 = -1$.

Recall that we used B_d to denote the grid with grid spacing d . In order to use the results in [35], we require the convergence assumption below.

Assumption B.3. Let $\bar{\mathbf{c}}$ be a local minimizer of an SIP. There exists a local solution \mathbf{c}_d of the discretized SIP with grid B_d such that

$$\|\mathbf{c}_d - \bar{\mathbf{c}}\| \rightarrow 0.$$

This assumption is satisfied for the SIP of interest as shown in the first part of the proof.

Assumption B.4. The following hold true:

- There is a neighborhood \bar{U} of $\bar{\mathbf{c}}$ such that the function $\frac{\partial^2}{\partial \lambda^2} g(\mathbf{c}, \lambda)$ is continuous on $\bar{U} \times B$.
- The set B is compact, non-empty and explicitly given as the solution set of a set of inequalities, $B = \{\lambda \in \mathbb{R} | v_i(\lambda) \leq 0, i \in I\}$, where I is a finite index set and $v_i \in C^2(B)$.
- For any $\bar{\lambda} \in B$, the vectors $\frac{\partial}{\partial \lambda} v_i(\bar{\lambda}), i \in \{i \in I | v_i(\bar{\lambda}) = 0\}$ are linearly independent.

Recall that our objective is of the form

$$g(\mathbf{c}, \lambda) = \mathbf{a}^T \mathbf{M}(\lambda) \mathbf{a} + \mathbf{a}^T \Lambda \Lambda^T \mathbf{a} - t,$$

where

$$\begin{aligned} \Lambda &\triangleq e^{-\lambda} (\lambda^0, \lambda^1, \dots, \lambda^L)^T, \quad \mathbf{c} = (\mathbf{a}; t), \\ \mathbf{M}(\lambda) &\triangleq \frac{e^{-\lambda}}{k} \text{Diag}(\lambda^0 0!, \lambda^1 1!, \dots, \lambda^L L!). \end{aligned}$$

It is straightforward to see that the first condition in Assumption B.4 holds. For the second condition, recall that $B = [\frac{n}{k}, 6.5L]$. Hence, the second condition can be satisfied by choosing $I = \{1\}$, $v_1(\lambda) = (\lambda - \frac{n}{k})(\lambda - 6.5L)$. Since we only have one variable v_1 , it is also easy to see that the third condition is met.

Assumption B.5. The set B satisfies Assumption B.4 and all the sets B_d contain the boundary points $\frac{n}{k}, 6.5L$.

This assumption also clearly holds for the grid of choice. Note that it is crucial to include the boundary points for the proof in [35] to be applicable.

Assumption B.6. $\nabla_{\mathbf{c}} g(\mathbf{c}, \lambda)$ is continuous on $\bar{U} \times B$, where \bar{U} is a neighborhood of $\bar{\mathbf{c}}$. Moreover, there exists a vector ξ such that

$$\nabla_{\mathbf{c}} g(\bar{\mathbf{c}}, \lambda)^T \xi \leq -1, \quad \forall \lambda \in B.$$

Note that $\nabla_{\mathbf{c}} g(\mathbf{c}, \lambda) = [\nabla_{\mathbf{a}} g(\mathbf{c}, \lambda); \nabla_t g(\mathbf{c}, \lambda)]$ and

$$\nabla_{\mathbf{a}} g(\mathbf{c}, \lambda) = 2(\mathbf{M}(\lambda) + \Lambda \Lambda^T) \mathbf{a}.$$

Also note that $\forall \lambda \in B$, $\mathbf{M}(\lambda) + \Lambda\Lambda^T$ is positive definite. Hence choosing ξ to be colinear with and of the same direction as $[-\mathbf{a}^T \ 1]^T$, as well as of sufficiently large norm, will allow us to satisfy the inequality

$$\nabla_{\mathbf{c}} g(\bar{\mathbf{c}}, \lambda)^T \xi \leq -1, \quad \forall \lambda \in B.$$

Hence, Assumption B.6 holds as well. The next results follow from the above assumptions and observations, and the results in [35].

Lemma B.1 (Corollary 1 in [35]). Let t_d be the optimal objective value of the discretized SIP used for support estimation with the grid B_d , and let t^* be the optimal objective value for the original SIP. Since Assumptions B.3, B.4, B.5 and B.6 hold, then for some $c_3 > 0$ and d sufficiently small, we have

$$0 \leq t^* - t_d \leq c_3 d^2.$$

Consequently, $t_d \rightarrow t^*$ with a convergence rate of $O(d^2)$.

Lemma B.2 (Theorem 2 in [35]). Assume that all assumptions in Lemma B.1 hold. If there exists a constant $c_4 > 0$ such that

$$t - \bar{t} \geq c_4 \|\mathbf{c} - \bar{\mathbf{c}}\|, \quad \forall \mathbf{c} \in M(B) \cap \bar{U},$$

then for sufficiently small d and $\sigma > 0$ we have

$$\|\mathbf{c}_d - \bar{\mathbf{c}}\| \leq \sigma d^2.$$

This result implies that if $\bar{\mathbf{c}}$ is also a strict minimum of order one, then the solution of the discretized SIP converges to that of the the original SIP with rate $O(d^2)$. Combining these results completes the proof.

APPENDIX C

THEORETICAL RESULTS SUPPORTING REMARK 3.1

The result described in the main text follows from Theorem 6.2 in [27], originally proved in [38, 39] and [40].

Theorem C.1 (Theorem 6.2 from [27]). Let $W(x) = \exp(-Q(x))$ be a weight function, where $Q : \mathbb{R} \mapsto [0, \infty)$ is even, convex, diverging for $x \rightarrow \infty$, and such that

$$0 = Q(0) < Q(x), \forall x \neq 0.$$

Then, for any polynomial $P(x)$ of degree $\leq L$, not identical to zero, one has

$$\begin{aligned} \sup_{x \in \mathbb{R}} |P(x)W(x)| &= \sup_{x \in [-M_L, M_L]} |P(x)W(x)|, \\ \sup_{x \in \mathbb{R} \setminus [-M_L, M_L]} |P(x)W(x)| &< \sup_{x \in [-M_L, M_L]} |P(x)W(x)|. \end{aligned}$$

Here, M_L stands for the *Mhaskar-Rakhmanov-Saff* (MSF) number, which is the smallest positive root of the integral equation

$$L = \frac{2}{\pi} \int_0^1 \frac{M_L t Q'(M_L t)}{\sqrt{1-t^2}} dt. \quad (\text{C.1})$$

In our setting, the weight equals $\exp(-x)$. Solving (C.1) gives us an MSF number equal to $M_L = \frac{\pi}{2}L$. Thus, we can restrict our optimization interval to $[\frac{n}{k}, \frac{\pi}{2}L + \frac{n}{k}]$. If there is no regularization term, the optimal interval reduces to $[\frac{n}{k}, \frac{\pi}{2}L + \frac{n}{k}]$.

APPENDIX D

CONSTRUCTION OF THE RWC-S ESTIMATOR

We introduce the optimization problem needed for minimizing the risk $\mathbb{E} \left(\frac{S - \hat{S}}{S} \right)^2$. Poissonization arguments once again establish that

$$\begin{aligned} \mathbb{E} \left(\frac{S - \hat{S}}{S} \right)^2 &= \frac{1}{S^2} \left\{ \sum_{i:\lambda_i > 0} \left(\sum_{l=0}^L e^{-\lambda_i} a_l^2 \lambda_i^l l! \right) \right. \\ &\quad \left. + \sum_{i \neq j: \lambda_i \lambda_j > 0} \left(e^{-\lambda_i} P_L(\lambda_i, \mathbf{a}) \right) \left(e^{-\lambda_j} P_L(\lambda_j, \mathbf{a}) \right) \right\}. \end{aligned}$$

Taking the supremum over D_k , one can further upper bound the risk as

$$\begin{aligned} &\leq \sup_{\lambda: \lambda_i \in [\frac{n}{k}, n]} \frac{1}{S^2} \left\{ \sum_{i:\lambda_i > 0} \left(\sum_{l=0}^L e^{-\lambda_i} a_l^2 \lambda_i^l l! \right) \right. \\ &\quad \left. + \sum_{i \neq j: \lambda_i \lambda_j > 0} \left(e^{-\lambda_i} P_L(\lambda_i, \mathbf{a}) \right) \left(e^{-\lambda_j} P_L(\lambda_j, \mathbf{a}) \right) \right\} \end{aligned} \tag{D.1}$$

$$\begin{aligned} &\leq \sup_{\lambda \in [\frac{n}{k}, n]} \left\{ \frac{1}{S} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\} \\ &\leq \sup_{\lambda \in [\frac{n}{k}, n]} \left\{ \frac{1}{\hat{S}_c} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\}, \end{aligned} \tag{D.2}$$

where the last inequality is due to the fact that $\hat{S}_c \leq S$. Note that the only difference between (D.2) and (3.1) is in terms of changing $1/k$ to $1/\hat{S}_c$ in the first term. In view of Theorem 3.1, (D.2) is optimized by the solution of the following problem:

$$\begin{aligned}
& \min_{t, a_1, \dots, a_L} t \quad \text{subject to} \\
& \left\{ \frac{1}{\widehat{S}_c} \left(\sum_{l=0}^L e^{-\lambda} a_l^2 \lambda^l l! \right) + \left(e^{-\lambda} P_L(\lambda, \mathbf{a}) \right)^2 \right\} \leq t \quad (\text{D.3}) \\
& \forall \lambda \in \text{Grid}\left(\left[\frac{n}{k}, 6.5L\right], s\right), \text{ with } a_0 = -1.
\end{aligned}$$

REFERENCES

- [1] R. A. Fisher, A. S. Corbet, and C. B. Williams, “The relation between the number of species and the number of individuals in a random sample of an animal population,” *The Journal of Animal Ecology*, pp. 42–58, 1943.
- [2] B. Efron and R. Thisted, “Estimating the number of unseen species: How many words did Shakespeare know?” *Biometrika*, vol. 63, no. 3, pp. 435–447, 1976.
- [3] S. Raskhodnikova, D. Ron, A. Shpilka, and A. Smith, “Strong lower bounds for approximating distribution support size and the distinct elements problem,” *SIAM Journal on Computing*, vol. 39, no. 3, pp. 813–842, 2009.
- [4] Z. Bar-Yossef, T. Jayram, R. Kumar, D. Sivakumar, and L. Trevisan, “Counting distinct elements in a data stream,” in *International Workshop on Randomization and Approximation Techniques in Computer Science*. Springer, 2002, pp. 1–10.
- [5] M. Charikar, S. Chaudhuri, R. Motwani, and V. Narasayya, “Towards estimation error guarantees for distinct values,” in *Proceedings of the nineteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2000, pp. 268–279.
- [6] I. Chien, C. Pan, and O. Milenkovic, “Query k-means clustering and the double Dixie cup problem,” in *Advances in Neural Information Processing Systems*, 2018, pp. 6650–6659.
- [7] H. Ashtiani, S. Kushagra, and S. Ben-David, “Clustering with same-cluster queries,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3216–3224.
- [8] I. Chien, H. Zhou, and P. Li, “HS² : Active learning over hypergraphs,” in *International Conference on Artificial Intelligence and Statistics*, 2019.

- [9] A. Orlitsky, N. P. Santhanam, and J. Zhang, “Always Good Turing: Asymptotically optimal probability estimation,” *Science*, vol. 302, no. 5644, pp. 427–431, 2003.
- [10] J. Acharya, H. Das, A. Jafarpour, A. Orlitsky, and S. Pan, “Competitive closeness testing,” in *Proceedings of the 24th Annual Conference on Learning Theory*, 2011, pp. 47–68.
- [11] L. Paninski, “Estimation of entropy and mutual information,” *Neural Computation*, vol. 15, no. 6, pp. 1191–1253, 2003.
- [12] J. Bunge and M. Fitzpatrick, “Estimating the number of species: A review,” *Journal of the American Statistical Association*, vol. 88, no. 421, pp. 364–373, 1993.
- [13] Z. Bar-Yossef, R. Kumar, and D. Sivakumar, “Sampling algorithms: Lower bounds and applications,” in *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*. ACM, 2001, pp. 266–275.
- [14] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, “Testing that distributions are close,” in *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*. IEEE, 2000, pp. 259–269.
- [15] M. R. Nelson, D. Wegmann, M. G. Ehm, D. Kessner, P. S. Jean, C. Verzilli, J. Shen, Z. Tang, S.-A. Bacanu, D. Fraser et al., “An abundance of rare functional variants in 202 drug target genes sequenced in 14,002 people,” *Science*, p. 1217876, 2012.
- [16] A. Keinan and A. G. Clark, “Recent explosive human population growth has resulted in an excess of rare genetic variants,” *Science*, vol. 336, no. 6082, pp. 740–743, 2012.
- [17] P. Valiant and G. Valiant, “Estimating the unseen: Improved estimators for entropy and other properties,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2157–2165.
- [18] Y. Wu and P. Yang, “Chebyshev polynomials, moment matching, and optimal estimation of the unseen,” *The Annals of Statistics*, vol. 47, no. 2, pp. 857–883, 2019.
- [19] J. Acharya, H. Das, A. Orlitsky, and A. T. Suresh, “A unified maximum likelihood approach for estimating symmetric properties of discrete distributions,” in *International Conference on Machine Learning*, 2017, pp. 11–21.
- [20] Y. Wu and P. Yang, “Sample complexity of the distinct elements problem,” *Mathematical Statistics and Learning*, vol. 1, no. 1, pp. 37–72, 2018.

- [21] D. S. Pavlichin, J. Jiao, and T. Weissman, “Approximate profile maximum likelihood,” *arXiv preprint arXiv:1712.07177*, 2017.
- [22] Y. Han, J. Jiao, and T. Weissman, “Local moment matching: A unified methodology for symmetric functional estimation and distribution estimation under wasserstein distance,” *arXiv preprint arXiv:1802.08405*, 2018.
- [23] H. Yi, A. Orlitsky, A. T. Suresh, and Y. Wu, “Data amplification: A unified and competitive approach to property estimation,” in *Advances in Neural Information Processing Systems*, 2018, pp. 8834–8843.
- [24] A. Orlitsky, N. P. Santhanam, K. Viswanathan, and J. Zhang, “On modeling profiles instead of values,” in *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2004, pp. 426–435.
- [25] J. Aldrich et al., “R.A. Fisher and the making of maximum likelihood 1912-1922,” *Statistical Science*, vol. 12, no. 3, pp. 162–176, 1997.
- [26] A. Orlitsky, A. T. Suresh, and Y. Wu, “Optimal prediction of the number of unseen species,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 47, pp. 13 283–13 288, 2016.
- [27] D. S. Lubinsky, “A survey of weighted polynomial approximation with exponential weights,” *Surveys in Approximation Theory*, vol. 3, pp. 1–105, 2007.
- [28] J. C. Mason and D. C. Handscomb, *Chebyshev Polynomials*. Chapman and Hall/CRC, 2002.
- [29] M. López and G. Still, “Semi-infinite programming,” *European Journal of Operational Research*, vol. 180, no. 2, pp. 491–518, 2007.
- [30] J. Peterson, S. Garges, M. Giovanni, P. McInnes, L. Wang, J. A. Schloss, V. Bonazzi, J. E. McEwen, K. A. Wetterstrand, C. Deal et al., “The NIH human microbiome project,” *Genome Research*, vol. 19, no. 12, pp. 2317–2323, 2009.
- [31] A. F. Timan, *Theory of Approximation of Functions of a Real Variable*. Elsevier, 2014, vol. 34.
- [32] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge university press, 2004.
- [33] R. Reemtsen and S. Görner, “Numerical methods for semi-infinite programming: A survey,” in *Semi-infinite Programming*. Springer, 1998, pp. 195–275.

- [34] R. Reemtsen, “Discretization methods for the solution of semi-infinite programming problems,” *Journal of Optimization Theory and Applications*, vol. 71, no. 1, pp. 85–103, 1991.
- [35] G. Still, “Discretization in semi-infinite programming: the rate of convergence,” *Mathematical Programming*, vol. 91, no. 1, pp. 53–69, 2001.
- [36] R. Sender, S. Fuchs, and R. Milo, “Revised estimates for the number of human and bacteria cells in the body,” *PLoS Biology*, vol. 14, no. 8, p. e1002533, 2016.
- [37] D. E. Wood and S. L. Salzberg, “Kraken: Ultrafast metagenomic sequence classification using exact alignments,” *Genome Biology*, vol. 15, no. 3, p. R46, 2014.
- [38] H. Mhaskar and E. Saff, “Extremal problems for polynomials with exponential weights,” *Transactions of the American Mathematical Society*, vol. 285, no. 1, pp. 203–234, 1984.
- [39] H. Mhaskar and E. Saff, “Where does the sup norm of a weighted polynomial live?” *Constructive Approximation*, vol. 1, no. 1, pp. 71–91, 1985.
- [40] E. A. Rakhmanov, “On asymptotic properties of polynomials orthogonal on the real axis,” *Sbornik: Mathematics*, vol. 47, no. 1, pp. 155–193, 1984.