

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

12-2021

Beyond smoothness : Incorporating low-rank analysis into nonparametric density estimation

Rob VANDERMEULEN

Antoine LEDENT

Singapore Management University, aledent@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Theory and Algorithms Commons](#)

Citation

VANDERMEULEN, Rob and LEDENT, Antoine. Beyond smoothness : Incorporating low-rank analysis into nonparametric density estimation. (2021). *Advances in Neural Information Processing Systems (NeurIPS 2021): December 7-10, Virtual: Proceedings*. 34, 12180-12193.

Available at: https://ink.library.smu.edu.sg/sis_research/7205

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Beyond Smoothness: Incorporating Low-Rank Analysis into Nonparametric Density Estimation

Robert A. Vandermeulen

Machine Learning Group
Technische Universität Berlin
Berlin, Germany
vandermeulen@tu-berlin.de

Antoine Ledent

Machine Learning Group
Technische Universität Kaiserslautern
Kaiserslautern, Germany
ledent@cs.uni-kl.de

Abstract

The construction and theoretical analysis of the most popular universally consistent nonparametric density estimators hinge on one functional property: smoothness. In this paper we investigate the theoretical implications of incorporating a multi-view latent variable model, a type of low-rank model, into nonparametric density estimation. To do this we perform extensive analysis on histogram-style estimators that integrate a multi-view model. Our analysis culminates in showing that there exists a universally consistent histogram-style estimator that converges to any multi-view model with a finite number of Lipschitz continuous components at a rate of $\tilde{O}(1/\sqrt[3]{n})$ in L^1 error. In contrast, the standard histogram estimator can converge at a rate slower than $1/\sqrt[4]{n}$ on the same class of densities. We also introduce a new nonparametric latent variable model based on the Tucker decomposition. A rudimentary implementation of our estimators experimentally demonstrates a considerable performance improvement over the standard histogram estimator. We also provide a thorough analysis of the sample complexity of our Tucker decomposition-based model and a variety of other results. Thus, our paper provides solid theoretical foundations for extending low-rank techniques to the nonparametric setting.

1 Introduction

Nonparametric density estimators are density estimators capable of estimating a density p while making few to no assumptions on p . Two commonly used nonparametric density estimators include the histogram estimator and the kernel density estimator (KDE). A common characteristic of these two estimators is that they make estimation tractable via a hyperparameter that relates to smoothness, namely bin width and bandwidth. Large bin width or bandwidth allows for good control of estimation error* at the cost of increased approximation error via large bin volume for the histogram and smoothing for KDEs. Selection of this parameter is crucial for estimator performance. In fact a recent survey on bandwidth selection for KDEs found at least 30 methods for setting this value along with a few surveys dedicated to the topic [18].

While nonparametric density estimation has been shown to be effective for many tasks, it has been observed empirically that estimator performance typically declines as data dimensionality increases, a manifestation of the curse of dimensionality. For the histogram and KDE this phenomenon has concrete mathematical analogs. For example, these estimators are only universally consistent† if $n \rightarrow \infty$ and $h \rightarrow 0$, with $nh^d \rightarrow \infty$, where n is the number of samples, d is the data dimension,

*For an estimator V restricted space \mathcal{P} of densities \mathcal{P} , the *estimation error* refers to the difference between $\|V - p\|$ and $\min_{q \in \mathcal{P}} \|p - q\|$, where p is the target density. This is similar to estimator variance.

†A density estimator is *universally consistent* if it asymptotically recovers *any* density.

and h is the bin width for the histogram and bandwidth parameter for the KDE [16]. One may then wonder whether there exists some other way to constrain model capacity so as to alleviate this exponential dependence on dimension.

In this paper we theoretically analyze the advantages of including a constraint akin to matrix/tensor rank in a nonparametric density estimator. We do so by analyzing histogram-style estimators (estimators that output a density that are piecewise constant on bins defined by a grid) that are enforced to have a low-rank PARAFAC or Tucker decomposition. Enforcing this low-rank constraint in the histogram estimator allows for much faster rates on $h \rightarrow 0$, while still controlling estimation error. This can remove the exponential penalty on rate of convergence that dimensionality produces with the standard histogram estimator.

The bulk of this work focuses on analysis of hypothetical estimators that, while offering large statistical advantages over the standard histogram estimator, are not computationally tractable. At the end of this work we include experiments demonstrating that low-rank histograms, computed using an off-the-shelf nonnegative tensor factorization library [27], consistently outperform the standard histogram estimator.

1.1 Density Models

Here we introduce the low-rank models that will be used to constrain our estimators. The first is a *multi-view model*. A multi-view model p is a finite mixture $p = \sum_{i=1}^k w_i f_i$ whose components are separable densities $f_i = p_{i,1} \otimes p_{i,2} \otimes \cdots \otimes p_{i,d}$ [‡] [21, 1, 41, 40]. A multi-view model has the following form

$$p(x_1, x_2, \dots, x_d) = \sum_{i=1}^k w_i p_{i,1}(x_1) p_{i,2}(x_2) \cdots p_{i,d}(x_d). \quad (1)$$

A multi-view model with one component is typically called a *naive Bayes model*. For the estimators we propose, the component marginals $p_{i,j}$ will have the form of one-dimensional histograms and the number of components k will be limited to restrict estimator capacity. When $k = 1$ the model is equivalent to a naive Bayes model, $p(x_1, x_2, \dots, x_d) = p_1(x_1) p_2(x_2) \cdots p_d(x_d)$. When the component marginals $p_{i,j}$ are histograms increasing k expands the set of potential estimates from naive Bayes models when $k = 1$ to all possible histogram estimates.

The models in this paper are motivated by nonnegative tensor factorizations so we term them generally as *nonparametric nonnegative tensor factorization* (NNTF) *models*. The previous model was related to nonnegative PARAFAC [38]. Our second model is based on the nonnegative Tucker decomposition [26]. A density in this model utilizes d collections, $\mathcal{F}_1, \dots, \mathcal{F}_d$, of k one-dimensional densities, $\mathcal{F}_i = \{p_{i,1}, \dots, p_{i,k}\}$, and some probability measure that randomly selects one density from each \mathcal{F}_i . This measure can be represented by a tensor $W \in \mathbb{R}^{k \times d}$ where the probability of selecting $(p_{1,i_1}, \dots, p_{d,i_d})$ from $\mathcal{F}_1 \times \cdots \times \mathcal{F}_d$ is W_{i_1, \dots, i_d} . To sample from this model we first randomly select the marginal distributions $p_{1,i_1}, \dots, p_{d,i_d}$ according to W , and an observation is sampled according to the d -dimensional distribution $p_{1,i_1} \otimes p_{2,i_2} \otimes \cdots \otimes p_{d,i_d}$. The density of this model is

$$p(x_1, x_2, \dots, x_d) = \sum_{i_1=1}^k \cdots \sum_{i_d=1}^k W_{i_1, \dots, i_d} p_{1,i_1}(x_1) p_{2,i_2}(x_2) \cdots p_{d,i_d}(x_d). \quad (2)$$

We are unaware of previous works investigating this model for general probability distributions so we will simply term it the *Tucker model*. Again we will investigate estimators where the component marginals are one-dimensional histograms and small k corresponds to reduced model capacity. We remark that Tucker decompositions typically have a rank *vector* $[k_1, \dots, k_d]$ which for our model would mean that each \mathcal{F}_i contains k_i vectors and W would lie in $\mathbb{R}^{k_1 \times \cdots \times k_d}$. This sort of rank restriction could be used in our methods however, for simplicity, we just set $k_1 = k_2 = \cdots = k_d \triangleq k$.

[‡]For two functions f, g let $f \otimes g : (x, y) \mapsto f(x)g(y)$. This is analogous to the tensor product of L^2 functions.

1.2 Overview of Results

The principal contribution of this paper is to analyze the advantage of incorporating rank restriction into density estimation. In Section 2.1 we precisely introduce the multi-view histogram and a histogram based on the Tucker model. With these models we investigate how quickly we can let bin width h go to zero and rank k grow, in relation to the amount of data n , while still being able to control estimation error and select an estimator that is nearly optimal for the allowable set. For the multi-view histogram we show that one can control estimation error so long as k/h is asymptotically dominated by n and for the Tucker histogram we need $k/h + k^d$ to be asymptotically dominated by n (we are omitting logarithmic factors here for convenience). This stands in stark contrast to standard space of histogram estimators which requires $1/h^d$ to be asymptotically dominated by n . We furthermore show that these estimators are universally consistent and that these rates cannot be significantly improved.

For a second style of analysis we provide finite-sample bounds on the convergence of NNTF histogram estimators. We then construct a class of universally consistent density estimators that converge at rate $\tilde{O}(1/\sqrt[n]{n})^\S$ on all densities that are a multi-view model with Lipschitz continuous component marginals. Note that the NNTF histogram estimators we construct do not require the target density to be an NNTF model to function well. Our estimators will select a good estimator so long as there exists any NNTF histogram that approximates the target density well; so our hypothetical estimators “fail elegantly” in some sense. We further show that the standard histogram can converge at a rate $\omega(1/\sqrt[n]{n})^\P$ on the same class of densities. In Section 3 we experimentally investigate the efficacy of using NNTF histograms on real-world data. In lieu of the computationally intractable methods we investigated in our theoretical analyses, we use an existing low-rank nonnegative Tucker factorization algorithm to fit an NNTF histogram to data. Surprisingly even this method outperforms the standard histogram estimator with very high statistical significance. Lastly we mention that this paper is an extension of [45] and contains a fair amount of overlap with that text.

1.3 Previous Work

Nonparametric density estimation has been extensively studied with the histogram estimator and KDE being some of the most well-known methods. There do exist, however, alternative methods for density estimation, e.g. the *forest density estimator* [28] and *k-nearest neighbor density estimator* [29]. The L^1 , L^2 , and L^∞ convergence of the histogram and KDE has been studied extensively [16, 10, 43, 19]. The KDE is generally regarded as the superior density estimator, with some mathematical justification [16, 39]. Numerous modifications and extensions of the KDE have been proposed including utilizing variable bandwidth [42], robust KDEs [24, 44, 47], methods for enforcing support boundary constraints [37], and a supervised variant [46]. Finally we mention [25] that demonstrated that uniform convergence of a KDE to its population estimate suffered when the intrinsic dimension of the data was lower than the ambient dimension, a phenomenon seemingly at odds with the curse of dimensionality.

For our review of NNTF models we also include a general review of tensor/matrix factorizations since both can be viewed being low-rank models. In particular, for the multi-view model we have the following analogy

$$\sum_{i=1}^k w_i p_{i,1}(x_1) p_{i,2}(x_2) \cdots p_{i,d}(x_d) \sim \sum_{i=1}^k \lambda_i \mathbf{v}_{i,1} \otimes \mathbf{v}_{i,2} \otimes \cdots \otimes \mathbf{v}_{i,d}. \quad (3)$$

A great deal of work has gone into leveraging low-rank assumptions to improve matrix estimation, particularly in the field of *compressed sensing* [11, 33]. The most basic version of compressed sensing is concerned with estimating a “tall” vector x from a “short” vector $y \triangleq Ax$ where A is a known “short and fat” matrix. One can recover y if it is sparse and A satisfies a property known as the restricted isometry property (RIP) [48, 8]. These methods can be extended to the estimation of matrices when x is a low-rank matrix [33, 30, 31]. In this extension A is an order-3 tensor which

^{\S} $f_n \in \tilde{O}(g_n) \iff \exists k \text{ s.t. } f_n \in O(g_n \log^k g_n)$

^{\P} $f_n \in \omega(g_n) \iff |f_n/g_n| \rightarrow \infty$ in probability. In d -dimensional space the standard histogram can converge slower than $1/\sqrt[n]{n}$.

acts as a linear operator on x and satisfies an adjusted form of RIP. RIP commonly arises from matrices and tensors whose entries are random. Because of this compressed sensing techniques are useful in settings where one wants to estimate x from random linear transforms of x . For example, in matrix completion observing the (i, j) -th entry of x can be represented as an inner product of x with an indicator matrix $\mathbf{e}_{i,j}$, i.e. $\langle x, \mathbf{e}_{i,j} \rangle_F$. Thus the random observed indices $(i_1, j_1), (i_2, j_2), \dots$ can be represented as random matrices $\mathbf{e}_{i_1, j_1}, \mathbf{e}_{i_2, j_2}, \dots$ which are then stacked into A . Now A is a random linear operator that is applied to x to represent the observation of random entries of x and the methods of compressed sensing can be used to recover x . Compressed sensing has also proven useful for multivariate regression and autoregressive models [30, 31]. Such techniques don't appear to be extensible to histogram estimation due to the lack of a linear sampling scheme.

General matrix/tensor factorization, including nonnegative matrix/tensor factorizations, has been extensively studied despite being inherently difficult due to non-convexity. The works [12, 5] present potential theoretical grounds for avoiding the computational difficulties of nonnegative matrix factorization. One notable approach to tensor factorization is to assume, in the tensor representation in (3), that $d \geq 3$ and the collections of vectors $\mathbf{v}_{1,j}, \dots, \mathbf{v}_{k,j}$ are linearly independent for all j . Under this assumption we are guaranteed that the factorization (3) is unique [1]. In [4] the authors present a method for recovering this factorization efficiently and demonstrate its utility for a variety of tasks. This work was extended in [40] to recover a multi-view KDE satisfying an analogous linear independence assumption and theoretically analyze the estimator's convergence to the true low-rank components. In [40] the authors investigate the sample complexity of their estimator but do not demonstrate that their technique has potential for improving rates for nonparametric density estimation in general. In [41] it was observed that using low-rank embeddings can improve density estimation. A multi-view histogram was investigated in [20] where the authors present an identifiability result and algorithm for recovering latent factors of the distribution. The related works [2, 3] consider a low-rank characteristic function as an approach to improving nonparametric density estimation. Though earlier works have observed that a low-rank approach improves nonparametric density estimation [41, 40, 20, 2, 3], we are the first to demonstrate this through theoretical analysis of sample complexity. Finally we note that the Tucker decomposition has been utilized in Bayesian statistics [36]. We are unaware of any literature on factoring functions in $\mathbb{R}^d \rightarrow \mathbb{R}$ in a Tucker-inspired as we do in (2).

2 Theoretical Results

In this section we mathematically demonstrate that histogram estimators can achieve greater performance by restricting to NNTF models and using a proper procedure to select a representative from these using data. To simplify analysis we will only consider densities on the unit cube $[0, 1]^d$ and analyze the number of bins per dimension b which is the inverse of the bin width, i.e. $b = 1/h$. To state our results precisely we must introduce a fair amount of notation.

2.1 Notation

We will denote the L^1 and L^2 Lebesgue space norms via a 1 or 2 subscript. Let \mathcal{D}_d be the set of all densities on $[0, 1]^d$. By *density* we mean probability measures that are absolutely continuous with respect to the d -dimensional Lebesgue measure on $[0, 1]^d$. We define a *probability vector* or *probability tensor* to simply mean a vector or tensor whose entries are nonnegative and sum to one. Let Δ_b denote the set of probability vectors in \mathbb{R}^b and $\mathcal{T}_{d,b}$ the set of probability tensors in $\mathbb{R}^{b \times d}$ ^{||}. The product symbol \prod will always mean the standard outer product, e.g. set product^{**} or tensor product, when the multiplicands are not real numbers^{††}. The natural numbers \mathbb{N} will always denote *positive* integers. For any natural number b let $[b] = \{1, \dots, b\}$. We will let $\mathbb{1}$ be the indicator function and Conv be the convex hull. Later we will use projection operator where $\text{Proj}_S x \triangleq \arg \min_{s \in S} \|x - s\|_2$; for every instance in this work this projection is unique.

^{||} $\mathbb{R}^{b \times d}$ is the set of $\underbrace{b \times \dots \times b}_{d \text{ times}}$ tensors. For example $\mathbb{R}^{b \times 2}$ is the set of $b \times b$ matrices.

^{**}For sets S_1, \dots, S_d we have $\prod_{i=1}^d S_i = S_1 \times \dots \times S_d = \{(s_1, \dots, s_d) : s_i \in S_i \forall i\}$.

^{††}For functions f_1, \dots, f_d then $\prod_{i=1}^d f_i = f_1 \otimes \dots \otimes f_d : (x_1, \dots, x_d) \mapsto \prod_{i=1}^d f_i(x_i)$.

We will now construct the space of histograms on $[0, 1)^d$. We begin with one-dimensional histograms, which will serve as the $p_{i,j}$ terms in (1) or (2). We define $h_{1,b,i}$ with $i \in [b]$ to be the one-dimensional histogram where all weight is allocated to the i th bin. Formally we define this as

$$h_{1,b,i}(x) \triangleq b \mathbb{1}\left(\frac{i-1}{b} \leq x < \frac{i}{b}\right).$$

Note that this is a valid density due to the leading b coefficient. We use these to construct higher-dimensional histograms. For a multi-index $A \in [b]^d$, let

$$h_{d,b,A} \triangleq \prod_{i=1}^d h_{1,b,A_i},$$

i.e. the d -dimensional histogram with b bins per dimension whose entire density is allocated to the bin indexed by A . Finally we define $\Lambda_{d,b,A}$ to be the support of $h_{d,b,A}$, i.e. the “bins” of a histogram estimator,

$$\Lambda_{d,b,A} \triangleq \prod_{i=1}^d \left[\frac{A_i-1}{b}, \frac{A_i}{b} \right).$$

For a dataset $\mathcal{X} = (X_1, \dots, X_n)$ in $[0, 1)^d$, the standard histogram estimator is

$$H_{d,b}(\mathcal{X}) \triangleq \frac{1}{n} \sum_{i=1}^n \sum_{A \in [b]^d} h_{d,b,A} \mathbb{1}(X_i \in \Lambda_{d,b,A}).$$

Let $\mathcal{H}_{d,b} \triangleq \text{Conv}\left(\left\{h_{d,b,A} \mid A \in [b]^d\right\}\right)$, the set of all d -dimensional histograms with b bins per dimension. Let $\mathcal{H}_{d,b}^k$ be the set of histograms with at most k separable components, i.e.

$$\mathcal{H}_{d,b}^k \triangleq \left\{ \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j} \mid w \in \Delta_k, p_{i,j} \in \mathcal{H}_{1,b} \right\}. \quad (4)$$

We will refer to elements in this space as *multi-view histograms*. Elements in this space have the same form as (1) in Section 1.1. Similarly we define the space of *Tucker histograms* to be

$$\tilde{\mathcal{H}}_{d,b}^k = \left\{ \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} \mid W \in \mathcal{T}_{d,k}, p_{i,j} \in \mathcal{H}_{1,b} \right\}.$$

These have the same form as (2) in Section 1.1.

We emphasize that the collections of densities $\mathcal{H}_{d,b}^k$ and $\tilde{\mathcal{H}}_{d,b}^k$ are primary objects of interest in this paper. The results we present are concerned with finding good density estimators restricted to these sets as k and b vary.

2.2 Estimator Theoretical Results

We present two approaches to the analysis of NNTF histogram estimators. All proofs and additional results are contained in the appendix.

First we provide an asymptotic analysis of NNTF histogram estimators in terms of estimation error control: how fast can we let b and k grow, with respect to n , while still controlling estimation error over all densities? An advantage of this analysis is that we can demonstrate that these rates are approximately optimal (up to logarithmic terms).

For our second approach we present finite-sample bound analysis. In this analysis we first present a distribution-dependent bound that, for an estimator restricted to $\mathcal{H}_{d,b}^k$, depends on n, b, k and $\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1$ where p is the data generating “target” distribution. We follow this up with distribution-free bounds.

Distribution-free bounds for nonparametric density estimation require that the target distribution belong to a well-behaved class of distributions (such as Sobolev or Hölder classes) to enable bounding of the approximation error [43]. Our distribution-free finite-sample analysis assumes that the target density is a multi-view model whose component marginals are Lipschitz continuous. We construct an estimator that converges at a rate of approximately $1/\sqrt[3]{n}$ on this class of densities, independent of d . For comparison we show that the standard histogram estimator can converge at a rate worse than $1/\sqrt[3]{n}$ on this same class of densities. We mention again that this estimator shouldn't fail catastrophically when these distributional assumptions aren't exactly met so long as its approximation error, $\min_{p \in \mathcal{H}_{d,b}^k} \|p - q\|_1$, isn't large. For brevity, and because the results are virtually direct analogues of their multi-view histogram counterparts, we reserve all finite-sample results for Tucker histogram for the appendix.

Main Technical Tools Our results rely on finding L^1 ε -coverings of the spaces of NNTF histograms and using Theorem 3.7 from [6] to select a good representative from that collection. The aforementioned theorem is a slight extension of Theorem 6.3 in [10] and is essentially proven in [49]. As was mentioned in [6], the application of these results typically does not yield a computationally practical algorithm. Likewise our results are simply meant to highlight the potential of NNTF models and are not practically implementable as is.

2.2.1 Asymptotic Error Control

The following theorem states that one can control the estimation error of multi-view histograms with k components and b bins per dimension so long as $n \sim bk$ (omitting logarithmic factors). Recall that the standard histogram estimator requires $n \sim b^d$, so we have removed the exponential dependence of bin rate on dimensionality. Here and elsewhere the \sim symbol is not a precise mathematical statement but rather signifies that the two values should be of the same order in a general sense. In the following b and k are functions of n so the space of histograms changes as one acquires more data.

Theorem 2.1. *For any pairs of sequences $b \rightarrow \infty$ and $k \rightarrow \infty$ satisfying*

$$n/(bk \log(b) + k \log(k)) \rightarrow \infty,$$

there exists an estimator $V_n \in \mathcal{H}_{d,b}^k$ such that, for all $\varepsilon > 0$

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

The sample complexity for the multi-view histogram is perhaps more accurately approximated as being on the order of dbk however the d disappears in the asymptotic analysis. The following theorem states that one can control the error of Tucker histogram estimates so long as $n \sim bk + k^d$ (omitting logarithmic factors).

Theorem 2.2. *For any pairs of sequences $b \rightarrow \infty$ and $k \rightarrow \infty$ satisfying*

$$n/(bk \log(b) + k^d \log(k^d)) \rightarrow \infty,$$

there exists an estimator $V_n \in \tilde{\mathcal{H}}_{d,b}^k$ such that, for all $\varepsilon > 0$

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Allowing b to grow as aggressively as possible we achieve consistent estimation so long as $n \sim b \log b$ and k grows sufficiently slowly, regardless of dimensionality.

Corollary 2.1. For all d, b, k fix $\mathcal{R}_{d,b}^k$ to be either $\mathcal{H}_{d,b}^k$ or $\tilde{\mathcal{H}}_{d,b}^k$ ^{‡‡}. For any sequence $b \rightarrow \infty$ with $n/(b \log b) \rightarrow \infty$, there exists a sequence $k \rightarrow \infty$ and estimator $V_n \in \mathcal{R}_{d,b}^k$ such that, for all $\varepsilon > 0$

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \mathcal{R}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

The following result shows that the approximation error of the estimators in Theorem 2.1, Theorem 2.2, and Corollary 2.1 go to zero for all densities. Thus these estimators are universally consistent even when the NNTF model assumption is not satisfied.

Lemma 2.1. Let $p \in \mathcal{D}_d$. If $k \rightarrow \infty$ and $b \rightarrow \infty$ then $\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 \rightarrow 0$.

A straightforward consequence of this is that the Tucker histogram approximation error also goes to zero.

Lemma 2.2. Let $p \in \mathcal{D}_d$. If $k \rightarrow \infty$ and $b \rightarrow \infty$ then $\min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 \rightarrow 0$.

The next theorem shows that the rate on bk in Theorem 2.1 cannot be made significantly faster.

Theorem 2.3. Let $d \geq 2$, $b \rightarrow \infty$, and $k \rightarrow \infty$ with $b \geq k$ and $n/(bk) \rightarrow 0$. There exists no estimator $V_n \in \mathcal{H}_{d,b}^k$ such that, for all $\varepsilon > 0$, the following limit holds

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Likewise the rate on $bk + k^d$ can also not be significantly improved in Theorem 2.2.

Theorem 2.4. Let $d \geq 2$, $b \rightarrow \infty$, and $k \rightarrow \infty$ with $b \geq k$ and $n/(bk + k^d) \rightarrow 0$. There exists no estimator $V_n \in \tilde{\mathcal{H}}_{d,b}^k$ such that, for all $\varepsilon > 0$, the following limit holds

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

2.2.2 Finite-Sample and Distribution-Independent Bounds

In this section we investigate the convergence of multi-view histogram estimators to densities that satisfy the multi-view assumption. This will be done via a standard bias/variance decomposition style of argument. We begin with the following distribution-dependent finite-sample bound.

Proposition 2.1. Let $d, b, k, n \in \mathbb{N}$ and $0 < \delta \leq 1$. There exists an estimator $V_n \in \mathcal{H}_{d,b}^k$ such that

$$\sup_{p \in \mathcal{D}_d} P \left(\|p - V_n\|_1 > 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + 7 \sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7 \sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \right) < \delta,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

To analyze the approximation error, $\|p - q\|_1$, we first consider the case where p has a single component with L -Lipschitz marginals, so $p = \prod_{i=1}^d p_i$. Using the indicator function over the unit cube with Hölder's Inequality we have that $\|p - q\|_1 = \|(p - q) \mathbb{1}\|_1 \leq \|p - q\|_2 \|\mathbb{1}\|_2 = \|p - q\|_2$. It is possible to show that the L^2 projection of p onto $\mathcal{H}_{d,b}^1$ is achieved by simply projecting each marginal to its best approximating one-dimensional marginal i.e. (see the appendix)

$$\arg \min_{q \in \mathcal{H}_{d,b}^1} \left\| \prod_{i=1}^d p_i - q \right\|_2 = \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} p_i.$$

^{‡‡} \mathcal{R} is fixed to \mathcal{H} or $\tilde{\mathcal{H}}$ and doesn't change as n, b, k vary.

Let Lip_L be the set of L -Lipschitz probability density functions $[0, 1]$ and let $m_L \triangleq \sup_{f \in \text{Lip}_L} \|f\|_2$. The following theorem characterizes the approximation error for separable densities with Lipschitz continuous marginals.

Theorem 2.5. *Let $b^2 \geq L^2/12$ and f_1, \dots, f_d be elements of Lip_L then*

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_1 \leq \sqrt{m_L^{2d} - \left(m_L^2 - \frac{L^2}{12b^2}\right)^d}.$$

We show in the appendix that this decays at a rate of $O(b^{-1})$ and that this rate is tight. To characterize the behavior of m_L we have the following.

Proposition 2.2. *Let $m_L = \sup_{f \in \text{Lip}_L} \|f\|_2$, then*

$$m_L^2 = \begin{cases} L^2/12 + 1 & 0 \leq L \leq 2 \\ \sqrt{\frac{8L}{9}} & L \geq 2 \end{cases}.$$

The following theorem is a finite-sample bound for a well-constructed multi-view histogram estimator.

Theorem 2.6. *Fix $n, d, k \in \mathbb{N}$, $L \geq 2$, and $1 \geq \delta > 0$. There exists b and an estimator $V_n \in \mathcal{H}_{d,b}^k$ such that*

$$\sup_{p \in Q} P \left(\|V_n - p\|_1 > \frac{21dk^{1/3}L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \right) < \delta,$$

where Q is the set of densities of the form $\sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j}$ with $p_{i,j} \in \text{Lip}_L$, $w \in \Delta_k$, and V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Here we analyze the asymptotic rate at which the estimator in Theorem 2.6 converges to a large class of multi-view models. To this end let

$$\Omega_d = \left\{ \sum_{i=1}^{k'} w_i p_{i,1} \otimes \dots \otimes p_{i,d} \mid k' \in \mathbb{N}, w \in \Delta_{k'}, L' \geq 0, p_{i,j} \in \text{Lip}_{L'} \right\},$$

i.e. Ω_d is the space of all multi-view models whose component marginals are all Lipschitz continuous. Consider letting $L \rightarrow \infty$, $k \rightarrow \infty$, and $\delta \rightarrow 0$ in Theorem 2.6 arbitrarily slowly as $n \rightarrow \infty$. For some element p in Ω_d its respective maximum Lipschitz constant and component number, L' and k' , are fixed, so for sufficiently large n we have $L > L'$ and $k > k'$ and the bound from Theorem 2.6 applies. From this it follows that $\|V_n - p\|_1 \in \tilde{O}(1/\sqrt[3]{n})$. So we can construct an estimator that converges at rate $\tilde{O}(1/\sqrt[3]{n})$ for *any* multi-view model in Ω_d , independent of dimension! This rate appears to be approximately optimal since “for smooth densities, the average L^1 error for the histogram estimate must vary at least as $n^{-1/3}$ ” ([16], p. 99). For comparison the histogram estimator’s convergence rate is hindered exponentially in dimension.

Proposition 2.3. *Let V_n be the standard histogram estimator with $n/b^d \rightarrow \infty$. There exists $p \in \Omega_d$ such that $\|V_n - p\|_1 \in \omega(1/\sqrt[3]{n})$.*

2.3 Discussion

While Theorem 2.6 gives an estimator with good convergence, the class of densities Q is somewhat restrictive and likely not realistic for many situations where one would want to apply a nonparametric density estimator. Proposition 2.1, on the other hand, is not so restrictive since it depends on $\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1$, and more clearly conveys the message of this paper. Typical works on multi-view nonparametric density estimation assume that the target density p is a multi-view model and are interested in recovering the model components $p_{i,j}$ and w from (1). Similarly to how the standard histogram estimator doesn’t assume $p \in \mathcal{H}_{d,b}$ our work isn’t meant to assume that p is a multi-view

model, but is instead meant to explore the benefits of including a hyperparameter k , in addition to b , to restrict rank of the estimator. Just as the histogram can approximate any density as $b \rightarrow \infty$, Lemmas 2.1 and 2.2 show that the inclusion of k does not hinder the approximation power of the estimator.

To explore the trade-off between b and k first observe that setting $k = b^d$ gives $\mathcal{H}_{d,b}^k = \mathcal{H}_{d,b}$ since one can allocate one component to each bin. Theorem 2.1 and Proposition 2.1 with $k = b^d$ gives a sample complexity of approximately $n \sim b^{d+1}$ which coincides with standard histogram estimator. Alternatively, setting $k = 1$ restricts the estimator to separable histograms $\mathcal{H}_{d,b}^1$, with a sample complexity of approximately $n \sim b$. Thus we have a span of k yielding different estimators with maximal k corresponding to the standard histogram and minimal k corresponding to a naive Bayes assumption. We observe in Section 3 that this trade-off is useful in practice: we virtually never want k to be maximized.

3 Experiments

The previous section proved the existence of estimators that select NNTF histograms that offer advantages over the standard histogram. Unfortunately these estimators are not computationally tractable and only demonstrate the potential benefit incorporating an NNTF model in nonparametric density estimation. It would be nonetheless interesting to observe the behavior of an NNTF histogram estimator, even if it lacks the theoretical guarantees developed in the previous section. To this end we consider an L^2 -minimizing NNTF histogram estimator. For all d, b, k fix $\mathcal{R}_{d,b}^k$ to be either $\mathcal{H}_{d,b}^k$ or $\tilde{\mathcal{H}}_{d,b}^k$. We consider an estimator U_n , that attempts to minimize

$$U_n = \operatorname{argmin}_{\hat{p} \in \mathcal{R}_{d,b}^k} \|\hat{p} - p\|_2^2.$$

Note that

$$\|\hat{p} - p\|_2^2 = \|\hat{p}\|_2^2 - 2 \langle p, \hat{p} \rangle + \|p\|_2^2. \quad (5)$$

The $\|p\|_2^2$ term in (5) is not relevant when minimizing over \hat{p} . Additionally for data $X_1, \dots, X_n \stackrel{iid}{\sim} p$ the law of large numbers gives us

$$\langle p, \hat{p} \rangle = \int_{[0,1]^d} p(x) \hat{p}(x) dx = \mathbb{E}_{X \sim p} [\hat{p}(X)] \approx \frac{1}{n} \sum_{i=1}^n \hat{p}(X_i)$$

so we may consider

$$U_n \triangleq \operatorname{argmin}_{\hat{p} \in \mathcal{R}_{d,b}^k} \|\hat{p}\|_2^2 - \frac{2}{n} \sum_{i=1}^n \hat{p}(X_i) \quad (6)$$

as a practical version of (5) which is conveniently equivalent to nonnegative tensor factorization (see the appendix).

Risk expressions for nonparametric density estimation based on L^2 -minimization similar to (6) have appeared in previous works related to kernel density estimation [44, 9, 35]. Proving optimal rates on bandwidth in these settings seems challenging, however.

For our experiments we used the Tensorly library [27] to perform the nonnegative Tucker decomposition [26] with Tucker rank $[k, k, \dots, k]$ which was then projected to the simplex of probability tensors using [13]. We also performed experiments with nonnegative PARAFAC decompositions using [38, 27]. These decompositions performed poorly. This is potentially because the PARAFAC optimization is more difficult or the additional flexibility of the Tucker decomposition was more appropriate for the experimental datasets.

3.1 Experimental Setup

Our experiments were performed on the Scikit-learn [32] datasets MNIST and Diabetes [32], with labels removed. We use the estimated risk from (6) to evaluate estimator performance (which may cause negative performance values, however lower values always indicate smaller estimated L^2 -distance). Our experiments considered estimating histograms in $d = 2, 3, 4, 5$ dimensional space.

Table 1: Experimental Results

Dataset	d Red.	Dim.	Hist. Perf.	Tucker Perf.	Hist. Bins	Tucker Bins	Tucker k	p -val.
MNIST	PCA	2	-1.455±0.089	-1.502±0.102	6.531±1.499	8.375±1.780	4.968±1.976	5e-4
		3	-2.040±0.196	-2.268±0.195	4.781±0.738	6.718±1.565	5.781±1.340	2e-4
		4	-3.532±0.996	-4.014±0.655	4.031±0.585	5.343±1.018	4.375±0.695	2e-3
		5	-4.673±1.026	-6.157±2.924	3.468±0.499	4.343±0.592	3.281±0.514	4e-5
	Rand.	2	-2.034±0.100	-2.099±0.102	6.062±1.197	7.562±1.657	2.062±1.784	3e-5
		3	-3.086±0.207	-3.331±0.387	4.812±0.526	6.843±1.227	2.687±1.959	1e-4
		4	-4.307±0.290	-5.731±0.435	3.500±0.559	5.656±0.642	2.593±1.497	8e-7
		5	-6.327±0.522	-9.539±1.053	3.250±0.433	4.718±0.571	2.562±1.087	8e-7
	PCA	2	-2.079±0.122	-2.212±0.132	5.718±1.304	7.468±1.478	1.062±0.242	8e-6
		3	-3.010±0.364	-3.606±0.420	3.593±0.860	7.062±1.058	1.843±1.543	2e-6
Diabetes		4	-4.002±0.415	-4.423±0.701	3.000±0.000	5.906±0.804	2.343±1.107	2e-3
		5	-6.139±0.661	-6.043±1.192	3.000±0.000	3.750±0.968	1.843±0.617	0.91
	Rand.	2	-3.074±0.224	-3.277±0.287	6.843±1.227	9.250±1.936	1.093±0.384	7e-5
		3	-4.726±0.483	-5.353±0.751	4.968±0.769	8.406±1.343	1.625±1.672	2e-5
		4	-6.017±0.873	-7.732±1.497	4.062±0.704	6.718±1.328	2.093±1.155	1e-5
		5	-8.986±1.292	-12.61±2.477	3.062±0.242	5.093±0.521	2.531±0.865	2e-6

We consider two forms of dimensionality reduction. First we consider projecting the dataset onto its top d principle components. We also performed experiments projecting each dataset onto a random subspace of dimension d . These random subspaces were constructed so that each additional dimension adds a new index without affecting the others. To do this, we first randomly select an orthonormal basis for each dataset that remains unchanged for all experiments v_1, v_2, \dots . Then to transform a point X to dimension d we perform the following transform: $X_{\text{reduced dim.}} = [v_1 \cdots v_d]^T X$. We consider both transforms since PCA may select dimensions where the features tend to be independent, e.g. a multivariate Gaussian. After dimensionality reduction we scale and translate the data to fit in the unit cube.

With our preprocessed dataset, each experiment consisted of randomly selecting 200 samples for training and using the rest to evaluate performance (again using (6)). For the estimators we tested all combinations using 1 to b_{\max} bins per dimension and k from 1 to k_{\max} . As d increased the best cross validated b and k value decreased, so we reduced b_{\max} and k_{\max} for larger d to reduce computational time, while still leaving a sizable gap between the best cross validated b and k and b_{\max} and k_{\max} across all runs of all experiment. For $d = 2, 3$ we have $b_{\max} = 15$ and $k_{\max} = 10$; for $d = 4$ we have $b_{\max} = 12$ and $k_{\max} = 8$; for $d = 5$ we have $b_{\max} = 8$ and $k_{\max} = 6$. For parameter fitting we used random subset cross validation repeated 80 times using 40 of the 200 samples to evaluate the performance of the estimator fit using the other 160 samples. Performing 80 folds of cross validation was necessary because of the high variance of the histogram’s estimated risk. This high variance is likely due to the noncontinuous nature of the histogram estimator itself and the noncontinuity of the histogram as a function of the data, i.e. slightly moving one training sample can potentially change histogram bin in which it lies. Each experiment was run 32 times and we report the mean and standard deviations of estimator performance as well as the best parameters found from cross validation. We additionally apply the two-tailed Wilcoxon signed rank test to the 32 pairs of performance results to statistically determine if the mean performance between the standard histogram and our algorithm are different and report the corresponding p -value.

3.2 Results

Our results are presented in Table 1. Apart from the density estimators’ performance (“Hist. Perf.” and “Tucker Perf.”) this table contains the mean and standard deviation over the 32 trials for the optimal cross validated parameters for the estimators. This includes the number of bins, corresponding to b from the earlier sections, and the Tucker rank k . We see that the Tucker histogram usually outperforms the standard histogram (“Tucker Perf.” is less than “Hist. Perf.”) with high statistical significance (small p -val). As one would expect, the Tucker histogram cross validates for more bins, presumably reducing the k in exchange for more bins. Note that the MNIST PCA experiments always cross validated for the largest k and yielded less statistically significant performance increases (outside of the unusual Diabetes PCA $d = 5$ experiment). Presumably this experiment fit the NNTF assumption the least well (hence the large k) and thus received the least benefit.

4 Conclusion

In this paper we have theoretically and experimentally demonstrated the advantages of including rank restriction into nonparametric density estimation. In particular, rank restriction may be a way to overcome the curse of dimensionality since it reduces the sample complexity penalty incurred from dimensionality from exponential to linear. This paper is an initial theoretical foray for demonstrating the potential of combining rank restriction with nonparametric methods.

Acknowledgments and Disclosure of Funding

Robert A. Vandermeulen acknowledges support by the Berlin Institute for the Foundations of Learning and Data (BIFOLD) sponsored by the German Federal Ministry of Education and Research (BMBF). Antoine Ledent acknowledges support by the German Research Foundation (DFG) award KL 2698/5-1 and the Federal Ministry of Science and Education (BMBF) award 01IS18051A. The authors thank Marius Kloft and Clayton Scott for useful discussions related to this work. The authors thank Nicholas D. Sidiropoulos for bringing to our attention some previous works on low-rank nonparametric density estimation.

Appendix Overview

Section A contains the basic theoretical tools we need for the rest of the appendix. We will also be using some notation from the paper without introduction. In Section B we have all the proofs for the results relating to asymptotic parameter rates of the histogram spaces in our estimators; this section corresponds to Section 2.2.1 in the main text. Section C contains the proofs relating to finite-sample rates, and several expanded or additional results, corresponding to Section 2.2.2 in the main text. Section 2.2.2 in the main text contains some results that are simplified for readability and this appendix has some results which are tighter, but less readable. Additionally Section C contains all the finite-sample analysis for Tucker histograms which were omitted in the main text. Experimental details, specifically the equivalency of L^2 minimization and nonnegative tensor factorization can be found in Section D. Section E is a bit different from the other sections and contains some discussion relevant to the future of this direction of research. There we show that not all densities can be written as a summation of separable densities with nonnegative coefficients.

A Theoretical Basics

A.1 Notation

In this section we include some notation that was not contained in the main text but will be necessary for the rest of the appendix. In particular we will need to introduce a fair amount of tensor notation which will be used to represent the various histograms. A histogram on the unit cube can naturally be represented as a tensor, which will be useful for deriving many of the results in this paper.

From the main text recall that $\mathcal{T}_{d,b}$ is the set of probability tensors in $\mathbb{R}^{b \times d}$, i.e. their entries are nonnegative and sum to one.

Let $\mathcal{T}_{d,b}^k$ be the set of tensors that are a convex combination of k separable probability tensors (which are analogous to multi-view models) i.e.

$$\mathcal{T}_{d,b}^k \triangleq \left\{ \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j} \mid w \in \Delta_k, p_{i,j} \in \Delta_b \right\}.$$

The following is the set of probability tensors constructed via a nonnegative Tucker factorization ($[k]^d$ represents a multi-index)

$$\tilde{\mathcal{T}}_{d,b}^k \triangleq \left\{ \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} \mid W \in \mathcal{T}_{d,k}, p_{i,j} \in \Delta_b \right\}.$$

For a multi-index $A \in [b]^d$ we define $\mathbf{e}_{d,b,A}$ as the element of $\mathcal{T}_{d,b}$ where the (A_1, \dots, A_d) -th entry is one and is zero elsewhere.

Note that there exists a $\ell^1 \rightarrow L^1$ linear isometry $U_{d,b} : \mathcal{T}_{d,b} \rightarrow \mathcal{H}_{d,b}$ with $U_{d,b}$ defined as

$$U_{d,b}(\mathbf{e}_{d,b,A}) = h_{d,b,A}.$$

The inverse function, $U_{d,b}^{-1}$, simply transforms a histogram to the tensor representing its bin weights and $U_{d,b}$ performs the reverse transformation. Note that $U_{d,b}$ is also a bijection between $\mathcal{T}_{d,b}^k \rightarrow \mathcal{H}_{d,b}^k$ and $\tilde{\mathcal{T}}_{d,b}^k \rightarrow \tilde{\mathcal{H}}_{d,b}^k$. Much of our analysis on histograms will be performed on the space of probability tensors with the analysis being translated to histograms via this operator.

For a set of vectors \mathcal{V} we define k -mix $(\mathcal{V}) \triangleq \left\{ \sum_{i=1}^k w_i v_i \mid w \in \Delta_k, v_i \in \mathcal{V} \right\}$, i.e. the set of convex combinations of collections of k vectors from \mathcal{V} . We define $N(\mathcal{V}, \varepsilon)$ to be the minimum cardinality for a subset of \mathcal{V} which ε -covers \mathcal{V} (with closed balls) with respect to the $\|\cdot\|_1$ metric. It will be clear from context whether $\|\cdot\|_1$ represents the ℓ^1 , L^1 , or total variation norm.

A.2 Preliminary Results

The following lemmas will be useful for all the theoretical results.

Lemma A.1. *For all $0 < \varepsilon \leq 1$ we have that $N(\Delta_b, \varepsilon) \leq \left(\frac{2b}{\varepsilon}\right)^b$.*

Lemma A.2. *For all $0 < \varepsilon \leq 1$ we have that $N(\mathcal{T}_{d,b}^1, \varepsilon) \leq \left(\frac{2bd}{\varepsilon}\right)^{bd}$.*

Lemma A.3. *Let \mathcal{P} be a set of probability measures, then*

$$N(k\text{-mix}(\mathcal{P}), \varepsilon + \delta) \leq N(\mathcal{P}, \varepsilon)^k N(\Delta_k, \delta).$$

Lemma A.4. *For all $0 < \varepsilon \leq 1$ the following holds $N(\mathcal{T}_{d,b}^k, \varepsilon) \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k}{\varepsilon}\right)^k$.*

Through application of the $U_{d,b}$ operator we now have a characterization of the complexity of the space $\mathcal{H}_{d,b}^k$.

Corollary A.1. *For all $0 < \varepsilon \leq 1$ following holds $N(\mathcal{H}_{d,b}^k, \varepsilon) \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k}{\varepsilon}\right)^k$.*

The following are analogous results for Tucker histograms.

Lemma A.5. *For all $0 < \varepsilon \leq 1$ the following holds $N(\tilde{\mathcal{T}}_{d,b}^k, \varepsilon) \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k^d}{\varepsilon}\right)^{k^d}$.*

Corollary A.2. *For all $0 < \varepsilon \leq 1$ following holds $N(\tilde{\mathcal{H}}_{d,b}^k, \varepsilon) \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k^d}{\varepsilon}\right)^{k^d}$.*

The following lemma from [6] provides us with a way to choose good estimators from finite collections of densities. It can be proven by applying a Chernoff bound to [10], Theorem 6.3.

Lemma A.6 (Thm 3.4 page 7 of [6], Thm 3.6 page 54 of [10]). *There exists a deterministic algorithm that, given a collection of distributions p_1, \dots, p_M , a parameter $\varepsilon > 0$ and at least $\frac{\log(3M^2/\delta)}{2\varepsilon^2}$ iid samples from an unknown distribution p , outputs an index $j \in [M]$ such that*

$$\|p_j - p\|_1 \leq 3 \min_{i \in [M]} \|p_i - p\|_1 + 4\varepsilon$$

with probability at least $1 - \frac{\delta}{3}$.

We present the following asymptotic version of the previous lemma. We highlight the use of finding sufficiently slow rates on parameters in order to establish asymptotic results, a technique which we will use in later proofs.

Lemma A.7. *Let $(\mathcal{P}_n)_{n \in \mathbb{N}}$ be a sequence of finite collections of densities in \mathcal{D}_d where $|\mathcal{P}_n| \rightarrow \infty$ with $n/\log(|\mathcal{P}_n|) \rightarrow \infty$. Then there exists a sequence of estimators $V_n \in \mathcal{P}_n$ such that, for all $\gamma > 0$,*

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \mathcal{P}_n} \|p - q\|_1 + \gamma \right) \rightarrow 0,$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Proof of Lemma A.7. Let $M = M(n) = |\mathcal{P}_n|$. Since $n/\log(M) \rightarrow \infty$ we have that for all $c > 0$ there exists a N_c such that, for all $n \geq N_c$ we have $n/\log(M) \geq c$ or equivalently $n \geq c \log(M)$. Because of this there exists sequence of positive values $C = C(n)$ such that $C \rightarrow \infty$ and $n \geq C \log(M)$.

We will be making use of the algorithm in Lemma A.6 as well as its notation. If we can show that there exist sequences of positive values $\varepsilon(n) \rightarrow 0, \delta(n) \rightarrow 0$ such that, for sufficiently large n , the following holds

$$\frac{\log(3M^2/\delta)}{2\varepsilon^2} \leq n,$$

then can simply set V_n equal to be the estimator from Lemma A.6 for sufficiently large n and, because the lemma holds independent of choice of p , the theorem statement follows.

Let $\varepsilon = (2/C)^{1/4}$ and $\delta = 3/\left(\exp\left(2\sqrt{\frac{C}{2}}\right)\right)$. Note that these are both positive sequences which converge to zero. Now we have

$$\begin{aligned} \frac{\log(3M^2/\delta)}{2\varepsilon^2} &= \frac{\log(M^2) + \log(3/\delta)}{2\varepsilon^2} \\ &= \frac{2\log(M) + \log(3/\delta)}{2\varepsilon^2} = \frac{\log(M) + \frac{1}{2}\log(3/\delta)}{\varepsilon^2} \\ &= \varepsilon^{-2} \left(\log(M) + \frac{1}{2}\log(3/\delta) \right) \\ &= \left((2/C)^{1/4} \right)^{-2} \left(\log(M) + \frac{1}{2}\log\left(\exp\left(2\sqrt{\frac{C}{2}}\right)\right) \right) \\ &= \sqrt{\frac{C}{2}} \left(\log(M) + \sqrt{\frac{C}{2}} \right) = \sqrt{\frac{C}{2}} \log(M) + \frac{C}{2}. \end{aligned} \tag{7}$$

For sufficiently large C and M we have that the RHS of (7) is less than or equal to

$$\begin{aligned} \frac{C}{2} \log(M) + \frac{C}{2} &\leq \frac{C}{2} \log(M) + \frac{C}{2} \log(M) \\ &= C \log(M) \leq n. \end{aligned}$$

which completes our proof. \square

A.3 Theoretical Basics Proofs

All norms are either the ℓ^1 , L^1 , or total variation norm, which are equivalent with respect to our analysis and the proper norm will be clear from context.

Proof of Lemma A.1. In Section 7.4 from [10], the authors show that for any collection of measures μ_1, \dots, μ_b , for all $\varepsilon > 0$, that

$$N(\text{Conv}(\{\mu_1, \dots, \mu_b\}), \varepsilon) \leq \left(b + \frac{b}{\varepsilon}\right)^b.$$

With the additional assumption that $\varepsilon \leq 1$ we have that $b + \frac{b}{\varepsilon} \leq \frac{b}{\varepsilon} + \frac{b}{\varepsilon} = \frac{2b}{\varepsilon}$ and thus

$$N(\text{Conv}(\{\mu_1, \dots, \mu_b\})) \leq \left(\frac{2b}{\varepsilon}\right)^b.$$

If we let $\mu_i = \mathbf{e}_i$, the indicator vector at index i , then the lemma follows. \square

Proof of Lemma A.2. From Lemma A.1 we know there exists a finite collection of probability vectors $\tilde{\mathcal{P}}$ such that $\tilde{\mathcal{P}}$ is an ε/d -covering of Δ_b and $|\tilde{\mathcal{P}}| \leq \left(\frac{2bd}{\varepsilon}\right)^b$. Note that the set $\{\tilde{p}_1 \otimes \dots \otimes \tilde{p}_d \mid \tilde{p}_i \in \tilde{\mathcal{P}}\}$ contains at most $\left(\left(\frac{2bd}{\varepsilon}\right)^b\right)^d = \left(\frac{2bd}{\varepsilon}\right)^{bd}$ elements. We will now show

that this set is an ε -cover of $\mathcal{T}_{d,b}^1$. Let $p_1 \otimes \dots \otimes p_d \in \mathcal{T}_{d,b}^1$ be arbitrary. From our construction of $\tilde{\mathcal{P}}$ there exist elements $\tilde{p}_1, \dots, \tilde{p}_d \in \tilde{\mathcal{P}}$ such that $\|p_i - \tilde{p}_i\|_1 \leq \frac{\varepsilon}{d}$.

We will now make use of Lemma 3.3.7 in [34], which states that, for any collection of probability vectors q_1, \dots, q_d and $\tilde{q}_1, \dots, \tilde{q}_d$, the following holds

$$\left\| \prod_{i=1}^d q_i - \prod_{j=1}^d \tilde{q}_j \right\|_1 \leq \sum_{i=1}^d \|q_i - \tilde{q}_i\|_1.$$

From this it follows that

$$\left\| \prod_{i=1}^d p_i - \prod_{j=1}^d \tilde{p}_j \right\|_1 \leq \sum_{i=1}^d \|p_i - \tilde{p}_i\|_1 \leq d \frac{\varepsilon}{d} = \varepsilon$$

thus completing our proof. \square

Proof of Lemma A.3. Let $\tilde{\mathcal{P}}$ be the finite collection of probability measures with $|\tilde{\mathcal{P}}| = N(\mathcal{P}, \varepsilon)$ which ε -covers \mathcal{P} . Similarly let $W \subset \Delta_k$ with $|W| = N(\Delta_k, \delta)$ such that W is a δ -cover of Δ_k . Consider the set

$$\Omega = \left\{ \sum_{i=1}^k \tilde{w}_i \tilde{p}_i \mid \tilde{w} \in W, \tilde{p}_i \in \tilde{\mathcal{P}} \right\}.$$

Note that this set contains at most $N(\mathcal{P}, \varepsilon)^k N(\Delta_k, \delta)$ elements. We will now show that it $(\delta + \varepsilon)$ -covers k -mix (\mathcal{P}) , which completes the proof. Let $\sum_{i=1}^k p_i w_i \in k$ -mix (\mathcal{P}) . We know there exists elements $\tilde{p}_1, \dots, \tilde{p}_k \in \tilde{\mathcal{P}}$ such that $\|\tilde{p}_i - p_i\|_1 \leq \varepsilon$ and $\tilde{w} \in W$ such that $\|w - \tilde{w}\|_1 \leq \delta$ and thus $\sum_{i=1}^k \tilde{p}_i \tilde{w}_i \in \Omega$. Now observe that

$$\begin{aligned} \left\| \sum_{i=1}^k \tilde{p}_i \tilde{w}_i - \sum_{j=1}^k p_j w_j \right\|_1 &= \left\| \sum_{i=1}^k \tilde{p}_i \tilde{w}_i - \sum_{j=1}^k p_j \tilde{w}_j + \sum_{l=1}^k p_l \tilde{w}_l - \sum_{r=1}^k p_r w_r \right\|_1 \\ &\leq \left\| \sum_{i=1}^k (\tilde{p}_i - p_i) \tilde{w}_i \right\|_1 + \left\| \sum_{i=1}^k p_i (\tilde{w}_i - w_i) \right\|_1 \\ &\leq \sum_{i=1}^k \tilde{w}_i \|\tilde{p}_i - p_i\|_1 + \sum_{i=1}^k |w_i| |\tilde{w}_i - w_i| \\ &\leq \sum_{i=1}^k \tilde{w}_i \varepsilon + \|\tilde{w} - w\|_1 \\ &\leq \varepsilon + \delta. \end{aligned}$$

\square

Proof of Lemma A.4. Note that $\mathcal{T}_{d,b}^k = k$ -mix $(\mathcal{T}_{d,b}^1)$. Applying Lemma A.3 followed by Lemmas A.1 and A.2 we have that

$$N(\mathcal{T}_{d,b}^k, \varepsilon) \leq N(\mathcal{T}_{d,b}^1, \varepsilon/2)^k N(\Delta_k, \varepsilon/2) \leq \left(\frac{4bd}{\varepsilon} \right)^{bdk} \left(\frac{4k}{\varepsilon} \right)^k.$$

\square

Proof of Lemma A.5. Fix k, d, b and $0 < \varepsilon \leq 1$. We are going to construct an ε -cover of $\tilde{\mathcal{T}}_{d,b}^k$. From Lemma A.1 we know that there exists a set $\mathcal{B} \subset \Delta_b$ which $(\frac{\varepsilon}{2d})$ -covers of Δ_b and contains no more

than $\left(\frac{4bd}{\varepsilon}\right)^b$ elements. Let \mathcal{P} be the collection of all $d \times k$ arrays whose entries are elements from \mathcal{B} . So we have that

$$|\mathcal{P}| = |\mathcal{B}|^{dk} \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk}.$$

From Lemma A.1 there exists \mathcal{W} which is an $\varepsilon/2$ -cover of $\mathcal{T}_{d,k}$ and contains no more than $(4k^d/\varepsilon)^{k^d}$ elements. Now let

$$\mathcal{L}_{d,b}^k = \left\{ \sum_{S \in [k]^d} \widetilde{W}_S \prod_{i=1}^d \widetilde{p}_{i,S_i} \mid \widetilde{W} \in \mathcal{W}, \widetilde{p} \in \mathcal{P} \right\}.$$

Note that

$$|\mathcal{L}_{d,b}^k| \leq |\mathcal{W}| |\mathcal{P}| \leq \left(\frac{4k^d}{\varepsilon}\right)^{k^d} \left(\frac{4bd}{\varepsilon}\right)^{bdk}.$$

We will now show that $\mathcal{L}_{d,b}^k$ is an ε -cover of $\widetilde{\mathcal{T}}_{d,b}^k$. To this end let $\sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} \in \widetilde{\mathcal{T}}_{d,b}^k$ be arbitrary, where $W \in \mathcal{T}_{d,k}$ and $p_{i,j} \in \Delta_b$. From our construction of \mathcal{W} , there exists $\widetilde{W} \in \mathcal{W}$ such that $\|W - \widetilde{W}\|_1 \leq \varepsilon/2$. There also exists $\widetilde{p} \in \mathcal{P}$ such that $\|\widetilde{p}_{i,j} - p_{i,j}\|_1 \leq \varepsilon/2$ for all i, j . Therefore we have that

$$\sum_{S \in [k]^d} \widetilde{W}_S \prod_{i=1}^d \widetilde{p}_{i,S_i} \in \mathcal{L}_{d,b}^k.$$

So finally

$$\begin{aligned} & \left\| \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} - \sum_{R \in [k]^d} \widetilde{W}_R \prod_{j=1}^d \widetilde{p}_{j,R_j} \right\|_1 \\ & \leq \left\| \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} - \sum_{R \in [k]^d} W_R \prod_{j=1}^d \widetilde{p}_{j,R_j} \right\|_1 + \left\| \sum_{S \in [k]^d} W_S \prod_{i=1}^d \widetilde{p}_{i,S_i} - \sum_{R \in [k]^d} \widetilde{W}_R \prod_{j=1}^d \widetilde{p}_{j,R_j} \right\|_1 \\ & \leq \sum_{S \in [k]^d} W_S \left\| \prod_{i=1}^d p_{i,S_i} - \prod_{j=1}^d \widetilde{p}_{j,S_j} \right\|_1 + \sum_{R \in [k]^d} |W_R - \widetilde{W}_R| \left\| \prod_{j=1}^d \widetilde{p}_{j,R_j} \right\|_1 \\ & \leq \sum_{S \in [k]^d} W_S \sum_{i=1}^d \|p_{i,S_i} - \widetilde{p}_{i,S_i}\|_1 + \sum_{R \in [k]^d} |W_R - \widetilde{W}_R| \left\| \prod_{j=1}^d \widetilde{p}_{j,R_j} \right\|_1 \\ & \leq \sum_{S \in [k]^d} W_S \frac{\varepsilon}{2} + \|W - \widetilde{W}\|_1 \\ & \leq \varepsilon/2 + \varepsilon/2 = \varepsilon. \end{aligned}$$

□

B Asymptotic Theoretical Results

This section contains results related to the asymptotic results related to the growth of model parameters with respect to the number of training samples. It corresponds to Section 2.2.1 in the main text.

Proof of Theorem 2.1. We will be applying the estimator from Lemma A.7 to a series of δ -covers of $\mathcal{H}_{d,b}^k$. We begin by constructing a series of δ -covers whose cardinality doesn't grow too quickly. Corollary A.1 states that, for all $0 < \delta \leq 1$, that $N(\mathcal{H}_{d,b}^k, \delta) \leq \left(\frac{4bd}{\delta}\right)^{bdk} \left(\frac{4k}{\delta}\right)^k$. For sufficiently large b and k and sufficiently small δ , the following holds

$$\begin{aligned} \log \left(\left(\frac{4bd}{\delta} \right)^{bdk} \left(\frac{4k}{\delta} \right)^k \right) &= bdk \log \left(\frac{4bd}{\delta} \right) + k \log \left(\frac{4k}{\delta} \right) \\ &= bdk \left[\log(b) + \log \left(\frac{4d}{\delta} \right) \right] + k \left[\log(k) + \log \left(\frac{4}{\delta} \right) \right] \\ &\leq bdk \left[\log(b) + \log(b) \log \left(\frac{4d}{\delta} \right) \right] + dk \left[\log(k) + \log(k) \log \left(\frac{4d}{\delta} \right) \right] \\ &= (bk \log(b) + k \log(k)) d \left(1 + \log \left(\frac{4d}{\delta} \right) \right). \end{aligned} \quad (8)$$

Using the argument from the proof of Lemma A.7 we have that, because $n/(bk \log(b) + k \log(k)) \rightarrow \infty$ there exists a sequence of positive values $C = C(n)$ such that $C \rightarrow \infty$ and $n > C[bk \log(b) + k \log(k)]$. If we let $\delta = \frac{4d}{\exp(\frac{C}{d}-1)}$ we have that $\delta \rightarrow 0$ and

$$(bk \log(b) + k \log(k)) d \left(1 + \log \left(\frac{4d}{\delta} \right) \right) \leq n.$$

Because of this we can construct collections of densities $\tilde{\mathcal{P}}_n \subset \mathcal{H}_{d,b}^k$ such that $\tilde{\mathcal{P}}_n$ is a δ -covering of $\mathcal{H}_{d,b}^k$ with $|\tilde{\mathcal{P}}| \rightarrow \infty$, $n/\log |\tilde{\mathcal{P}}_n| \rightarrow \infty$ and $\delta \rightarrow 0$. Let V_n be the estimator from Lemma A.7 applied to the sequence $\tilde{\mathcal{P}}_n$.

Let $\varepsilon > 0$ be arbitrary. Due to the way that we have constructed the sequence $\tilde{\mathcal{P}}_n$, for sufficiently large n , we have that $3 \sup_{q \in \mathcal{H}_{d,b}^k} \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|q - \tilde{q}\|_1 \leq \varepsilon/2$. It therefore follows that, for sufficiently large n , the following holds for all $p \in \mathcal{D}_d$

$$\begin{aligned} 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + \varepsilon &\geq 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + 3 \sup_{q \in \mathcal{H}_{d,b}^k} \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|q - \tilde{q}\|_1 + \varepsilon/2 \\ &\geq 3 \min_{q \in \mathcal{H}_{d,b}^k} \left[\|p - q\|_1 + \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|q - \tilde{q}\|_1 \right] + \varepsilon/2 \\ &= 3 \min_{q \in \mathcal{H}_{d,b}^k} \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|p - q\|_1 + \|q - \tilde{q}\|_1 + \varepsilon/2 \\ &\geq 3 \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|p - \tilde{q}\|_1 + \varepsilon/2. \end{aligned}$$

From this we have that, for sufficiently large n

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_i - p\|_1 > 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \leq \sup_{p \in \mathcal{D}_d} P \left(\|V_i - p\|_1 > 3 \min_{\tilde{q} \in \tilde{\mathcal{P}}_n} \|p - \tilde{q}\|_1 + \varepsilon/2 \right)$$

and the right side goes to zero due to Lemma A.7, thus completing the proof. \square

Proof of Theorem 2.2. This proof is very similar to the proof of Theorem 2.1. We will be applying the estimator from Lemma A.7 to a series of δ -covers of $\tilde{\mathcal{H}}_{d,b}^k$. We begin by constructing a series of δ -covers whose cardinality doesn't grow too quickly. Corollary A.2 states that, for all $0 < \delta \leq 1$,

that $N\left(\tilde{\mathcal{H}}_{d,b}^k, \delta\right) \leq \left(\frac{4bd}{\delta}\right)^{bdk} \left(\frac{4k^d}{\delta}\right)^{k^d}$. For sufficiently large b and k and sufficiently small δ , the following holds

$$\begin{aligned} \log \left(\left(\frac{4bd}{\delta} \right)^{bdk} \left(\frac{4k^d}{\delta} \right)^{k^d} \right) &= bdk \log \left(\frac{4bd}{\delta} \right) + k^d \log \left(\frac{4k^d}{\delta} \right) \\ &\leq d \left(bk \log \left(\frac{4bd}{\delta} \right) + k^d \log \left(\frac{4k^d}{\delta} \right) \right) \\ &= d \left(bk \left(\log(b) + \log \left(\frac{4d}{\delta} \right) \right) + k^d \left(\log(k^d) + \log \left(\frac{4}{\delta} \right) \right) \right) \\ &\leq d \left(bk \left(\log(b) + \log \left(\frac{4d}{\delta} \right) \right) + k^d \left(\log(k^d) + \log \left(\frac{4d}{\delta} \right) \right) \right) \\ &= (bk \log(b) + k^d \log(k^d)) d \left(1 + \log \left(\frac{4d}{\delta} \right) \right). \end{aligned}$$

Note that replacing $bk \log(b) + k \log(k)$ with $bk \log(b) + k^d \log(k^d)$ in the last line is exactly (8) in our proof of Theorem 2.1. From here we can proceed exactly as in the proof of Theorem 2.1 by replacing $\mathcal{H}_{d,b}^k$ with $\tilde{\mathcal{H}}_{d,b}^k$ and $bk \log(b) + k \log(k)$ with $bk \log(b) + k^d \log(k^d)$. \square

Proof of Corollary 2.1. This follows directly from Theorems 2.1 and 2.2 and selecting an appropriately slow rate for $k \rightarrow \infty$. \square

Proof of Lemma 2.1. Let $\varepsilon > 0$. Theorem 5 in Chapter 2 of [16]^{§§} states that, for any $p \in \mathcal{D}_d$, that $\min_{h \in \mathcal{H}_{d,b}} \|p - h\|_1 \rightarrow 0$ as $b \rightarrow \infty$, i.e. the bias of a histogram estimator goes to zero as the number of bins per dimension goes to infinity. Thus there exists a sufficiently large B such that there exists a histogram $h \in \mathcal{H}_{d,B}$ which is a good approximation of p , $\|p - h\|_1 < \varepsilon/2$. In this proof we will argue that once $k \geq B^d$ and b is sufficiently large, we can find an element of $\mathcal{H}_{d,b}^k$ where the multi-view components can approximate the B^d bins of h .

We have that, for some $w \in \mathcal{T}_{d,B}$

$$h = \sum_{A \in [B]^d} w_A h_{d,B,A}.$$

From the same theorem in [16] there exists a_0 such that, for all $a \geq a_0$, for all i , there exists $\tilde{h}_{1,a,i} \in \mathcal{H}_{1,a}$ such that $\|h_{1,B,i} - \tilde{h}_{1,a,i}\|_1 < \varepsilon/(2d)$ for all $i \in [B]$. For any multi-index $A \in [B]^d$, we define

$$\tilde{h}_{d,a,A} = \prod_{j=1}^d \tilde{h}_{1,a,A_j}.$$

Now we have that, for all $a \geq a_0$ and $A \in [B]^d$,

$$\begin{aligned} \|h_{d,B,A} - \tilde{h}_{d,a,A}\|_1 &= \left\| \prod_{i=1}^d h_{1,B,A_i} - \prod_{j=1}^d \tilde{h}_{1,a,A_j} \right\|_1 \\ &\leq \sum_{i=1}^d \|h_{1,B,A_i} - \tilde{h}_{1,a,A_i}\|_1 \\ &\leq d \frac{\varepsilon}{2d} \\ &= \varepsilon/2, \end{aligned} \tag{9}$$

^{§§}See p. 20 in this text for the application to histograms.

where we use the previously mentioned product measure inequality for (9). As soon as $k \geq B^d$ and $a \geq a_0$ the set $\mathcal{H}_{d,a}^k$ contains the element,

$$\tilde{h} \triangleq \sum_{A \in [B]^d} w_A \tilde{h}_{d,a,A}.$$

Now we have that, for all $a \geq a_0$.

$$\begin{aligned} \|h - \tilde{h}\|_1 &= \left\| \sum_{A \in [B]^d} w_A h_{d,B,A} - \sum_{Q \in [B]^d} w_Q \tilde{h}_{d,a,Q} \right\|_1 \\ &\leq \sum_{A \in [B]^d} w_A \|h_{d,B,A} - \tilde{h}_{d,a,A}\|_1 \\ &\leq \varepsilon/2. \end{aligned}$$

From the triangle inequality we have that

$$\|p - \tilde{h}\|_1 \leq \|p - h\|_1 + \|h - \tilde{h}\|_1 \leq \varepsilon.$$

So we have that, for sufficiently large b and k

$$\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 \leq \varepsilon$$

which completes our proof. \square

Proof of Lemma 2.2. We will show that $\mathcal{H}_{d,b}^k \subset \tilde{\mathcal{H}}_{d,b}^k$ and the lemma clearly follows due to Lemma 2.1. Any element of $\mathcal{H}_{d,b}^k$ will have the following representation

$$\sum_{i=1}^k w_i \prod_{j=1}^d f_{i,j} : w \in \Delta_k, f_{i,j} \in \mathcal{H}_{1,b}. \quad (10)$$

Letting $W \in \mathcal{T}_{d,k}$ with $W_{i,\dots,i} = w_i$ for all i , the rest of the entries of W be zero, and $\tilde{f}_{j,i} = f_{i,j}$ for all i, j we have that

$$\begin{aligned} \sum_{S \in [k]^d} W_S \prod_{j=1}^d \tilde{f}_{j,S_j} &= \sum_{i=1}^k W_{i,\dots,i} \prod_{j=1}^d \tilde{f}_{j,i} \\ &= \sum_{i=1}^k w_i \prod_{j=1}^d f_{i,j} \end{aligned}$$

so we have that (10) is an element of $\tilde{\mathcal{H}}_{d,b}^k$ and we are done. \square

Proof of Theorem 2.3. We will proceed by contradiction. Suppose V_n is an estimator violating the theorem statement, i.e. there exist sequences $b \rightarrow \infty$ and $k \rightarrow \infty$ with $n/(bk) \rightarrow 0$ and $b \geq k$ such that, for all $\varepsilon > 0$,

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0.$$

Let $(p_n)_{n=1}^\infty$ be a sequence of probability vectors $p_n \in \Delta_{b(n) \times k(n)}$ which represent distributions over $[b(n)] \times [k(n)]$. Let $\mathcal{X}_n \triangleq (X_{n,1}, \dots, X_{n,n})$ with $X_{n,1}, \dots, X_{n,n} \stackrel{iid}{\sim} p_n$.

We will now construct a series of estimators for p_n using V_n . Let $\tilde{\mathcal{X}}_n = (\tilde{X}_{n,1}, \dots, \tilde{X}_{n,n})$ which are independent random variables with $\tilde{X}_{n,i} \sim h_{d,b,(X_{n,i,1}, \dots, 1)}$, so $\tilde{X}_{n,i}$ is uniformly distributed over the bin designated by $X_{n,i}, 1, \dots, 1$. For this proof we will assume $d > 2$ but the proof can be

simplified in a straightforward manner to the $d = 2$ case by ignoring the indices and modes beyond the second. Note that $X_{n,i}$ contains two indices. Now we have the following for the densities of $\tilde{X}_{n,i}$

$$\begin{aligned}
p_{\tilde{X}_{n,i}} &= \sum_{(j,\ell) \in [b] \times [k]} p_{\tilde{X}_{n,i}|X_{n,i}=(j,\ell)} P(X_{n,i} = (j,\ell)) \\
&= \sum_{(j,\ell) \in [b] \times [k]} h_{d,b,(j,\ell,1,\dots,1)} p_n(j,\ell) \\
&= \sum_{\ell \in [k]} \sum_{j \in [b]} h_{d,b,(j,\ell,1,\dots,1)} p_n(j,\ell) \\
&= \sum_{\ell \in [k]} \sum_{j \in [b]} p_n(j,\ell) h_{1,b,j} \otimes h_{1,b,\ell} \otimes \prod_{a \in [d-2]} h_{1,b,1} \\
&= \sum_{\ell \in [k]} \left(\sum_{j \in [b]} p_n(j,\ell) h_{1,b,j} \right) \otimes h_{1,b,\ell} \otimes \prod_{a \in [d-2]} h_{1,b,1} \tag{11}
\end{aligned}$$

$$= \sum_{\ell \in [k]} \left(\sum_{q \in [b]} p_n(q,\ell) \right) \left(\sum_{j \in [b]} \frac{p_n(j,\ell)}{\sum_{q \in [b]} p_n(q,\ell)} h_{1,b,j} \right) \otimes h_{1,b,\ell} \otimes \prod_{a \in [d-2]} h_{1,b,1}. \tag{12}$$

For (12) we let $0/0$ be equal to zero as is common for discrete conditioning. This last line is in the form of (4) in the main text and is thus an element of $\mathcal{H}_{d,b}^k$. To see this we will show the correspondence between the terms in (12) from here and the terms in (4) in the main text:

$$\begin{aligned}
w_\ell &:= \left(\sum_{q \in [b]} p_n(q,\ell) \right) \\
f_{\ell,1} &:= \left(\sum_{j \in [b]} \frac{p_n(j,\ell)}{\sum_{q \in [b]} p_n(q,\ell)} h_{1,b,j} \right) \\
f_{\ell,2} &:= h_{1,b,\ell} \\
f_{i,j} &:= h_{1,b,1}, \forall j > 2, \forall i.
\end{aligned}$$

Let V_n estimate $\tilde{P}_n \triangleq p_{\tilde{X}_{n,i}}$ so $\tilde{X}_{n,1}, \dots, \tilde{X}_{n,n} \stackrel{iid}{\sim} \tilde{P}_n$. We will use V_n to construct an estimator v_n for p_n .

Because $\tilde{P}_n \in \mathcal{H}_{d,b}^k$ ^{¶¶} for all n and our contradiction hypothesis we have that $\|V_n - \tilde{P}_n\|_1 \xrightarrow{p} 0$.

From this it follows that $\|U_{d,b}^{-1}(V_n) - U_{d,b}^{-1}(\tilde{P}_n)\|_1 \xrightarrow{p} 0$. Note that $[U_{d,b}^{-1}(\tilde{P}_n)]_{j,\ell,A} = p_n(j,\ell)$ when $A = (1, \dots, 1)$ and zero otherwise (see (11)). We define the linear operator $B_n : \mathcal{T}_{d,b} \rightarrow \Delta_{b \times k}$ as

$$[B_n(T)]_{j,\ell} \triangleq \sum_{A \in [b]^{d-2}} T_{j,\ell,A}$$

i.e. the linear operator which sums out all modes except for the first two. We have that $B_n(U_{d,b}^{-1}(\tilde{P}_n)) = p_n$. Now let $v_n = B_n(U_{d,b}^{-1}(V_n))$ be the estimator for p_n . Now we have that

$$\|v_n - p_n\|_1 = \|B_n(U_{d,b}^{-1}(\tilde{P}_n)) - B_n(U_{d,b}^{-1}(V_n))\|_1 = \|B_n(U_{d,b}^{-1}(\tilde{P}_n - V_n))\|_1.$$

We have that B_n is an ℓ^1 -nonexpansive operator due to the triangle inequality,

$$\|B_n(T)\|_1 = \sum_{j,\ell} \left| \sum_{A \in [b]^{d-2}} T_{j,\ell,A} \right| \leq \sum_{j,\ell} \sum_{A \in [b]^{d-2}} |T_{j,\ell,A}| = \|T\|_1,$$

^{¶¶}We will use this portion of the proof again for our proof of Theorem 2.4

so the operator norm of B_n is less than or equal to one. We also know that $U_{d,b}^{-1}$ an isometry and $\|\tilde{P}_n - V_n\|_1 \xrightarrow{P} 0$, so it follows that $\|v_n - p_n\|_1 \xrightarrow{P} 0$ for any sequence of $p_n \in \Delta_{[b(n)] \times [k(n)]}$. We will now use following theorem from [17] to show that no such estimator v_n can exist.

Theorem B.1 ([17] Theorem 2.). *For any $\zeta \in (0, 1]$, we have*

$$\inf_{\hat{p}} \sup_{p \in \Delta_a} \mathbb{E}_p \|\hat{p} - p\|_1 \geq \frac{1}{8} \sqrt{\frac{ea}{(1+\zeta)n}} \mathbb{1} \left(\frac{(1+\zeta)n}{a} > \frac{e}{16} \right) \\ + \exp \left(-\frac{2(1+\zeta)n}{a} \right) \mathbb{1} \left(\frac{(1+\zeta)n}{a} \leq \frac{e}{16} \right) - \exp \left(-\frac{\zeta^2 n}{24} \right) - 12 \exp \left(-\frac{\zeta^2 a}{32(\log a)^2} \right)$$

where the infimum is over all estimators.

Our estimator is equivalent to estimating a categorical distribution with $a = bk$ categories. Letting $\zeta = 1$, $bk \rightarrow \infty$, and $n \rightarrow \infty$, with $n/(bk) \rightarrow 0$, we get that for sufficiently large n

$$\inf_{\hat{p}} \sup_{p \in \Delta_{bk}} \mathbb{E}_p \|\hat{p} - p\|_1 \geq \exp \left(-\frac{4n}{bk} \right) - \exp \left(-\frac{n}{24} \right) - 12 \exp \left(-\frac{bk}{32(\log bk)^2} \right)$$

whose right hand side converges to 1. From this we get that

$$\liminf_{n \rightarrow \infty} \sup_{p_n \in \Delta_{bk}} \mathbb{E}_{p_n} \|v_n - p_n\|_1 > \frac{1}{2}$$

which contradicts $\|v_n - p_n\|_1 \xrightarrow{P} 0$ for arbitrary sequences p_n . \square

Proof of Theorem 2.4. We will proceed by contradiction. Suppose V_n is an estimator violating the theorem statement, i.e. there exist sequences $b \rightarrow \infty$ and $k \rightarrow \infty$ with $n/(bk + k^d) \rightarrow 0$ and $b \geq k$ such that, for all $\varepsilon > 0$,

$$\sup_{p \in \mathcal{D}_d} P \left(\|V_n - p\|_1 > 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 + \varepsilon \right) \rightarrow 0.$$

Since $n/(bk + k^d) \rightarrow 0$ we have that $(bk + k^d)/n \rightarrow \infty$ so there is a subsequence n_i such that $b(n_i)k(n_i)/n_i \rightarrow \infty$ or $k(n_i)^d/n_i \rightarrow \infty$, or equivalently $n_i/(b(n_i)k(n_i)) \rightarrow 0$ or $n_i/k(n_i)^d \rightarrow 0$. We will show that both cases lead to a contradiction. We will let b and k be functions of n_i implicitly when defining limits.

Case $n_i/(bk) \rightarrow 0$: We proceed similarly to the proof of Theorem 2.3. Let $(p_n)_{n=1}^\infty$, \tilde{P}_n , and \mathcal{X}_n be defined as in the proof of Theorem 2.3. Note that $\mathcal{H}_{d,b}^k \subset \tilde{\mathcal{H}}_{d,b}^k$ (see proof of Lemma 2.2) and thus $\tilde{P}_n \in \tilde{\mathcal{H}}_{d,b}^k$. We can proceed exactly as in our proof of Theorem 2.3 at footnote [11], by simply replacing $\mathcal{H}_{d,b}^k$ with $\tilde{\mathcal{H}}_{d,b}^k$ and n with n_i which finishes this case.

Case $n_i/k^d \rightarrow 0$: Let $(p_n)_{n=1}^\infty$ be a sequence of elements in $\mathcal{T}_{d,k}$ which represents distributions over $[k]^d$. Let $\mathcal{X}_n \triangleq (X_{n,1}, \dots, X_{n,n})$ with $X_{n,1}, \dots, X_{n,n} \stackrel{iid}{\sim} p_n$. Let $\tilde{\mathcal{X}}_n = (\tilde{X}_{n,1}, \dots, \tilde{X}_{n,n})$ which are independent random variables with $\tilde{X}_{n,i} \sim h_{d,b,X_{n,i}}$. Let \tilde{P}_n be the density for $\tilde{X}_{n,i}$. Note that $k \leq b$. So we have that

$$\begin{aligned} \tilde{P}_n &= \sum_{S \in [k]^d} p_{\tilde{X}_{n,i} | X_{n,i}=S} P(X_{n,i} = S) \\ &= \sum_{S \in [k]^d} h_{d,b,S} p_n(S) \\ &= \sum_{S \in [k]^d} p_n(S) \prod_{i=1}^d h_{1,b,S_i} \end{aligned}$$

and thus $\tilde{P}_n \in \tilde{\mathcal{H}}_{d,b}^k$. We proceed as in Theorem 2.3 to find an estimator for elements of $\mathcal{T}_{d,k}$ which is equivalent to estimating elements of Δ^{k^d} which is impossible since $n_i/k^d \rightarrow 0$. \square

C Finite Sample Theoretical Results

In this section we cover the proofs for the finite-sample results in our paper, Section 2.2.2 in the main text. This includes proofs related to estimator bias on Lipschitz continuous functions. We note that all projections in this paper are in their respective L^2 space. As noted in the main text we will be using *set projection*, $\text{Proj}_S x \triangleq \arg \min_{s \in S} \|x - s\|_2$ [7]. When S is a (closed) linear subspace this is equivalent to a linear projection. Every instance of set projection in this work yields a unique minimizer. As in the main text let Lip_L be the set of L -Lipschitz continuous densities on $[0, 1]$ and let $mL = \sup_{f \in \text{Lip}_L} \|f\|_2$.

C.1 Bias

Theorem 2.5 in the main text but is not used as is in any of our other proofs and is meant simply to be illustrative of the behavior of the bias as in the main text. We will get the proof of this theorem out of the way before moving on to the core results of this portion of the proofs. Note that λ is the standard Lebesgue measure.

Proof of Theorem 2.5. From Hölder's Inequality we have that for any function $f : [0, 1]^d \rightarrow \mathbb{R}$ that

$$\|f\|_1 = \|f \cdot \mathbb{1}\|_1 \leq \|f\|_2 \|\mathbb{1}\|_2 = \|f\|_2. \quad (13)$$

Applying this directly to the inequality from Theorem C.1 we have

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 \leq m_L^{2d} - \left(m_L^2 - \frac{L^2}{12b^2} \right)^d.$$

□

We now develop the core results needed for the paper. Because of the consequence of Hölder's Inequality (13) we will focus mainly on integrated squared distance between functions.

Lemma C.1. *Let $f : [a, b] \rightarrow \mathbb{R}$ be an L -Lipschitz continuous function. Then*

$$\min_{\alpha \in \mathbb{R}} \int_a^b (\alpha - f(x))^2 dx \leq \frac{L^2 (b-a)^3}{12}.$$

Proof of Lemma C.1. Let $\alpha = f\left(\frac{a+b}{2}\right)$. Then we have that

$$\begin{aligned} \int_a^b (\alpha - f(x))^2 dx &\leq \int_a^b \left(L \left| x - \frac{a+b}{2} \right| \right)^2 dx \\ &= L^2 \int_a^b \left(x - \frac{a+b}{2} \right)^2 dx \\ &= L^2 \int_{a-\frac{a+b}{2}}^{b-\frac{a+b}{2}} x^2 dx \\ &= L^2 \int_{\frac{a-b}{2}}^{\frac{b-a}{2}} x^2 dx \\ &= \frac{L^2}{3} [x^3]_{\frac{a-b}{2}}^{\frac{b-a}{2}} \\ &= \frac{L^2}{3} 2 \left(\frac{b-a}{2} \right)^3 \\ &= \frac{L^2 (b-a)^3}{12}. \end{aligned}$$

□

Lemma C.2. *Let f be an L -Lipschitz function on $[0, 1]$, then*

$$\left\| f - \text{Proj}_{\text{span}(\mathcal{H}_{1,b})} f \right\|_2^2 \leq \frac{L^2}{12b^2}.$$

Proof of Lemma C.2. Applying Lemma C.1 we have that

$$\begin{aligned} \left\| f - \text{Proj}_{\text{span}(\mathcal{H}_{1,b})} f \right\|_2^2 &= \min_{w \in \mathbb{R}^b} \left\| f - \sum_{i=1}^b w_i h_{1,b,i} \right\|_2^2 \\ &= \min_{w \in \mathbb{R}^b} \left\| f - \sum_{i=1}^b w_i b \mathbb{1}_{\left[\frac{i-1}{b}, \frac{i}{b}\right)} \right\|_2^2 \\ &= \min_{w \in \mathbb{R}^b} \int_{[0,1]} \left(f(x) - \sum_{i=1}^b w_i \mathbb{1} \left(\frac{i-1}{b} \leq x < \frac{i}{b} \right) \right)^2 dx \\ &= \min_{w \in \mathbb{R}^b} \int_{[0,1]} \left(\sum_{i=1}^b (f(x) - w_i) \mathbb{1} \left(\frac{i-1}{b} \leq x < \frac{i}{b} \right) \right)^2 dx \quad (14) \end{aligned}$$

$$\begin{aligned} &= \min_{w \in \mathbb{R}^b} \int_{[0,1]} \sum_{i=1}^b \left((f(x) - w_i) \mathbb{1} \left(\frac{i-1}{b} \leq x < \frac{i}{b} \right) \right)^2 dx \quad (15) \\ &= \min_{w \in \mathbb{R}^b} \sum_{i=1}^b \int_{\frac{i-1}{b}}^{\frac{i}{b}} (f(x) - w_i)^2 dx \\ &\leq b \frac{L^2}{12b^3} = \frac{L^2}{12b^2}, \end{aligned}$$

where (14) to (15) is justified since, when distributing the square, the cross terms of the form $\mathbb{1} \left(\frac{i-1}{b} \leq x < \frac{i}{b} \right) \mathbb{1} \left(\frac{j-1}{b} \leq x < \frac{j}{b} \right)$ are equal to zero when $i \neq j$. \square

Lemma C.3. *Let $f_1, \dots, f_d \in L^2([0, 1])$. Then*

$$\text{Proj}_{\text{span}(\mathcal{H}_{d,b})} \prod_{i=1}^d f_i = \prod_{i=1}^d \text{Proj}_{\text{span}(\mathcal{H}_{1,b})} f_i.$$

Proof of Lemma C.3. We will be using a multi-index $A = (i_1, \dots, i_d) \in [b]^d$ so that $\prod_{j=1}^d h_{1,b,i_j} = h_{d,b,A}$ and note that $h_{d,b,A} \perp h_{d,b,A'}$ for $A \neq A'$, so $h_{d,b,A}$ is an *orthogonal* basis for $\mathcal{H}_{d,b}$. We have the following

$$\begin{aligned} \prod_{i=1}^d \text{Proj}_{\text{span}(\mathcal{H}_{1,b})} f_i &= \prod_{i=1}^d \sum_{j=1}^b \frac{h_{1,b,j}}{\|h_{1,b,j}\|_2^2} \langle h_{1,b,j}, f_i \rangle \\ &= \left(\frac{h_{1,b,1}}{\|h_{1,b,1}\|_2^2} \langle h_{1,b,1}, f_1 \rangle + \dots + \frac{h_{1,b,b}}{\|h_{1,b,b}\|_2^2} \langle h_{1,b,b}, f_1 \rangle \right) \otimes \\ &\quad \dots \otimes \left(\frac{h_{1,b,1}}{\|h_{1,b,1}\|_2^2} \langle h_{1,b,1}, f_d \rangle + \dots + \frac{h_{1,b,b}}{\|h_{1,b,b}\|_2^2} \langle h_{1,b,b}, f_d \rangle \right). \quad (16) \end{aligned}$$

Now, distributing the terms in (16) and consolidating the subscripts into A we get that (16) is equal to

$$\begin{aligned}
\sum_{A \in [b]^d} \prod_{i=1}^d \frac{h_{1,b,A_i}}{\|h_{1,b,A_i}\|_2^2} \langle h_{1,b,A_i}, f_i \rangle &= \sum_{A \in [b]^d} \frac{\prod_{i=1}^d h_{1,b,A_i}}{\prod_{i=1}^d \|h_{1,b,A_i}\|_2^2} \prod_{i=1}^d \langle h_{1,b,A_i}, f_i \rangle \\
&= \sum_{A \in [b]^d} \frac{h_{d,b,A}}{\|h_{d,b,A}\|_2^2} \left\langle h_{d,b,A}, \prod_{i=1}^d f_i \right\rangle \\
&= \text{Proj}_{\text{span}(\mathcal{H}_{d,b})} \prod_{i=1}^d f_i.
\end{aligned}$$

□

Lemma C.4. *Let $f \in L^2([0, 1]^d)$ be a probability density. Then*

$$\text{Proj}_{\text{span}(\mathcal{H}_{d,b})} f = \text{Proj}_{\mathcal{H}_{d,b}} f.$$

Proof of Lemma C.4. Since $h_{d,b,A} \perp h_{d,b,A'}$ for $A \neq A'$ we have

$$\text{Proj}_{\text{span}(\mathcal{H}_{d,b})} f = \sum_{A \in [b]^d} \frac{h_{d,b,A}}{\|h_{d,b,A}\|_2^2} \langle h_{d,b,A}, f \rangle$$

and

$$\text{Proj}_{\text{span}(\mathcal{H}_{d,b})} f = \sum_{A \in [b]^d} w_A h_{d,b,A}$$

where

$$w_A = \frac{\langle h_{d,b,A}, f \rangle}{\|h_{d,b,A}\|_2^2}.$$

Clearly $w_A \geq 0$ so to finish we need only show that $\sum_{A \in [b]^d} w_A = 1$. To this end we have

$$\begin{aligned}
\sum_{A \in [b]^d} w_A &= \sum_{A \in [b]^d} \frac{\langle h_{d,b,A}, f \rangle}{\|h_{d,b,A}\|_2^2} \\
&= \sum_{A \in [b]^d} \int_{[0,1]^d} b^d \mathbb{1}(x \in \Lambda_{d,b,A}) f(x) dx / b^d \\
&= \int_{[0,1]^d} \sum_{A \in [b]^d} \mathbb{1}(x \in \Lambda_{d,b,A}) f(x) dx \\
&= \int_{[0,1]^d} f(x) dx \\
&= 1.
\end{aligned}$$

□

Corollary C.1. *Let $f_1, \dots, f_d \in L^2([0, 1])$ be probability densities, then*

$$\text{Proj}_{\mathcal{H}_{d,b}} \prod_{i=1}^d f_i = \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} f_i = \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i.$$

Proof of Corollary C.1. The first equality follows from Lemmas C.3 and C.4. To second equality follows from the first equality, the observation that $\mathcal{H}_{d,b}^1 \subset \mathcal{H}_{d,b}$, and that Lemma C.4 implies $\prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} f_i \in \mathcal{H}_{d,b}^1$. □

Theorem C.1. Let $m_L = \sup_{f \in \text{Lip}_L} \|f\|_2$ and let $b^2 \geq L^2/12$ then, for any $f_1, \dots, f_d \in \text{Lip}_L$, we have that

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 \leq m_L^{2d} - \left(m_L^2 - \frac{L^2}{12b^2} \right)^d.$$

Proof of Theorem C.1. We have

$$\begin{aligned} \left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 &= \left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}} \prod_{i=1}^d f_i \right\|_2^2 && \text{Corollary C.1} \\ &= \left\| \prod_{i=1}^d f_i \right\|_2^2 - \left\| \text{Proj}_{\mathcal{H}_{d,b}} \prod_{i=1}^d f_i \right\|_2^2 && \text{(Lemma C.4 with linear projection)} \\ &= \left\| \prod_{i=1}^d f_i \right\|_2^2 - \left\| \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 && \text{Corollary C.1} \\ &= \prod_{i=1}^d \|f_i\|_2^2 - \prod_{i=1}^d \left\| \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 && \text{Def'n of tensor product norm.} \end{aligned} \tag{17}$$

Noting that $\left\| \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 = \|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2$ (rearrangement of the property used in (17)) we have that

$$\left\| \prod_{i=1}^d f_i - \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{d,b}} f_i \right\|_2^2 = \prod_{i=1}^d \|f_i\|_2^2 - \prod_{i=1}^d \left(\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \right). \tag{18}$$

We will now turn our attention to the subtrahend on the right hand side of the previous equation. Note that the product terms $\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 = \left\| \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 > 0$ so $\prod_{i=1}^d \left(\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \right)$ is a product of positive values. From Lemmas C.2 and C.4 we have that $\left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \leq \frac{L^2}{12b^2}$. So we have

$$\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \geq \|f_i\|_2^2 - \frac{L^2}{12b^2}.$$

From Hölder's Inequality we have that

$$\|f \cdot \mathbb{1}\|_1 \leq \|f\|_2 \|\mathbb{1}\|_2 \Rightarrow \|f\|_2 \geq 1.$$

Combining this with the hypothesis $\frac{L^2}{12b^2} \leq 1$ we have that

$$\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \geq \|f_i\|_2^2 - \frac{L^2}{12b^2} \geq 0.$$

For a pair of tuples $b_i \geq a_i \geq 0$ it follows that $\prod b_i \geq \prod a_i$ and thus we have that

$$\prod_{i=1}^d \left(\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \right) \geq \prod_{i=1}^d \left(\|f_i\|_2^2 - \frac{L^2}{12b^2} \right).$$

Plugging this back into the RHS of (18) we get

$$\prod_{i=1}^d \|f_i\|_2^2 - \prod_{i=1}^d \left(\|f_i\|_2^2 - \left\| f_i - \text{Proj}_{\mathcal{H}_{1,b}} f_i \right\|_2^2 \right) \leq \prod_{i=1}^d \|f_i\|_2^2 - \prod_{i=1}^d \left(\|f_i\|_2^2 - \frac{L^2}{12b^2} \right). \tag{19}$$

For some arbitrary j we perform the following derivative

$$\frac{d}{d\|f_j\|_2^2} \left(\prod_{i=1}^d \|f_i\|_2^2 - \prod_{i=1}^d \left(\|f_i\|_2^2 - \frac{L^2}{12b^2} \right) \right) = \prod_{i \neq j} \|f_i\|_2^2 - \prod_{i \neq j} \left(\|f_i\|_2^2 - \frac{L^2}{12b^2} \right) \geq 0.$$

Thus we can find an upper bound to the RHS of (19) by maximizing $\|f_i\|_2^2$ over f_i , thus yielding

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}} \prod_{i=1}^d f_i \right\|_2^2 \leq m_L^{2d} - \left(m_L^2 - \frac{L^2}{12b^2} \right)^d.$$

□

Calculating the value of M_L is quite involved and is done in the next subsection. However we will first present the following result which gives a more tractable bound on bias.

Proposition C.1. *Let $L \geq 2$, $b^2 \geq L^2/12$, and let f_1, \dots, f_d be elements of Lip_L . Then*

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 \leq \frac{dL^{\frac{d+3}{2}}}{12b^2} \left[\frac{\sqrt{8}}{3} \right]^{d-1}.$$

If $0 \leq L \leq 2$ instead of $L \geq 2$, then

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 \leq d \frac{L^2}{12b^2} \exp \left(\frac{(d-1)L^2}{12} \right).$$

Proof of Proposition C.1. For the $L \geq 2$ case, applying Theorem C.1 and Proposition C.3 gives us

$$\left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 \leq \left[\frac{2\sqrt{2L}}{3} \right]^d - \left[\frac{2\sqrt{2L}}{3} - \frac{L^2}{12b^2} \right]^d.$$

We will use the identity $a^n - b^n = (a-b) \sum_{i=0}^{n-1} a^i b^{n-1-i}$ with $a = \frac{2\sqrt{2L}}{3}$ and $b = \frac{2\sqrt{2L}}{3} - \frac{L^2}{12b^2}$. Note that $a > b > 0$. From $a^d - b^d = (a-b) \sum_{i=0}^{d-1} a^i b^{d-1-i}$ we have

$$a^d - b^d = (a-b) \sum_{i=0}^{d-1} a^i b^{d-1-i} < (a-b) \sum_{i=0}^{d-1} a^i a^{d-1-i} = d(a-b)a^{d-1}$$

and thus

$$\begin{aligned} \left[\frac{2\sqrt{2L}}{3} \right]^d - \left[\frac{2\sqrt{2L}}{3} - \frac{L^2}{12b^2} \right]^d &\leq d \frac{L^2}{12b^2} \left[\frac{2\sqrt{2L}}{3} \right]^{d-1} \\ &= \frac{dL^{\frac{d+3}{2}}}{12b^2} \left[\frac{\sqrt{8}}{3} \right]^{d-1}. \end{aligned}$$

The other case proceeds analogously. If $L = 0$ the proposition statement is trivial, so we can assume $L > 0$ and thus $a > b > 0$ as before. Now we can finish with

$$\begin{aligned} \left\| \prod_{i=1}^d f_i - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d f_i \right\|_2^2 &\leq \left[\frac{L^2}{12} + 1 \right]^d - \left[\frac{L^2}{12} + 1 - \frac{L^2}{12b^2} \right]^d \\ &\leq d \frac{L^2}{12b^2} \left[\frac{L^2}{12} + 1 \right]^{d-1} \\ &\leq d \frac{L^2}{12b^2} \exp \left(\frac{(d-1)L^2}{12} \right). \end{aligned}$$

□

We find that the rate of the convergence of this bias term doesn't depend on the Lipschitz constants nor dimension and always has order $O(b^{-2})$. Applying a square root and applying Hölder's Inequality as we did in (13) we find that the L^1 bias convergence rate is $O(b^{-1})$, regardless of dimension.

C.1.1 Maximizing the L^2 norm of a Lipschitz density

In this section, we compute the value of m_L , the largest possible L^2 norm of an L -Lipschitz density function on $[0, 1]$. We will show that if $L \leq 2$, $m_L^2 = L^2/12 + 1$, and if $L \geq 2$, $m_L^2 = 2\sqrt{2L}/3$ (i.e. Proposition 2.2). This will follow directly from the more general Proposition C.3 below.

We will need several lemmas, starting with a discretization of the problem.

For the following results we are going to use a notion of *monotonic rearrangement*. For a finite sequence of real numbers (y_1, \dots, y_n) its rearrangement $R(y) = \tilde{y} = (\tilde{y}_1, \dots, \tilde{y}_n)$ which is a re-ordering of y so that it is increasing, it essentially sorts y to be ascending. Interestingly this concept can also be applied to functions [22]. It is usually defined so that the function is decreasing for the functional case. Interestingly if one monotonically rearranges a Lipschitz continuous function, $f : \mathbb{R}^+ \rightarrow \mathbb{R}$ for example, then its rearrangement will also be Lipschitz continuous with a Lipschitz constant no larger than that for f (and potentially smaller). The following lemma says that the discrete equivalent of the “Lipschitz constant” of a sequence is always higher than that of its monotone reordering.

Lemma C.5. *Let $y = (y_1, \dots, y_n)$ be a sequence of real numbers, and let $\tilde{y} = (\tilde{y}_1, \dots, \tilde{y}_n)$ be its monotone reordering so that $\tilde{y}_1 \leq \tilde{y}_2 \leq \dots \leq \tilde{y}_n$. We have*

$$\max_{i \leq n-1} |y_{i+1} - y_i| \geq \max_{i \leq n-1} \tilde{y}_{i+1} - \tilde{y}_i.$$

Proof of Lemma C.5. If y is constant or $n = 1$ then we are done, so we will assume y is nonconstant. Let i^* be such that $\tilde{y}_{i^*+1} - \tilde{y}_{i^*} = \max_{i \leq n-1} \tilde{y}_{i+1} - \tilde{y}_i$. Define the sets

$$\begin{aligned} A &:= \{i \in \{1, 2, \dots, n\} : y_i \leq \tilde{y}_{i^*}\} \\ B &:= \{i \in \{1, 2, \dots, n\} : y_i \geq \tilde{y}_{i^*+1}\}. \end{aligned}$$

These sets partition $[n]$ so $A \sqcup B = \{1, 2, \dots, n\}$ and A and B are non empty. Thus there must exist $j \in [n]$ such that $j \in A$ and $j+1 \in B$ or vice versa. For the first case we have that

$$y_{j+1} - y_j \geq \tilde{y}_{i^*+1} - \tilde{y}_{i^*} = \max_{i \leq n-1} \tilde{y}_{i+1} - \tilde{y}_i$$

and for the vice versa case we have

$$|y_{j+1} - y_j| = y_j - y_{j+1} \geq \tilde{y}_{i^*+1} - \tilde{y}_{i^*} = \max_{i \leq n-1} \tilde{y}_{i+1} - \tilde{y}_i.$$

So we have demonstrated that there exists a j such that $|y_{j+1} - y_j| \geq \max_{i \leq n-1} \tilde{y}_{i+1} - \tilde{y}_i$. \square

We are now in a position to prove the following discrete version of Proposition C.3 below.

Proposition C.2. *Let $n \in \mathbb{N}_{\geq 2}$, $U > 0$, $L > 0$, and let S_n denote the set of sequences $y = (y_1, \dots, y_n)$ such that the following conditions are satisfied:*

$$y_i \geq 0 \quad (\forall i) \tag{20}$$

$$|y_{i+1} - y_i| \leq L \quad (\forall 1 \leq i \leq n-1) \tag{21}$$

$$\sum_{i=1}^n y_i = U. \tag{22}$$

A sequence $y^ = (y_1, \dots, y_n) \in S_n$, which maximizes the quantity $\sum_{i=1}^n y_i^2$ subject to the conditions above is given by the following formulae.*

If $L \leq \frac{2U}{n(n-1)}$, for all i ,

$$y_i^* = (i-1)L + [U/n - (n-1)L/2]. \tag{23}$$

If $L \geq \frac{2U}{n(n-1)}$, for all i ,

$$y_i^* = ((i-n+m)L + U/m - L(m+1)/2)_+, \tag{24}$$

where x_+ denotes the positive part of x and m is the smallest integer such that $m(m+1)L \geq 2U$.

Proof of Proposition C.2. First note that S_n is compact since the intersection of the closed sets denoted by (20) and (22) is compact and the set denoted by (21) is closed and thus further intersecting with that set gives us compact set. Because $y \mapsto \sum_{i=1}^n y_i^2 = \|y\|_2^2$ is continuous it must therefore attain a maximum on S_n . Let $\tilde{\zeta} \in S_n$ be some sequence which attains the maximum. Define the sequence ζ as a monotone increasing (non-decreasing) rearrangement of $\tilde{\zeta}$. By Lemma C.5, we have $\zeta \in S_n$ and clearly $\|\zeta\|_2^2 = \|\tilde{\zeta}\|_2^2$ so ζ is also a maximizer.

Consider the following class of sequences.

\mathfrak{L} -sequences: A sequence $x = (x_1, \dots, x_n)$ is an \mathfrak{L} -sequence if

1. $x \in S_n$.
2. x is non-decreasing.
3. For all $i < n$, $x_i = 0$ or $x_{i+1} - x_i = L$.

We are first going to show that y^* is the only element in this class.

Consider some arbitrary \mathfrak{L} -sequence y be and let $k = \min(i : y_i \neq 0)$. We have $y_i - y_k = L(i - k)$ for all $i \geq k$ and therefore

$$\begin{aligned} U &= \sum_{i=1}^n y_i \\ &= \sum_{i=k}^n y_k + L \sum_{i=k+1}^n (i - k) \end{aligned} \tag{25}$$

$$\Rightarrow U = (n - k + 1)y_k + \frac{L(n - k)(n - k + 1)}{2}. \tag{26}$$

One can check that (26) holds when $k = n$, even though the summation in (25) is ill-posed. We will now split into two cases, corresponding to (23) and (24), to show that there is only one \mathfrak{L} -sequence for fixed U, L, n and that this sequence is given by y^* .

Case 1: $\frac{Ln(n-1)}{2} < U$.

If $\frac{Ln(n-1)}{2} < U$ then (26) cannot hold unless $k = 1$ because if $k \geq 2$ then $y_k \leq L$ (since $y_{k-1} = 0$) and

$$\begin{aligned} U &= (n - k + 1)y_k + \frac{L(n - k)(n - k + 1)}{2} \\ &\leq (n - 1)L + \frac{L(n - 2)(n - 1)}{2} \\ &= \frac{Ln(n - 1)}{2} < U, \end{aligned}$$

a contradiction.

Solving for y_k in (26) with $k = 1$ we get

$$y_k = y_1 = \frac{U}{n} - \frac{L(n - 1)}{2},$$

so k and y_k are both unique, thus S_n contains only one element. We see that this coincides with the expression for y^* in (23) by noting that the derived y_1 is equal to y_1^* . The situation where $L = \frac{2U}{n(n-1)}$ in (23) will be addressed at the end of Case 2.

Case 2: $\frac{Ln(n-1)}{2} \geq U$.

For this case we need $k \geq 2$ since $y_k > 0$ and letting $k = 1$ in (26) yields

$$\begin{aligned} U &= (n - k + 1)y_k + \frac{L(n - k)(n - k + 1)}{2} \\ &> \frac{Ln(n - 1)}{2} \geq U \end{aligned}$$

a contradiction. This implies $y_{k-1} = 0$, so $y_k \leq L$. Define $m \triangleq n - k + 1$. From (26) we have

$$\begin{aligned} U &= (n - k + 1)y_k + \frac{L(n - k)(n - k + 1)}{2} \\ &= my_k + \frac{L(m - 1)m}{2} \\ &= m \left(y_k + \frac{L(m - 1)}{2} \right). \end{aligned} \tag{27}$$

Since $y_k > 0$ we have $U \geq \frac{Lm(m-1)}{2}$. Additionally since $y_k \leq L$,

$$\begin{aligned} U &\leq m \left(L + \frac{L(m - 1)}{2} \right) \\ &= \frac{Lm(m + 1)}{2}. \end{aligned}$$

It now follows that m must be the smallest integer such that $\frac{Lm(m+1)}{2} \geq U$.

Solving for y_k in (27) we have

$$y_k = \frac{U}{m} - \frac{L(m - 1)}{2},$$

so both k and y_k are unique so there is only one element in S_n . Further note that

$$\begin{aligned} y_k &= \frac{U}{m} - \frac{L(m - 1)}{2} \\ &= \frac{U}{m} - \frac{L(m + 1)}{2} + L. \end{aligned}$$

From the definition of S_n it follows that

$$\begin{aligned} y_i &= (L(i - k) + y_k)_+ \\ &= \left(L(i - n + m - 1) + \frac{U}{m} - \frac{L(m + 1)}{2} + L \right)_+ \\ &= \left(L(i - n + m) + \frac{U}{m} - \frac{L(m + 1)}{2} \right)_+ \end{aligned}$$

which coincides with y^* from (24).

To finish this case we will show that when $\frac{Ln(n-1)}{2} = U$ then (23) equals (24). To see this first note that $\frac{Ln(n-1)}{2} = U$ implies $m = n - 1$. Therefore (24) equals

$$\begin{aligned} y_i^* &= \left((i - n + m)L + \frac{U}{m} - \frac{L(m + 1)}{2} \right)_+ \\ &= \left((i - 1)L + \frac{U}{n - 1} - \frac{Ln}{2} \right)_+ \\ &= \left((i - 1)L + \frac{Ln}{2} - \frac{Ln}{2} \right)_+ \\ &= (i - 1)L \end{aligned}$$

and accordingly (23) equals

$$\begin{aligned} y_i^* &= (i - 1)L + \left[\frac{U}{n} - \frac{(n - 1)L}{2} \right] \\ &= (i - 1)L + \left[\frac{(n - 1)L}{2} - \frac{(n - 1)L}{2} \right] \\ &= (i - 1)L. \end{aligned}$$

This finishes Case 2.

We now have that y^* satisfies property 1 of \mathcal{L} -sequences (in particular (22) was nontrivial) as well as properties 2 and 3. This concludes the proof that y^* is the only \mathcal{L} -sequence.

Now we are going to show that ζ is also a \mathcal{L} -sequence which will complete this proof, since y^* is the only \mathcal{L} -sequence. For sake of contradiction assume that ζ is not an \mathcal{L} -sequence. From this there must exist $i_* \in \{1, 2, \dots, n-1\}$ such that $\zeta_{i_*+1} - \zeta_{i_*} < L$ and $\zeta_{i_*} \neq 0$. We then define a sequence t by:

$$t_i = \begin{cases} \zeta_i & \text{if } i \neq i_*, i_* + 1 \\ \zeta_{i_*} - \delta & \text{if } i = i_* \\ \zeta_{i_*+1} + \delta & \text{if } i = i_* + 1, \end{cases}$$

where $\delta = \min\left(\frac{L - \zeta_{i_*+1} + \zeta_{i_*}}{2}, \zeta_{i_*}\right)$. Note that $0 < \delta \leq L/2$, a fact which we will be using extensively.

We will now show that $t \in S_n$.

t satisfies (20): Note that by construction, since $\delta \leq \zeta_{i_*}$ and $t_{i_*} \geq 0$ (and clearly, $t_i \geq 0$ for $i \neq i_*$).

t satisfies (21): The sequences t and ζ differ on only two indices so (21) holds trivially for any pair of indices not involving i_* or $i_* + 1$. We will address the remaining cases:

Case 1 $|t_{i_*+1} - t_{i_*}|$: We now have that

$$\begin{aligned} t_{i_*+1} - t_{i_*} &= \zeta_{i_*+1} - \zeta_{i_*} + 2\delta \\ &\leq \zeta_{i_*+1} - \zeta_{i_*} + L - \zeta_{i_*+1} + \zeta_{i_*} \\ &= L. \end{aligned}$$

Since ζ is monotonic and $\delta > 0$ we further have that $\zeta_{i_*+1} - \zeta_{i_*} + 2\delta > 0$ thereby finishing this case.

Case 2 $|t_{i_*+2} - t_{i_*+1}|$: We have that

$$\begin{aligned} t_{i_*+2} - t_{i_*+1} &= \zeta_{i_*+2} - \zeta_{i_*+1} - \delta \\ &\geq -\delta \\ &\geq -L/2 \end{aligned}$$

and

$$t_{i_*+2} - t_{i_*+1} = (\zeta_{i_*+2} - \zeta_{i_*+1}) - \delta \leq L$$

thereby completing this case.

Case 3: $|t_{i_*} - t_{i_*-1}|$: This is virtually identical to the previous case so we omit it.

Thus, we have that t satisfies condition (21).

t satisfies (22): finally, it is clear that $t_{i_*} + t_{i_*+1} = \zeta_{i_*} + \zeta_{i_*+1}$ and thus $\sum_i t_i = \sum_i \zeta_i = U$.

Hence, t satisfies (22) and we have proved the claim that $t \in S_n$.

Now, writing θ for $\frac{\zeta_{i_*+1} + \zeta_{i_*}}{2}$ and Δ for $\frac{\zeta_{i_*+1} - \zeta_{i_*}}{2}$:

$$\begin{aligned} \sum_{i=1}^n (t_i^2 - \zeta_i^2) &= t_{i_*+1}^2 + t_{i_*}^2 - \zeta_{i_*+1}^2 - \zeta_{i_*}^2 \\ &= (\zeta_{i_*+1} + \delta)^2 + (\zeta_{i_*} - \delta)^2 - \zeta_{i_*+1}^2 - \zeta_{i_*}^2 \\ &= (\theta + \Delta + \delta)^2 + (\theta - \Delta - \delta)^2 - (\theta + \Delta)^2 - (\theta - \Delta)^2 \\ &= 2\theta^2 + 2(\Delta + \delta)^2 - [2\theta^2 + 2\Delta^2] = 4\Delta\delta + 2\delta^2 \geq 2\delta^2 > 0, \end{aligned}$$

where at the last line, we have used the fact that since by construction, since $\zeta \neq y^*$, we have $\delta > 0$. This shows that:

$$\sum_{i=1}^n t_i^2 > \sum_{i=1}^n \zeta_i^2,$$

a contradiction. Hence we conclude that ζ must equal y^* . \square

We can now proceed with the statement and proof of the continuous case.

Proposition C.3. *Let $L, U > 0$ be given and let S be the set of functions $f : [0, 1] \rightarrow \mathbb{R}^+$ such that the following conditions are satisfied:*

1. *f is Lipschitz continuous with Lipschitz constant L .*
2. *$\int_0^1 f(x)dx = U$.*

Let $f_^{U,L}$ be defined as follows*

$$f_*^{U,L}(x) = \left(x - \frac{1}{2}\right) L + U \quad \text{if } L \leq 2U \quad (28)$$

$$f_*^{U,L}(x) = L \left(x - 1 + \sqrt{\frac{2U}{L}}\right)_+ \quad \text{if } L > 2U. \quad (29)$$

Then $f_^{U,L} \in \arg \max_{f \in S} \int_0^1 f^2(x)dx$ and*

$$\max_{f \in S} \int_0^1 f^2(x)dx = M(U, L) \triangleq \frac{W^3 L^2}{12} + \frac{U^2}{W},$$

where $W = 1$ if $L \leq 2U$ and $W = \sqrt{\frac{2U}{L}}$ if $L > 2U$. Equivalently, $M = \frac{L^2}{12} + U^2$ if $L \leq 2U$ and $M = \frac{2\sqrt{2}}{3} \sqrt{L} U^{3/2}$ if $L \geq 2U$.

Remark (Proposition C.3 applied to probability densities). The highly general formation of Proposition C.3 is useful for its proof. However in density estimation only the case where f_* is a pdf and $U = 1$ is particularly meaningful. This gives

$$\begin{aligned} f_*(x) &= \left(x - \frac{1}{2}\right) L & \|f_*\|_2^2 &= \frac{L^2}{12} + 1 & \text{if } L \leq 2 \\ f_*(x) &= L \left(x - 1 + \sqrt{\frac{2}{L}}\right)_+ & \|f_*\|_2^2 &= \frac{2\sqrt{2}}{3} \sqrt{L} & \text{if } L > 2. \end{aligned}$$

Proof of Proposition C.3. We first introduce some notation. For $n \in \mathbb{N}$ and $U', L' \geq 0$ we will denote by $\mathcal{G}^n(U', L')$ the nonnegative sequence (y_1, y_2, \dots, y_n) maximizing $\sum_{i=1}^n y_i^2$ subject to $\sum_{i=1}^n y_i \frac{1}{n} = U'$ and $|y_{i+1} - y_i| \leq \frac{L'}{n}$ (i.e. $\mathcal{G}^n(U', L') = y^*$ from Proposition C.2 with $U \leftarrow U'n$ and $L \leftarrow \frac{L'}{n}$). We similarly write $M_n(U', L')$ for the optimal value, i.e., $M_n(U', L') = \sum_{i=1}^n (\mathcal{G}^n(U', L'))_i^2 \frac{1}{n}$. For ease of notation we set $\mathcal{G}^n \triangleq \mathcal{G}^n(U', L')$, $M_n \triangleq M_n(U', L')$, and $f_* = f_*^{U', L'}$. Later we will set $U' \leftarrow U$ and $L' \leftarrow L$ to prove the proposition statement, but it will be useful to establish some results for general U' and L' . We use the prime symbol to help avoid confusion between proving the final proposition statement and proving supporting results. Unless otherwise specified, all limits are taken as $n \rightarrow \infty$.

Let $\tilde{\mathcal{G}}^n$ be the piecewise linear function from $[0, 1]$ to \mathbb{R}^+ with $\tilde{\mathcal{G}}^n(i/n) = \mathcal{G}_i^n$ for $i \in [n]$ and $\tilde{\mathcal{G}}^n(0) = \mathcal{G}_1^n$.

Lemma C.6. *The following limits hold,*

$$\begin{aligned} \int_0^1 \tilde{\mathcal{G}}^n(x)dx &\rightarrow U' \\ \int_0^1 \tilde{\mathcal{G}}^n(x)^2 dx - M_n &\rightarrow 0. \end{aligned}$$

Proof of Lemma C.6. Note that the function $\tilde{\mathcal{G}}^n$ is L' -Lipschitz. The Lipschitz continuity also implies the following bounds on the difference between the Riemann sums below and their correspond-

ing integrals:

$$\begin{aligned}
\left| \int_0^1 \tilde{\mathcal{G}}^n(x) dx - U' \right| &= \left| \int_0^1 \tilde{\mathcal{G}}^n(x) dx - \sum_{i=1}^n \mathcal{G}_i^n \frac{1}{n} \right| \\
&\leq \sum_{i=1}^n \left| \int_{(i-1)/n}^{i/n} \tilde{\mathcal{G}}^n(x) dx - \mathcal{G}_i^n \frac{1}{n} \right| \\
&= \sum_{i=1}^n \left| \int_{(i-1)/n}^{i/n} \tilde{\mathcal{G}}^n(x) - \mathcal{G}_i^n dx \right| \\
&\leq \sum_{i=1}^n \int_{(i-1)/n}^{i/n} |\tilde{\mathcal{G}}^n(x) - \mathcal{G}_i^n| dx \\
&\leq \sum_{i=1}^n \int_{(i-1)/n}^{i/n} \frac{L'}{n} dx \\
&= n \frac{L'}{n^2} \rightarrow 0,
\end{aligned}$$

and also

$$\begin{aligned}
\left| \int_0^1 \tilde{\mathcal{G}}^n(x)^2 dx - M_n \right| &= \left| \int_0^1 \tilde{\mathcal{G}}^n(x)^2 dx - \sum_{i=1}^n (\mathcal{G}_i^n)^2 \frac{1}{n} \right| \\
&= \left| \sum_{i=1}^n \left[\int_{(i-1)/n}^{i/n} \tilde{\mathcal{G}}^n(x)^2 - (\mathcal{G}_i^n)^2 dx \right] \right| \\
&= \left| \sum_{i=1}^n \left[\int_{(i-1)/n}^{i/n} [\tilde{\mathcal{G}}^n(x) - (\mathcal{G}_i^n)] [\tilde{\mathcal{G}}^n(x) + (\mathcal{G}_i^n)] dx \right] \right| \\
&\leq n \frac{1}{n} \frac{L'}{n} \max_i [2\mathcal{G}_i^n] \\
&\leq \frac{L'}{n} [2(U' + L')] \rightarrow 0,
\end{aligned}$$

where at the last line we have used the fact that $\mathcal{G}_i^n \leq U' + L'$ for all i . This is because $\sum_{i=1}^n \mathcal{G}_i^n = U'n$ implies there exists an i_* such that $\mathcal{G}_{i_*}^n \leq U'$ and the Lipschitz condition further implies (for all i):

$$\mathcal{G}_i^n \leq \mathcal{G}_{i_*}^n + \frac{L'}{n} |i - i_*| \leq U' + L'. \quad (30)$$

□

Lemma C.7. *The sequence of functions $\tilde{\mathcal{G}}^n \rightarrow f_*$ pointwise.*

Proof of Lemma C.7. Given a continuous function $f : [0, 1] \rightarrow \mathbb{R}$, that is also differentiable on $(0, 1)$, for all $x \in [0, 1]$ the fundamental theorem of calculus implies

$$\begin{aligned}
\int_0^x \partial_y f(y) dy &= f(x) - f(0) \\
\Rightarrow f(x) &= \int_0^x \partial_y f(y) dy + f(0) \\
\Rightarrow |f(x)| &\leq \int_0^x |\partial_y f(y)| dy + |f(0)| \leq \int_0^1 |\partial_y f(y)| dy + |f(0)|.
\end{aligned}$$

Suppose f was not differentiable, but was still continuous, at some $c \in (0, x)$. We would still have

$$\begin{aligned} f(x) &= f(x) - f(c) + f(c) \\ &= \int_c^x \partial_y f(y) dy + \int_0^c \partial_y f(y) dy + f(0) \\ \Rightarrow |f(x)| &\leq \int_0^1 |\partial_y f(y)| dy + |f(0)|. \end{aligned}$$

Both f_* and $\tilde{\mathcal{G}}^n$ are piecewise linear. They are therefore continuous and their derivatives exist on all but a finite set of points, so for all $x \leq 1$,

$$|f_*(x) - \tilde{\mathcal{G}}^n(x)| \leq |f_*(0) - \tilde{\mathcal{G}}^n(0)| + \int_0^1 |\partial_y f_*(y) - \partial_y \tilde{\mathcal{G}}^n(y)| dy. \quad (31)$$

We now go on with the proof of Lemma C.7 in two cases.

Case 1: $L' \leq 2U'$

Note that if $L' \leq 2U'$, then for all $n \geq 2$ we have that $\frac{L'}{n} \leq \frac{2U'n}{n(n-1)}$ and thus \mathcal{G}^n is defined by (23). Then for any n we have $\tilde{\mathcal{G}}^n(i/n) - \tilde{\mathcal{G}}^n((i-1)/n) = L'/n$ for all $i \geq 2$, and $\tilde{\mathcal{G}}^n(1/n) - \tilde{\mathcal{G}}^n(0/n) = 0$. In particular, we have under these conditions that $\partial_x \tilde{\mathcal{G}}^n(x) = L'$ for all $x > 1/n$ and $\partial_x \tilde{\mathcal{G}}^n(x) = 0$ for $x < 1/n$. Hence by (31) we have for all $x \leq 1$,

$$\begin{aligned} |f_*(x) - \tilde{\mathcal{G}}^n(x)| &\leq |f_*(0) - \tilde{\mathcal{G}}^n(0)| + \int_0^1 |\partial_y f_*(y) - \partial_y \tilde{\mathcal{G}}^n(y)| dy \\ &\leq \left| U' - \frac{1}{2}L' - \left[U' - \frac{n-1}{n} \frac{L'}{2} \right] \right| + \int_0^{1/n} |\partial_y f_*(y) - \partial_y \tilde{\mathcal{G}}^n(y)| dy \\ &\leq \frac{L'}{2n} + \frac{L'}{n} \rightarrow 0. \end{aligned}$$

Case 2: $L' > 2U'$

Since $L' > 2U'$, for sufficiently large n we have that $\frac{L'}{n} > \frac{2U'n}{n(n-1)}$ and \mathcal{G}^n is defined using (24). Since we are interested the limit as $n \rightarrow \infty$ we will proceed using (24). As in the proof of Proposition C.2 we let k be the smallest natural number with $\mathcal{G}_k^n > 0$, noting also that $k \geq 2$ (see Case 2 in that proof).

For all $i \leq k-1$ we have $\tilde{\mathcal{G}}^n(i/n) - \tilde{\mathcal{G}}^n((i-1)/n) = 0$ so $\partial_x \tilde{\mathcal{G}}^n(x) = 0$ for all $x < (k-1)/n$. Similarly for all $i \geq k+1$ we have $\tilde{\mathcal{G}}^n(i/n) - \tilde{\mathcal{G}}^n((i-1)/n) = L'/n$ so $\partial_x \tilde{\mathcal{G}}^n(x) = L'$ for all $x > (k+1)/n$. We also have that $\tilde{\mathcal{G}}^n(0) = 0$.

By definition (29) we have $\partial_x f_*(x) = 0$ for all $x < \left[1 - \sqrt{\frac{2U'}{L'}}\right]$ and $\partial_x f_*(x) = L'$ for all $x > 1 - \sqrt{\frac{2U'}{L'}}$. We also have $f_*(0) = 0$.

Again from the proof of Proposition C.2 we know $k = n - m + 1$, where m is the smallest integer such that $m(m+1)L'/n \geq 2U'n$. It follows that $m(m+1)/n^2 \rightarrow 2U'/L'$. Since $m \rightarrow \infty$ and m^2/n^2 doesn't diverge we have that $m/n^2 \rightarrow 0$ and thus $m/n \rightarrow \sqrt{2U'/L'}$. Substituting in $k = n - m + 1$ gives $k/n \rightarrow 1 - \sqrt{2U'/L'}$.

Let $\varepsilon > 0$. For sufficiently large n both $(k+1)/n$ and $(k-1)/n$ lie in $\left[1 - \sqrt{2U'/L'} - \varepsilon, 1 - \sqrt{2U'/L'} + \varepsilon\right]$. Letting $x \leq 1$, for large enough n the following holds

by (31)

$$\begin{aligned}
|f_*(x) - \tilde{\mathcal{G}}^n(x)| &\leq |f_*(0) - \tilde{\mathcal{G}}^n(0)| + \int_0^1 |\partial_y f_*(y) - \partial_y \tilde{\mathcal{G}}^n(y)| dy \\
&= |f_*(0) - \tilde{\mathcal{G}}^n(0)| + \int_{1-\sqrt{2U'/L'}-\varepsilon}^{1-\sqrt{2U'/L'}+\varepsilon} |\partial_y f_*(y) - \partial_y \tilde{\mathcal{G}}^n(y)| dy \\
&\leq 2\varepsilon \left(\max_y |\partial_y f_*(y)| + |\partial_y \tilde{\mathcal{G}}^n(y)| \right) \\
&\leq 2\varepsilon (L' + L' + U') \quad \text{using (30).}
\end{aligned}$$

Since this holds for all $\varepsilon > 0$, $|f_*(x) - \tilde{\mathcal{G}}^n(x)|$ goes to zero, and we have pointwise convergence for Case 2. \square

Lemma C.8. *We have that $M_n \rightarrow \int_0^1 f_*(x)^2 dx$.*

Proof of Lemma C.8. From (30) we can bound $\max_x \tilde{\mathcal{G}}^n(x)^2 \leq (U' + L')^2$. This bound, along with the pointwise convergence of Lemma C.7, allows us to apply the dominated convergence theorem, thus giving $\int_0^1 \tilde{\mathcal{G}}^n(x)^2 dx \rightarrow \int_0^1 f_*(x)^2 dx$. Simply applying Lemma C.6 finishes our proof. \square

We can now proceed with the proof of the proposition. First we will calculate the integral $\int_0^1 f_*(x)^2 dx$, thus establishing that $M(U', L') = \int_0^1 f_*(x)^2 dx$. To avoid confusion and to be fully precise: M is defined by the \triangleq symbol in the proposition statement; the left equality in that equation being proven.

We define W' from U' and L' analogously to the way W is defined from U and L in the proposition statement. Note that f_* is linear on the interval $[1 - W', 1]$ and zero elsewhere. From Lemmas C.6 and C.7 and using the dominated convergence theorem as before we have that $\int_0^1 f_*(x) dx = U'$. Thus we can write

$$\begin{aligned}
M(U', L') &= \int_0^1 [f_*(x)]^2 dx \\
&= \int_{1-W'}^1 [f_*(x)]^2 dx \\
&= \int_{1-W'}^1 \left[f_*(x) - \frac{U'}{W'} + \frac{U'}{W'} \right]^2 dx \\
&= \int_{1-W'}^1 \left[f_*(x) - \frac{U'}{W'} \right]^2 dx + \int_{1-W'}^1 \left[\frac{U'}{W'} \right]^2 dx + 2 \int_{1-W'}^1 \left[f_*(x) - \frac{U'}{W'} \right] \left[\frac{U'}{W'} \right] dx \\
&= \int_{1-W'}^1 \left[f_*(x) - \frac{U'}{W'} \right]^2 dx + \int_{1-W'}^1 \left[\frac{U'}{W'} \right]^2 dx,
\end{aligned}$$

where at the last line we have used the fact that $\int_{1-W'}^1 f_*(x) - \frac{U'}{W'} dx = 0$.

Clearly $\int_{1-W'}^1 \left[\frac{U'}{W'} \right]^2 dx = \frac{U'^2}{W'}$. Further, $f_*|_{[1-W', 1]} - \frac{U'}{W'}$ is linear with slope L' and $f_*(1 - W'/2) - \frac{U'}{W'} = 0$, so $f_*|_{[1-W', 1]}(x) - \frac{U'}{W'}$ is antisymmetric about $x = 1 - \frac{W'}{2}$. Using

these facts we can continue,

$$\begin{aligned}
\int_0^1 [f_*(x)]^2 dx &= \int_{1-W'}^1 [f_*(x)]^2 dx \\
&= \frac{U'^2}{W'} + 2 \int_{1-\frac{W'}{2}}^1 \left[f_*(x) - \frac{U'}{W'} \right]^2 dx \\
&= \frac{U'^2}{W'} + 2 \int_{1-\frac{W'}{2}}^1 L'^2 \left(x - 1 + \frac{W'}{2} \right)^2 dx \\
&= \frac{U'^2}{W'} + 2L'^2 \int_0^{\frac{W'}{2}} x^2 dx \\
&= \frac{U'^2}{W'} + 2L'^2 \frac{(W'/2)^3}{3} = \frac{W'^3 L'^2}{12} + \frac{U'^2}{W'},
\end{aligned}$$

as expected.

We now fix $U' \leftarrow U$ and $L' \leftarrow L$. We have shown $\int_0^1 f_*(x) dx = U$ which, along direct inspection of the definition of f_* , establishes that f_* satisfies the properties of S . Let $g \in S$ be arbitrary. We will show that $\int_0^1 g(x)^2 dx \leq \int_0^1 f_*(x)^2 dx$, which demonstrates that $f_* \in \arg \max_{f \in S} \int_0^1 f^2(x) dx$ and finishes our proof.

For all $n \geq 2$, define the sequence $g^n = (g_1^n, \dots, g_n^n)$ with $g_i^n = g(\frac{i}{n})$. From the Lipschitz continuity of g we have as $n \rightarrow \infty$,

$$\sum_{i=1}^n g_i^n \frac{1}{n} \rightarrow \int_0^1 g(x) dx = U, \quad (32)$$

and

$$\sum_{i=1}^n (g_i^n)^2 \frac{1}{n} \rightarrow \int_0^1 g(x)^2 dx. \quad (33)$$

Let $\varepsilon > 0$ be arbitrary. By (32) we can set N such that $\sum_{i=1}^n g_i^n \frac{1}{n} \leq U + \varepsilon$ for all $n \geq N$. Note that $|g_{i+i}^n - g_i^n| \leq L/n$ and $g_i^n \geq 0$ for all i . Let $Q_n \triangleq \frac{1}{n} (n\varepsilon + nU - \sum_{i=1}^n g_i^n)$. For $n \geq N$ we have that $Q_n \geq 0$ and the following hold,

$$\begin{aligned}
g_i^n + Q_n &\geq 0 && \text{for all } i \\
|(g_i^n + Q_n) - (g_{i+1}^n + Q_n)| &\leq L/n && \text{for all } i \\
\frac{1}{n} \left(\sum_{i=1}^n g_i^n + Q_n \right) &= U + \varepsilon.
\end{aligned}$$

It follows that

$$\begin{aligned}
\int_0^1 g(x)^2 dx &= \lim_{n \rightarrow \infty} \sum_{i=1}^n (g_i^n)^2 \frac{1}{n} \\
&\leq \limsup_{n \rightarrow \infty} \sum_{i=1}^n (g_i^n + Q_n)^2 \frac{1}{n} \\
&\leq \lim_{n \rightarrow \infty} M_n(U + \varepsilon, L) \\
&= M(U + \varepsilon, L). \quad \text{Lemma C.8}
\end{aligned}$$

Since $M(\cdot, L)$ is continuous and ε was arbitrary we have that $\int_0^1 g(x)^2 dx \leq M(U, L)$. \square

C.1.2 Lower bound for Lipschitz densities in L^1 .

Lemma C.9. *Let $f : [a, b] \rightarrow \mathbb{R}$ be the function $f(x) = Lx + c$ for some $a, b, c, L \in \mathbb{R}$ with $a < b$. Then*

$$\min_{\alpha \in \mathbb{R}} \int_a^b |\alpha - f(x)| dx = \frac{L(b-a)^2}{4}, \quad (34)$$

and

$$\arg \min_{\alpha \in \mathbb{R}} \int_a^b |\alpha - f(x)| dx = L \frac{a+b}{2} + c = \frac{1}{b-a} \int_a^b f(x) dx. \quad (35)$$

Proof of Lemma C.9. When $L = 0$ the result follows from just setting $\alpha = c$. We will assume that $L \neq 0$.

Suppose that $\alpha > \max(f(a), f(b))$, then we have

$$\begin{aligned} \int_a^b |\alpha - f(x)| dx &= \int_a^b \alpha - f(x) dx \\ &= \int_a^b \alpha - \max(f(a), f(b)) + \max(f(a), f(b)) - f(x) dx \\ &= \int_a^b |\alpha - \max(f(a), f(b))| + |\max(f(a), f(b)) - f(x)| dx \\ &> \int_a^b |\max(f(a), f(b)) - f(x)| dx. \end{aligned}$$

Therefore $\alpha > \max(f(a), f(b))$ cannot be the minimizer since we can simply let $\alpha = \max(f(a), f(b))$ and we have a better minimizer. So we have that $\alpha \leq \max(f(a), f(b))$ and a similar argument gives us that $\alpha \geq \min(f(a), f(b))$. Now we have that $\alpha \in [\min_{x \in [a,b]} f(x), \max_{x \in [a,b]} f(x)]$. From this and the continuity of f there exists $r \in [a, b]$ such that $f(r) = \alpha$, specifically $\alpha = Lr + c$. We now assume that $L > 0$ as the other case ($L < 0$) is analogous.

Continuing with this r we get

$$\begin{aligned} \int_a^b |\alpha - f(x)| dx &= \int_a^b |Lr + c - (Lx + c)| dx \\ &= \int_a^b L|r - x| dx \\ &= L \int_a^r (r - x) dx + L \int_r^b (x - r) dx \\ &= L(r - a)^2/2 + L(b - r)^2/2 \\ &= L(r - a)^2/2 + L(r - b)^2/2 \\ &= \frac{L(\delta - A)^2}{2} + \frac{L(\delta + A)^2}{2} \quad \text{letting } \delta := r - \frac{a+b}{2} \text{ and } A := \frac{a-b}{2} \\ &= L(A^2 + \delta^2) \\ &= \frac{L(b-a)^2}{4} + L\delta^2. \end{aligned}$$

Upon noting that the last line is minimized for $\delta = 0$ we arrive at (34) and the first equality in (35). The second equality in (35) follows from noting

$$\int_a^b f(x) dx = \left[\frac{1}{2} Lx^2 + cx \right]_{x=a}^{x=b} = \frac{1}{2} L(b^2 - a^2) + c(b - a) = \left(L \frac{a+b}{2} + c \right) (b - a).$$

□

The following lemma addresses the point in the main text “[w]e show in the appendix that this decays at a rate of $O(b^{-1})$ and that *this rate is tight*.” We use it later to show a lower bound for the rate of convergence of the standard histogram.

Lemma C.10. Let $b \in \mathbb{N}$ and $f_L : [0, 1] \rightarrow \mathbb{R}$ be a collection of pdfs indexed by $L \in [0, \infty)$ via

$$f_L(x) = \begin{cases} 1 + (x - \frac{1}{2})L & L \leq 2 \\ (xL - L + \sqrt{2L})_+ & L \geq 2. \end{cases}$$

We have that

- (a) If $L \leq 2$ then $\min_{g \in \mathcal{H}_{1,b}} \|f_L - g\|_1 = \frac{L}{4b}$.
(b) If $L \geq 2$ and b is a multiple of $\sqrt{L/2}$, we have $\min_{g \in \mathcal{H}_{1,b}} \|f_L - g\|_1 = \frac{\sqrt{2L}}{4b}$.

Proof of Lemma C.10. We will begin by optimizing over $\text{span}(\mathcal{H}_{1,b})$ and will then show that the optimum lies in $\mathcal{H}_{1,b}$. To begin

$$\begin{aligned} \min_{g \in \text{span}(\mathcal{H}_{1,b})} \|f_L - g\|_1 &= \min_w \int_0^1 \left| f_L(x) - \sum_{i=1}^b w_i b \mathbb{1}((i-1)/b \leq x < i/b) \right| dx \\ &= \min_{\tilde{w}} \sum_{i=1}^b \int_{(i-1)/b}^{i/b} |f_L(x) - \tilde{w}_i| dx. \end{aligned} \quad (36)$$

For case (b) there is a breakpoint at x satisfying $xL - L + \sqrt{2L} = 0 \iff x = 1 - \sqrt{\frac{2}{L}}$. Since b is a multiple of $\sqrt{L/2}$, there exists an integer z such that

$$\begin{aligned} b &= z\sqrt{L/2} \\ \iff z/b &= \sqrt{2/L} \\ \iff (b-z)/b &= 1 - \sqrt{2/L} \end{aligned} \quad (37)$$

and we have for both cases (a) and (b) that f_L is linear on the bins $[(i-1)/b, i/b)$.

Applying Lemma C.9 we have the following for (a),

$$\min_{\tilde{w}} \sum_{i=1}^b \int_{(i-1)/b}^{i/b} |f_L(x) - \tilde{w}_i| dx = \sum_{i=1}^b \frac{L(i/b - (i-1)/b)^2}{4} = \frac{L}{4b}.$$

For (b) we split the summation between the breakpoint

$$\begin{aligned} &\min_{\tilde{w}} \sum_{i=1}^b \int_{(i-1)/b}^{i/b} |f_L(x) - \tilde{w}_i| dx \\ &= \min_{\tilde{w}} \sum_{i \in [b]: \frac{i}{b} \leq 1 - \sqrt{\frac{2}{L}}} \int_{(i-1)/b}^{i/b} |f_L(x) - \tilde{w}_i| dx + \sum_{i \in [b]: \frac{i}{b} > 1 - \sqrt{\frac{2}{L}}} \int_{(i-1)/b}^{i/b} |f_L(x) - \tilde{w}_i| dx \\ &= \sum_{i \in [b]: \frac{i}{b} \leq 1 - \sqrt{\frac{2}{L}}} 0 + \sum_{i \in [b]: \frac{i}{b} > 1 - \sqrt{\frac{2}{L}}} \frac{L}{4b^2} \\ &= \left| \left\{ i \in [b] : \frac{i}{b} > 1 - \sqrt{\frac{2}{L}} \right\} \right| \frac{L}{4b^2}. \end{aligned} \quad (38)$$

From (37) and because $1 - \sqrt{2/L} \geq 0$ it follows that $b(1 - \sqrt{2/L})$ is a nonnegative integer so

$$\begin{aligned} \left| \left\{ i \in [b] : \frac{i}{b} > 1 - \sqrt{\frac{2}{L}} \right\} \right| &= b - \left| \left\{ i \in [b] : \frac{i}{b} \leq 1 - \sqrt{\frac{2}{L}} \right\} \right| \\ &= b - \left| \left\{ i \in [b] : i \leq b \left(1 - \sqrt{\frac{2}{L}} \right) \right\} \right| \\ &= b - \left(b \left(1 - \sqrt{\frac{2}{L}} \right) \right) \\ &= b\sqrt{\frac{2}{L}}. \end{aligned}$$

Now (38) is equal to

$$b\sqrt{\frac{2}{L}} \frac{L}{4b^2} = \frac{\sqrt{2L}}{4b}.$$

We will show that the argument for the minimum w (with $bw = \tilde{w}$) from (36) lies in Δ_b to finish the proof. From the second equality in (35) we clearly get that $\tilde{w}_i \geq 0$ for all i . Again using the second equality in (35) we have

$$\begin{aligned} \sum_{i=1}^b w_i &= \frac{1}{b} \sum_{i=1}^b \tilde{w}_i = \frac{1}{b} \sum_{i=1}^b \frac{1}{i/b - (i-1)/b} \int_{(i-1)/b}^{i/b} f_L(x) dx \\ &= \sum_{i=1}^b \int_{(i-1)/b}^{i/b} f_L(x) dx \\ &= \int_0^1 f_L(x) dx = 1. \end{aligned}$$

□

C.2 Finite Sample Rate: Multi-view

This section contains the finite-sample bounds from Section 2.2.2 in the main text. The results here are a bit stronger than those from the main text at the cost of being a bit less concise. In the main text we weakened the results by assuming that Lipschitz constants were greater than or equal to 2. The next section contains corresponding results for Tucker models. First we will prove the following finite-sample bound.

Proposition C.4 (Proposition 2.1 in main text). *Let $d, b, k, n \in \mathbb{N}$ and $0 < \delta \leq 1$. There exists an estimator $V_n \in \mathcal{H}_{d,b}^k$ such that for all densities $p \in \mathcal{D}_d$ the following holds with probability at least $1 - \delta$*

$$\|p - V_n\|_1 \leq \min_{q \in \mathcal{H}_{d,b}^k} 3\|p - q\|_1 + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \quad (39)$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Proof of Proposition C.4. We begin by showing that (39), and therefore the proposition, holds if $n = 1$, and $d, b, k \geq 1$. Bounding the second term in the next summation with $b, d, k \geq 1$ and $n = 1$ yields the following inequality

$$\begin{aligned} 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} &\geq 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} \\ &\geq 7\sqrt{2 \log(4)} \\ &\geq 7. \end{aligned}$$

The triangle inequality gives us $\|p - V_n\|_1 \leq 2 \leq 7$. Therefore Proposition C.5 holds for $n = 1$. We will now proceed assuming that $n \geq 2$.

Let $1 \geq \varepsilon > 0$ be arbitrary. By Corollary A.1 there exists $p_1, \dots, p_M \in \mathcal{H}_{d,b}^k$ such that $M \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k}{\varepsilon}\right)^k$ and for all $q \in \mathcal{H}_{d,b}^k$ there exists $i \leq M$ with $\|p_i - q\|_1 \leq \varepsilon$. Applying Lemma A.6 with the same ε gives a deterministic algorithm V_n that, given at least $\frac{\log(3M^2/\delta)}{2\varepsilon^2}$ samples from a density p , outputs an index $j \in [M]$ where, with probability at least $1 - \delta$, the following holds

$$\begin{aligned} \|p_j - p\|_1 &\leq 3 \min_{i \in [M]} \|p_i - p\|_1 + 4\varepsilon \\ &\leq 3 \min_{q \in \mathcal{H}_{d,b}^k} \min_{i \in [M]} (\|p_i - q\|_1 + \|q - p\|_1) + 4\varepsilon \\ &= \min_{q \in \mathcal{H}_{d,b}^k} 3 \left(\min_{i \in [M]} \|p_i - q\|_1 \right) + 3\|q - p\|_1 + 4\varepsilon \\ &\leq 7\varepsilon + 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 \end{aligned}$$

(we are loosing $\delta/3$ from Lemma A.6 to δ for convenience).

Note that

$$\frac{\log(3M^2/\delta)}{2\varepsilon^2} = \frac{\log(M)}{\varepsilon^2} + \frac{\log(3/\delta)}{2\varepsilon^2}. \quad (40)$$

We will now bound $\log(M)$ which, because ε is positive and \log is strictly increasing, will give us an upper bound on the previous term. The following follows from the fact that b, d and k are all greater than or equal to 1 and $1 \geq \varepsilon > 0$,

$$\begin{aligned} M &\leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k}{\varepsilon}\right)^k \\ &\leq \left(\frac{4bdk}{\varepsilon}\right)^{bdk} \left(\frac{4bdk}{\varepsilon}\right)^{bdk} \\ &= \left(\frac{4bdk}{\varepsilon}\right)^{2bdk}. \end{aligned}$$

Applying this to (40) we have

$$\frac{\log(3M^2/\delta)}{2\varepsilon^2} \leq \frac{2bdk \log(\frac{4bdk}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2}. \quad (41)$$

The rest of the proof will be primarily concerned with choosing $\varepsilon \in (0, 1]$ so that the RHS of (41) is less than or equal to n so the hypotheses of Lemma A.6 are satisfied. We begin by eliminating some settings where selecting V_n trivial; we will then apply Lemma A.6 for the remaining settings.

Observe that if $n < 4bdk \log(4bdkn)$, then

$$7\sqrt{\frac{2bdk \log(4bdkn)}{n}} > 2,$$

and inequality (39) holds trivially. Similarly, if $n < \log(\frac{3}{\delta})$, then

$$7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \geq \frac{7}{\sqrt{2}} > 2,$$

and again inequality (39) holds trivially.

Thus, we can proceed with the setting

$$\begin{aligned} n &\geq 2 \max\left(2bdk \log(4bdkn), \frac{\log(\frac{3}{\delta})}{2}\right) \\ &\geq 2bdk \log(4bdkn) + \frac{\log(\frac{3}{\delta})}{2}. \end{aligned} \quad (42)$$

Defining the function $\rho(\varepsilon) := \frac{2bdk \log(\frac{4bdk}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2}$, inequality (42) implies that

$$\rho(1) = \frac{2bdk \log(\frac{4bdk}{1})}{1} + \frac{\log(\frac{3}{\delta})}{2} \leq n. \quad (43)$$

Furthermore,

$$\lim_{x \rightarrow 0} \rho(x) = \infty. \quad (44)$$

Together with the mean value theorem, (43) and (44) imply that we can now pick $1 \geq \varepsilon > 0$ such that

$$\rho(\varepsilon) = \frac{2bdk \log(\frac{4bdk}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2} = n. \quad (45)$$

We can now apply the estimator from Lemma A.6 to select the estimator V_n . As we have shown before the estimator in Lemma A.6 outputs a density in $\mathcal{H}_{d,b}^k$ such that

$$\|V_n - p\|_1 \leq 7\varepsilon + 3 \min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1. \quad (46)$$

By (45), we have

$$\varepsilon = \sqrt{\frac{2bdk \log(\frac{4bdk}{\varepsilon})}{n} + \frac{\log(\frac{3}{\delta})}{2n}} \geq \sqrt{\frac{1}{2n}} \geq \frac{1}{n},$$

since $0 < \delta, \varepsilon \leq 1$ and $n \geq 2$. Using this in (45), we obtain

$$\begin{aligned} \varepsilon &= \sqrt{\frac{2bdk \log(\frac{4bdk}{\varepsilon})}{n} + \frac{\log(\frac{3}{\delta})}{2n}} \\ &\leq \sqrt{\frac{2bdk \log(4bdkn)}{n} + \frac{\log(\frac{3}{\delta})}{2n}} \\ &\leq \sqrt{\frac{2bdk \log(4bdkn)}{n}} + \sqrt{\frac{\log(\frac{3}{\delta})}{2n}}. \end{aligned}$$

The result follows upon plugging this back into inequality (46). \square

Now we can prove the key result from the paper.

Proposition C.5 (Theorem 2.6 in main text). *Let $L \geq 2$, $0 < \delta \leq 1$ and $k, n \in \mathbb{N}$. Then there exists b and an estimator $V_n \in \mathcal{H}_{d,b}^k$ such that for any density $p \triangleq \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j}$ where $p_{i,j} \in \text{Lip}_L$ and w is in the probability simplex, the following holds with probability at least $1 - \delta$,*

$$\|V_n - p\|_1 \leq \frac{21dk^{1/3}L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \quad (47)$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

This also holds with “ $L \leq 2$ ” replacing “ $L \geq 2$ ” and the following inequality replacing (47)

$$\|V_n - p\|_1 \leq \sqrt{d} \frac{k^{1/3}}{n^{1/3}} \left[L^{1/3} \exp\left(\frac{L^2(d-1)}{24}\right) + 20\sqrt{\log(7dnk)} \right] + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}. \quad (48)$$

Proof of Proposition C.5. We begin with the $L \geq 2$ case and the other case will follow with some minor adjustments. From Hölder’s Inequality followed by Proposition C.1, if $b^2 \geq L^2/12$ and $L \geq 2$, then for any collection f_1, \dots, f_d in Lip_L we have that

$$\left\| \prod_{j=1}^d f_j - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{j=1}^d f_j \right\|_1^2 \leq \left\| \prod_{j=1}^d f_j - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{j=1}^d f_j \right\|_2^2 \leq \frac{dL^{\frac{d+3}{2}}}{12b^2} \left[\frac{\sqrt{8}}{3} \right]^{d-1} \leq \frac{dL^{\frac{d+3}{2}}}{9b^2}.$$

Taking the square root, we have

$$\left\| \prod_{j=1}^d f_j - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{j=1}^d f_j \right\|_1 \leq \frac{\sqrt{d}L^{\frac{d+3}{4}}}{3b}.$$

Consider some density p like that from the theorem statement

$$p = \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j}.$$

Since $\sum_{i=1}^k w_i \prod_{j=1}^d \text{Proj}_{\mathcal{H}_{1,b}} p_{i,j}$ is an element of $\mathcal{H}_{d,b}^k$, using Corollary C.1 gives us

$$\begin{aligned}
\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 &\leq \left\| \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j} - \sum_{i=1}^k w_i \prod_{j=1}^d \text{Proj}_{\mathcal{H}_{1,b}} p_{i,j} \right\|_1 \\
&= \left\| \sum_{i=1}^k w_i \prod_{j=1}^d p_{i,j} - \sum_{i=1}^k w_i \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{j=1}^d p_{i,j} \right\|_1 \\
&\leq \sum_{i=1}^k w_i \left\| \prod_{j=1}^d p_{i,j} - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{j=1}^d p_{i,j} \right\|_1 \\
&\leq \sum_{i=1}^k w_i \frac{\sqrt{d} L^{\frac{d+3}{4}}}{3b} = \frac{\sqrt{d} L^{\frac{d+3}{4}}}{3b}.
\end{aligned}$$

Invoking the estimator V_n from Proposition C.4 and using the previous bound it follows that, for any choice of k and b such that $b^2 \geq L^2/12$, there exists an estimator $V_n \in \mathcal{H}_{d,b}^k$ where for all densities p from our theorem statement, with probability at least $1 - \delta$, the following holds

$$\|p - V_n\|_1 \leq \frac{\sqrt{d} L^{\frac{d+3}{4}}}{b} + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}. \quad (49)$$

We now set $b = \left\lceil n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3} \right\rceil$. Note that if $n < kL$ then (47) is trivially satisfied since $L \geq 2$ and

$$\frac{21dk^{1/3} L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \geq \frac{21dk^{1/3} L^{\frac{1+3}{12}}}{n^{\frac{1}{3}}} \geq \frac{21dk^{1/3} L^{\frac{1}{3}}}{n^{\frac{1}{3}}} \geq 21.$$

Thus we can proceed with the assumption that $n \geq kL$ for the $L \geq 2$ case of this proof. Under this assumption, noting again that $L \geq 2$, we have the following bound on b ,

$$b \geq n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3} \geq L^{1/3} L^{\frac{1+3}{6}} = L,$$

which implies $b^2 \geq \frac{L^2}{12}$. So for the remainder of the $L \geq 2$ case of this proof we will proceed with the assumption $b^2 \geq \frac{L^2}{12}$ and use the estimator from (49).

Since $n \geq kL \geq k$ and $L \geq 2$, we also have $1 \leq n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3} \leq b \leq 2n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3}$. Thus

$$\begin{aligned}
7\sqrt{\frac{2bdk \log(4bdkn)}{n}} &\leq 7\sqrt{\frac{4dk^{2/3} n^{\frac{1}{3}} L^{\frac{d+3}{6}} \log\left(8L^{\frac{d+3}{6}} dk^{2/3} n^{4/3}\right)}{n}} \\
&= 14 \frac{\sqrt{d} k^{1/3} L^{\frac{d+3}{12}}}{n^{1/3}} \sqrt{\log\left(8L^{\frac{d+3}{6}} dk^{2/3} n^{4/3}\right)} \\
&\leq 14 \frac{\sqrt{d} k^{1/3} L^{\frac{d+3}{12}}}{n^{1/3}} \sqrt{\log(3^{2d} L^{2d} d^{2d} k^{2d} n^{2d})} \\
&= 14 \frac{\sqrt{d} k^{1/3} L^{\frac{d+3}{12}}}{n^{1/3}} \sqrt{2d \log(3Ldkn)} \\
&\leq 20 \frac{dk^{1/3} L^{\frac{d+3}{12}}}{n^{1/3}} \sqrt{\log(3Ldkn)}.
\end{aligned}$$

Using this with (49) and applying $n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3} \leq b$ to the first summand in (49) we have

$$\begin{aligned}
\|p - V_n\|_1 &\leq \frac{\sqrt{d} k^{1/3} L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} + \frac{20dk^{1/3} L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \\
&\leq \frac{21dk^{1/3} L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}},
\end{aligned}$$

as expected. This finishes the $L \geq 2$ case.

For the $L \leq 2$ case, we obtain (using Proposition C.1 and taking the square root along with some simple manipulations) instead,

$$\min_{q \in \mathcal{H}_{d,b}^k} \|p - q\|_1 \leq \sqrt{d} \frac{L}{3b} \exp\left(\frac{L^2(d-1)}{24}\right),$$

and instead of (49),

$$\|p - V_n\|_1 \leq \sqrt{d} \frac{L}{b} \exp\left(\frac{L^2(d-1)}{24}\right) + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}. \quad (50)$$

Since $L \leq 2$, we clearly have $\frac{L^2}{12} < b^2$ so we can use Proposition C.1 without issue. Set $b = \lceil n^{1/3} k^{-1/3} L^{2/3} \rceil$ (we can assume $L > 0$ and thus $b \geq 1$ since the $L = 0$ case can be solved by simply setting V_n to output the uniform distribution). If $k > n$ then (48) is trivially satisfied so we can proceed with the assumption that $k \leq n$.

Since $n \geq k$ and $n^{1/3} k^{-1/3} \geq 1$ it follows that $b \geq n^{1/3} k^{-1/3} L^{2/3}$ and $b \leq \lceil n^{1/3} k^{-1/3} 2^{2/3} \rceil \leq \lceil n^{1/3} k^{-1/3} 2 \rceil \leq 3n^{1/3} k^{-1/3}$. Letting $C = 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}$ we can bound (50) as follows

$$\begin{aligned} \|p - V_n\|_1 &\leq \sqrt{d} \frac{L}{b} \exp\left(\frac{L^2(d-1)}{24}\right) + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + C \\ &\leq \sqrt{d} \frac{L}{b} \exp\left(\frac{L^2(d-1)}{24}\right) + 7\sqrt{\frac{2(3n^{1/3} k^{-1/3})k}{n}} \sqrt{d \log(4(3n^{1/3} k^{-1/3})dkn)} + C \\ &\leq \sqrt{d} \frac{L^{1/3} k^{1/3}}{n^{1/3}} \exp\left(\frac{L^2(d-1)}{24}\right) + \frac{7k^{1/3}}{n^{1/3}} \sqrt{6d \log(12dn^{4/3} k^{2/3})} + C \\ &\leq \sqrt{d} \frac{L^{1/3} k^{1/3}}{n^{1/3}} \exp\left(\frac{L^2(d-1)}{24}\right) + \frac{7k^{1/3}}{n^{1/3}} \sqrt{6d \log(7^{4/3} d^{4/3} n^{4/3} k^{4/3})} + C \\ &\leq \sqrt{d} \frac{L^{1/3} k^{1/3}}{n^{1/3}} \exp\left(\frac{L^2(d-1)}{24}\right) + \frac{7k^{1/3}}{n^{1/3}} \sqrt{8d \log(7dnk)} + C \\ &\leq \sqrt{d} \frac{L^{1/3} k^{1/3}}{n^{1/3}} \exp\left(\frac{L^2(d-1)}{24}\right) + \frac{20k^{1/3} \sqrt{d}}{n^{1/3}} \sqrt{\log(7dnk)} + C \\ &= \sqrt{d} \frac{k^{1/3}}{n^{1/3}} \left[L^{1/3} \exp\left(\frac{L^2(d-1)}{24}\right) + 20\sqrt{\log(7dnk)} \right] + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}, \end{aligned}$$

which finishes the proof for the $L \leq 2$ case. \square

C.3 Finite Sample Rate: Tucker

Here we present results for the Tucker decomposition that are analogous to those in the last section. The results here were omitted from the main text “[f]or brevity, and because the results are virtually direct analogues of their multi-view histogram counterparts...” We begin again with a proof of a finite-sample bound, which will then be used to prove a distribution-free bound.

Proposition C.6. *Let $d, b, k, n \in \mathbb{N}$ and $0 < \delta \leq 1$. There exists an estimator $V_n \in \tilde{\mathcal{H}}_{d,b}^k$ such that for all densities $p \in \mathcal{D}_d$, the following holds with probability at least $1 - \delta$*

$$\|p - V_n\|_1 \leq 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \quad (51)$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

Proof of Proposition C.6. This proof is very similar to the proof of Proposition C.4 and we will provide fewer intermediate steps when they are virtually identical to those that proof.

Similarly to the proof of Proposition C.4, we can assume $n \geq 2$. Let $1 \geq \varepsilon > 0$ be arbitrary. By Corollary A.2 there exists $p_1, \dots, p_M \in \tilde{\mathcal{H}}_{d,b}^k$ such that $M \leq \left(\frac{4bd}{\varepsilon}\right)^{bdk} \left(\frac{4k^d}{\varepsilon}\right)^{k^d}$ and for all $q \in \tilde{\mathcal{H}}_{d,b}^k$ there exists $i \leq M$ with $\|p_i - q\|_1 \leq \varepsilon$.

Now, by applying Lemma A.6 with the same ε , there exists a deterministic algorithm which, for all densities p , can output an index $j \in [M]$ such that

$$\|p_j - p\|_1 \leq 7\varepsilon + 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1,$$

with probability at least $1 - \delta$ given at least N samples from the distribution, where

$$N \geq \frac{\log(3M^2/\delta)}{2\varepsilon^2} = \frac{\log(M)}{\varepsilon^2} + \frac{\log(3/\delta)}{2\varepsilon^2}.$$

Now we can bound

$$\frac{\log(3M^2/\delta)}{2\varepsilon^2} \leq \frac{2bdk \log(\frac{4bd}{\varepsilon})}{\varepsilon^2} + \frac{2k^d \log(\frac{4k^d}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2}.$$

Note that:

- If $n \leq 6bdk \log(4bdkn)$, then $7\sqrt{\frac{2bdk \log(4bdkn)}{n}} \geq 7/\sqrt{3} \geq 2$, making (51) trivial.
- If $n \leq 6k^d \log(4k^d n)$ then $7\sqrt{\frac{2k^d \log(4k^d n)}{n}} \geq 7/\sqrt{3} \geq 2$, also making (51) trivial.
- Similarly, if $n \leq 3 \log(\frac{3}{\delta})$, then $7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \geq 2$, making (51) trivial yet again.

Thus we can assume that

$$\begin{aligned} n &\geq 3 \max \left(2bdk \log(4bdkn), 2k^d \log(4k^d n), \log \left(\frac{3}{\delta} \right) \right) \\ &\geq 2bdk \log(4bdkn) + 2k^d \log(4k^d n) + \log \left(\frac{3}{\delta} \right). \end{aligned} \quad (52)$$

We now define $\rho(\varepsilon)$ as

$$\rho(\varepsilon) := \frac{2bdk \log(\frac{4bd}{\varepsilon})}{\varepsilon^2} + \frac{2k^d \log(\frac{4k^d}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2}. \quad (53)$$

So (52) gives us $\rho(1) \leq n$ and, as before, $\lim_{x \rightarrow 0} \rho(x) = \infty$. Together with the mean value theorem it follows we can now pick $1 \geq \varepsilon > 0$ such that

$$\rho(\varepsilon) = \frac{2bdk \log(\frac{4bd}{\varepsilon})}{\varepsilon^2} + \frac{2k^d \log(\frac{4k^d}{\varepsilon})}{\varepsilon^2} + \frac{\log(\frac{3}{\delta})}{2\varepsilon^2} = n. \quad (54)$$

We can now apply the estimator from Lemma A.6 to select the estimator V_n . As we have shown before the estimator in Lemma A.6 outputs a density in $\mathcal{H}_{d,b}^k$ such that

$$\|V_n - p\|_1 \leq 7\varepsilon + 3 \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1. \quad (55)$$

Now, note that by (54), we have

$$\varepsilon = \sqrt{\frac{2bdk \log(\frac{4bd}{\varepsilon})}{n} + \frac{2k^d \log(\frac{4k^d}{\varepsilon})}{n} + \frac{\log(\frac{3}{\delta})}{2n}} \geq \sqrt{\frac{1}{2n}} \geq \frac{1}{n},$$

since $0 < \delta, \varepsilon \leq 1$ and $n \geq 2$. Plugging this back into (54), we obtain

$$\begin{aligned} \varepsilon &\leq \sqrt{\frac{2bdk \log(4bdkn)}{n} + \frac{2k^d \log(4k^d n)}{n} + \frac{\log(\frac{3}{\delta})}{2n}} \\ &\leq \sqrt{\frac{2bdk \log(4bdkn)}{n}} + \sqrt{\frac{2k^d \log(4k^d n)}{n}} + \sqrt{\frac{\log(\frac{3}{\delta})}{2n}}. \end{aligned}$$

The result follows upon plugging this back into inequality (55). \square

Now we can prove our distribution-free bound.

Proposition C.7. *Let $L \geq 2$, $0 < \delta \leq 1$ and $k, n \in \mathbb{N}$. Then there exists $b \in \mathbb{N}$ and an estimator $V_n \in \tilde{\mathcal{H}}_{d,b}^k$ such that for any density $p \triangleq \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i}$ where $p_{i,j} \in \text{Lip}_L$ and W is a probability tensor, the following holds with probability at least $1 - \delta$,*

$$\|p - V_n\|_1 \leq \frac{21dk^{1/3}L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}} \quad (56)$$

where V_n is a function of $X_1, \dots, X_n \stackrel{iid}{\sim} p$.

This also holds with “ $L \leq 2$ ” replacing “ $L \geq 2$ ” and the following inequality replacing (56)

$$\|p - V_n\|_1 \leq \sqrt{d} \frac{k^{1/3}}{n^{1/3}} \left[L^{1/3} \exp\left(\frac{L^2(d-1)}{24}\right) + 20\sqrt{\log(7dnk)} \right] + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}.$$

Proof of Proposition C.7. This proof is very similar to the proof of Proposition C.5 and we will provide fewer intermediate steps when they are virtually identical to those that proof. We begin with the $L \geq 2$ case. Consider some density p like that from the theorem statement

$$p = \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i}.$$

Since $\sum_{S \in [k]^d} W_S \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} p_{i,S_i}$ is an element of $\tilde{\mathcal{H}}_{d,b}^k$, using Corollary C.1, Hölder’s Inequality, and Proposition C.1 yields the following when $b^2 \geq L^2/12$

$$\begin{aligned} \min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 &\leq \left\| \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} - \sum_{S \in [k]^d} W_S \prod_{i=1}^d \text{Proj}_{\mathcal{H}_{1,b}} p_{i,S_i} \right\|_1 \\ &= \left\| \sum_{S \in [k]^d} W_S \prod_{i=1}^d p_{i,S_i} - \sum_{S \in [k]^d} W_S \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d p_{i,S_i} \right\|_1 \\ &\leq \sum_{S \in [k]^d} W_S \left\| \prod_{i=1}^d p_{i,S_i} - \text{Proj}_{\mathcal{H}_{d,b}^1} \prod_{i=1}^d p_{i,S_i} \right\|_1 \\ &\leq \sum_{S \in [k]^d} W_S \frac{\sqrt{d}L^{\frac{d+3}{4}}}{3b} = \frac{\sqrt{d}L^{\frac{d+3}{4}}}{3b}. \end{aligned} \quad (57)$$

Combining this with the estimator from Proposition C.6 we get that, for any b such that $b^2 \geq L^2/12$, we have

$$\|p - V_n\|_1 \leq \frac{\sqrt{d}L^{\frac{d+3}{4}}}{b} + 7\sqrt{\frac{2bdk \log(4bdkn)}{n}} + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}.$$

If $n < kL$, the RHS of (56) is greater than 2, which means that (56) holds trivially. Thus we assume $n \geq kL$. Since b doesn’t appear in the third summand in the previous inequality, and the first, second, and fourth summand are exactly the same as those from (49) in the proof Proposition C.5, we can set $b = \left\lceil n^{\frac{1}{3}} L^{\frac{d+3}{6}} k^{-1/3} \right\rceil$ and proceed exactly as we did in that proof (again $b^2 \geq L^2/12$ so Proposition C.1 holds). We then obtain

$$\|p - V_n\|_1 \leq \frac{21dk^{1/3}L^{\frac{d+3}{12}}}{n^{\frac{1}{3}}} \sqrt{\log(3Ldkn)} + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}},$$

as expected.

For the $L \leq 2$ case, (57) becomes

$$\min_{q \in \tilde{\mathcal{H}}_{d,b}^k} \|p - q\|_1 \leq \sqrt{d} \frac{L}{3b} \exp\left(\frac{L^2(d-1)}{24}\right),$$

from which we get

$$\|p - V_n\|_1 \leq \sqrt{d} \frac{L}{b} \exp\left(\frac{L^2(d-1)}{24}\right) + 7\sqrt{\frac{2bdk \log(4bdk)}{n}} + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}}.$$

Since b doesn't appear in the third summand in the previous inequality, and the rest of the inequality is exactly the same as in the proof of Proposition C.5, we can again let $b = \lceil n^{1/3} k^{-1/3} L^{2/3} \rceil$, assume $n \geq k$ and, proceed identically as we did in the proof of Proposition C.5. This yields

$$\|p - V_n\|_1 \leq \sqrt{d} \frac{k^{1/3}}{n^{1/3}} \left[L^{1/3} \exp\left(\frac{L^2(d-1)}{24}\right) + 20\sqrt{\log(7dnk)} \right] + 7\sqrt{\frac{2k^d \log(4k^d n)}{n}} + 7\sqrt{\frac{\log(\frac{3}{\delta})}{2n}},$$

as expected. \square

C.4 Lower Bound: Standard Histogram

Proof of Proposition 2.3. Let $p \in \mathcal{D}_d$ with $p = \prod_{i=1}^d p_i$ where p_1 is the density from Lemma C.10 for $L = 2$, i.e. $p_1(x) = 2x$, and $p_i(x) \equiv 1$ for $i > 1$. Let $(Y_1, \dots, Y_d) \sim V_n$ and $(X_1, \dots, X_d) \sim p$. Because total variation distance is never increased through mappings of the random variables (see Theorem 5.2 in [10]) we have that $\|V_n - p\|_1 \geq \|f - p_1\|_1$ where f is the probability density associated with Y_1 . We will now show that f is an element of $\mathcal{H}_{1,b}$. Let $S \subset [0, 1]$ be an arbitrary (Borel) measurable set and note that V_n has the form $\sum_{A \in [b]^d} \hat{w}_A h_{d,b,A}$. Then we have that

$$\begin{aligned}
P(Y \in S) &= \int_S f d\lambda \\
&= \int_{S \times [0,1] \times \dots \times [0,1]} V_n d\lambda \\
&= \int_{S \times [0,1] \times \dots \times [0,1]} \sum_{A \in [b]^d} \hat{w}_A h_{d,b,A} d\lambda \\
&= \sum_{A \in [b]^d} \hat{w}_A \int_{S \times [0,1] \times \dots \times [0,1]} h_{d,b,A} d\lambda \\
&= \sum_{A \in [b]^d} \hat{w}_A \int_{S \times [0,1] \times \dots \times [0,1]} \prod_{i=1}^d h_{1,b,A_i} d\lambda \\
&= \sum_{A \in [b]^d} \hat{w}_A \left(\int_S h_{1,b,A_1} d\lambda \right) \left(\int_{[0,1]} h_{1,b,A_2} d\lambda \right) \dots \left(\int_{[0,1]} h_{1,b,A_d} d\lambda \right) \\
&= \sum_{A \in [b]^d} \hat{w}_A \int_S h_{1,b,A_1} d\lambda \\
&= \int_S \sum_{A \in [b]^d} \hat{w}_A h_{1,b,A_1} d\lambda.
\end{aligned}$$

note that $\sum_{A \in [b]^d} \hat{w}_A h_{1,b,A_1}$ is a histogram and thus the density associated with f is a histogram. Using this fact with the earlier mentioned inequality we have that

$$\begin{aligned}
\|V_n - p\|_1 &\geq \|f - p_1\|_1 \\
&\geq \min_{h \in \mathcal{H}_{1,b}} \|h - p_1\|_1.
\end{aligned}$$

From Lemma C.10 it then follows that

$$\|V_n - p\|_1 \geq \frac{1}{2b}.$$

Let $D > 0$ be arbitrary. From our assumption that $n/b^d \rightarrow \infty$ it follows that, for sufficiently large n , that $n/b^d \geq (2D)^d$ and furthermore

$$n/b^d \geq (2D)^d \iff \sqrt[d]{n}/b \geq 2D \iff 1/(2b) \geq D/\sqrt[d]{n} \Rightarrow \|V_n - p\|_1 \geq D/\sqrt[d]{n}$$

so $\|V_n - p\|_1 \in \omega(1/\sqrt[d]{n})$ by definition. \square

D Experimental Setting

Consider the problem of finding some density estimator \hat{p} with minimal L_2 distance to an unknown density p (p is the various projections of MNIST and Diabetes from the main text). This is equivalent to minimizing the squared L^2 loss:

$$\begin{aligned} & \int_{[0,1]^d} (p(x) - \hat{p}(x))^2 dx \\ &= \int_{[0,1]^d} \hat{p}(x)^2 dx - 2 \int_{[0,1]^d} p(y)\hat{p}(y)dy + \int_{[0,1]^d} p(z)^2 dz. \end{aligned} \quad (58)$$

Because the right term in (58) does not depend on \hat{p} it can be ignored when finding optimal \hat{p} . The left term in (58) is known. The middle term in (58) can be estimated with the following approximation

$$\int_{[0,1]^d} p(x)\hat{p}(x)dx = \mathbb{E}_{X \sim p} [\hat{p}(X)] \approx \frac{1}{n} \sum_{i=1}^n \hat{p}(X_i)$$

where $\mathcal{X} = X_1, \dots, X_n \stackrel{iid}{\sim} p$. We can use this to find a good estimate $\hat{H} \in \mathcal{R}_{d,b}^k$ for p which represents $\mathcal{H}_{d,b}^k$ or $\tilde{\mathcal{H}}_{d,b}^k$:

$$\begin{aligned} & \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \int_{[0,1]^d} (\hat{H}(x) - \hat{p}(x))^2 dx = \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \langle \hat{H}, \hat{H} \rangle - 2 \int_{[0,1]^d} \hat{H}(x)p(x)dx \\ & \approx \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \langle \hat{H}, \hat{H} \rangle - 2 \frac{1}{n} \sum_{i=1}^n \hat{H}(X_i). \end{aligned} \quad (59)$$

Recall that the standard histogram estimator is $H = H_{d,b}(\mathcal{X}) = \frac{1}{n} \sum_{i=1}^n \sum_{A \in [b]^d} h_{d,b,A} \mathbb{1}(X_i \in \Lambda_{d,b,A})$ and let $\hat{H} = \sum_{A \in [b]^d} \hat{w}_A h_{d,b,A}$. We have the following

$$\begin{aligned} \langle \hat{H}, H \rangle &= \left\langle \sum_{A \in [b]^d} \hat{w}_A h_{d,b,A}, \frac{1}{n} \sum_{i=1}^n \sum_{B \in [b]^d} h_{d,b,B} \mathbb{1}(X_i \in \Lambda_{d,b,B}) \right\rangle \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{A \in [b]^d} \hat{w}_A \mathbb{1}(X_i \in \Lambda_B) b^d = \frac{1}{n} \sum_{i=1}^n \hat{H}(X_i). \end{aligned}$$

As a consequence (59) is equal to

$$\begin{aligned} \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \langle \hat{H}, \hat{H} \rangle - 2 \langle H, \hat{H} \rangle &= \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \langle \hat{H}, \hat{H} \rangle - 2 \langle H, \hat{H} \rangle + \langle H, H \rangle \\ &= \arg \min_{\hat{H} \in \mathcal{R}_{d,b}^k} \|H - \hat{H}\|_2^2. \end{aligned}$$

Using the $U_{d,b}$ operator we can reformulate this into a tensor factorization problem

$$\min_{\hat{T} \in \mathcal{Q}_{d,b}^k} \|H - U_{d,b}(\hat{T})\|_2^2 = \min_{\hat{T} \in \mathcal{Q}_{d,b}^k} b^d \|U_{d,b}^{-1}(H) - \hat{T}\|_2^2$$

where $\mathcal{Q}_{d,b}^k$ could be either $\mathcal{T}_{d,b}^k$ or $\tilde{\mathcal{T}}_{d,b}^k$. Because of this equivalence, to find estimates in $\mathcal{H}_{d,b}^k$ or $\tilde{\mathcal{H}}_{d,b}^k$ we can simply use nonnegative tensor decomposition algorithms, which minimize ℓ^2 loss, to find NNTF histograms that approximate H .

E Nonexistence of Infinite Tensor Decomposition

Let $p : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a probability density function. Consider the possibility of decomposing p as follows

$$p(x, y) = \sum_{i=1}^{\infty} w_i f_i(x) g_i(y) \quad (60)$$

where, for all i , $w_i \geq 0$ and f_i and g_i are probability densities. We are going to show that this is not always possible, which we will do by contradiction. Let λ be the Lebesgue measure (dimensionality will be left implicit). We are going to use the following proposition which we will prove later.

Proposition E.1. *There exists a set $E \subset [0, 1] \times [0, 1]$ such that $\lambda(E) > 0$ and for all non-null measurable sets $A, B \subset [0, 1]$ we have that $\lambda(E \cap A \times B) < \lambda(A \times B)$.*

Let E be a set satisfying the property in Proposition E.1. Let $\mathbb{1}_E$ be the indicator function of a set E . We will let $p = \mathbb{1}_E$ and assume that p has a decomposition as in (60).

We will assume that $w_1 > 0$. Clearly we have that $p - w_1 f_1 g_1$ is an almost everywhere (a.e.) nonnegative function (all products of functions here are outer products). Let $\varepsilon > 0$ such that $A \triangleq f_1^{-1}([\varepsilon, \infty))$ and $B \triangleq g_1^{-1}([\varepsilon, \infty))$ have positive measure. Such an ε must exist otherwise f_1 and g_1 are 0 a.e.. Now we have that $\varepsilon^2 \mathbb{1}_A \mathbb{1}_B \leq f_1 g_1$. And thus $p - w_1 \varepsilon^2 \mathbb{1}_A \mathbb{1}_B \geq 0$ a.e. or equivalently $\lambda(E)^{-1} \mathbb{1}_E - w_1 \varepsilon^2 \mathbb{1}_{A \times B} \geq 0$ a.e.. From our definition of E we know that $\lambda(A \times B \setminus E) = \lambda(A \times B) - \lambda(E \cap A \times B) > 0$ so $\lambda(E)^{-1} \mathbb{1}_E - w_1 \varepsilon^2 \mathbb{1}_{A \times B}$ is negative on a set of positive measure, a contradiction.

We will now address the existence of the set E . The most direct statement of the existence of such an E can be found in [23], the following is the exact statement from the text.

Theorem E.1 ([23] Theorem 2.1). *There exist Borel measurable subsets $E \subset [0, 1]^2$ of positive measure which are rectangle free, so that if $A \times B \subseteq E$ then $\text{area}(A \times B) = 0$.*

That paper builds the set E via a random construction and contains an image which showing an example that approximates a randomly sampled E . Their construction seems to imply that the condition “ $A \times B \subseteq E$ ” was intended to be interpreted measure theoretically, i.e. “ $\text{area}(A \times B \setminus E) = 0$ ”; it is not particularly difficult to construct a measurable subset of $[0, 1] \times [0, 1]$ which contains all but a null set of $[0, 1] \times [0, 1]$ and is “rectangle free” as described in the theorem statement (see [14] and references therein). If the measure theoretic strengthening is true it would imply the existence of the set E from Proposition E.1 above. Since we are not *totally* certain that this strengthening is possible we include a proof of the existence of E above.

For a topological space (Ω, τ) equipped with a Borel measure μ , a set $S \subseteq \Omega$ is called *essentially dense* if, for any nonempty open set I , $\mu(I \cap S) > 0$. For any measurable set in \mathbb{R}^d we will equip it with the standard subspace topology and measure induced by the standard Lebesgue measure. There exists a measurable set $D \subset \mathbb{R}$ such that D and D^C are essentially dense (see [15] 134J(a)). The following is a simplification of Theorem 1 in [14] that we will use to construct E .

Theorem E.2 ([14] Theorem 1). *For a measurable set of the form $E = \{(x, y) : x - y \in D\}$ the following two conditions are equivalent*

1. D is essentially dense on \mathbb{R} .
2. $\lambda(E \cap A \times B) > 0$ for all A, B such that $\lambda(A)\lambda(B) > 0$.

From this we have that $E \triangleq \{(x, y) : x - y \in D\} \cap [0, 1]^2$ and $E^C = \{(x, y) : x - y \in D^C\} \cap [0, 1]^2$ (we let E live in $[0, 1]^2$) have that property that for non null sets $A, B \subset [0, 1]$, $\lambda(E \cap A \times B) > 0$, $\lambda(E^C \cap A \times B) > 0$. Note that $\lambda(E \cap A \times B) + \lambda(E^C \cap A \times B) = \lambda(A \times B)$ and thus $\lambda(E \cap A \times B) < \lambda(A \times B)$ so we have constructed E .

We mention that the E we have constructed contains a non-null rectangle when rotated by 45 degrees. Thusly rotating p to give \tilde{p} allows us to find f, g and $w > 0$ such that $\tilde{p} - wfg$ is a.e. nonnegative. In [14] the authors discuss the existence of sets E where, for all non null A, B , we have that $\lambda(f(E) \cap A \times B) > 0$, for all f in certain classes of transforms. These results hint towards research directions of finding transforms so that our data is better approximated by nice NNTF model.

References

- [1] Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Ann. Statist.*, 37(6A):3099–3132, 12 2009.
- [2] Magda Amiridi, Nikos Kargas, and Nicholas D. Sidiropoulos. Low-rank Characteristic Tensor Density Estimation Part I: Foundations. *arXiv e-prints*, page arXiv:2008.12315, August 2020.
- [3] Magda Amiridi, Nikos Kargas, and Nicholas D. Sidiropoulos. Low-rank Characteristic Tensor Density Estimation Part II: Compression and Latent Density Estimation. *arXiv e-prints*, page arXiv:2106.10591, June 2021.
- [4] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.
- [5] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization – provably. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pages 145–162, New York, NY, USA, 2012. ACM.
- [6] Hassan Ashtiani, Shai Ben-David, Nicholas Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Nearly tight sample complexity bounds for learning mixtures of gaussians via sample compression schemes. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 3412–3421. Curran Associates, Inc., 2018.
- [7] H.H. Bauschke and P.L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics, Ouvrages de mathématiques de la SMC. Springer New York, 2011.
- [8] Emmanuel J Candes and Terence Tao. Reflections on compressed sensing. *IEEE Information Theory Society Newsletter*, 58(4):20–23, 2008.
- [9] Efrén Cruz Cortés and Clayton Scott. Consistent Kernel Density Estimation with Non-Vanishing Bandwidth. *arXiv e-prints*, page arXiv:1705.08921, May 2017.
- [10] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation*. Springer, New York, 2001.
- [11] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, April 2006.
- [12] David Donoho and Victoria Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 1141–1148. MIT Press, 2004.
- [13] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning*, ICML ’08, page 272–279, New York, NY, USA, 2008. Association for Computing Machinery.
- [14] P. Erdős and John C. Oxtoby. Partitions of the plane into sets having positive measure in every non-null measurable product set. *Transactions of the American Mathematical Society*, 79(1):91–102, 1955.
- [15] D.H. Fremlin. *Measure Theory*. Number v. 1 in Measure theory. Torres Fremlin, 2000.
- [16] László Györfi, Luc Devroye, and Laszlo Györfi. *Nonparametric density estimation: the L1 view*. John Wiley & Sons, New York; Chichester, 1985.
- [17] Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of discrete distributions under ℓ_1 loss. *CoRR*, abs/1411.1467, 2014.

- [18] Nils-Bastian Heidenreich, Anja Schindler, and Stefan Sperlich. Bandwidth selection for kernel density estimation: a review of fully automatic selectors. *AStA Advances in Statistical Analysis*, 97(4):403–433, Oct 2013.
- [19] Heinrich Jiang. Uniform convergence rates for kernel density estimation. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1694–1703, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [20] Nikos Kargas and Nicholas D. Sidiropoulos. Learning mixtures of smooth product distributions: Identifiability and algorithm. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 388–396. PMLR, 16–18 Apr 2019.
- [21] Hiroyuki Kasahara and Katsumi Shimotsu. Non-parametric identification and estimation of the number of components in multivariate mixtures. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):97–111, 2014.
- [22] Bernhard Kawohl. *Rearrangements*, pages 7–99. Springer Berlin Heidelberg, Berlin, Heidelberg, 1985.
- [23] Wilfrid S. Kendall and Giovanni Montana. Small sets and markov transition densities. *Stochastic Processes and their Applications*, 99(2):177 – 194, 2002.
- [24] J. Kim and C. Scott. Robust kernel density estimation. *J. Machine Learning Res.*, 13:2529–2565, 2012.
- [25] Jisu Kim, Jaehyeok Shin, Alessandro Rinaldo, and Larry Wasserman. Uniform Convergence Rate of the Kernel Density Estimator Adaptive to Intrinsic Volume Dimension. *arXiv e-prints*, page arXiv:1810.05935, Oct 2018.
- [26] Yong-Deok Kim and Seungjin Choi. Nonnegative tucker decomposition. *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [27] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. TensorLy: Tensor Learning in Python. *arXiv e-prints*, page arXiv:1610.09555, Oct 2016.
- [28] Han Liu, Min Xu, Haijie Gu, Anupam Gupta, John Lafferty, and Larry Wasserman. Forest density estimation. *J. Mach. Learn. Res.*, 12:907–951, July 2011.
- [29] Y. P. Mack and M. Rosenblatt. Multivariate k-nearest neighbor density estimates. *Journal of Multivariate Analysis*, 9(1):1–15, March 1979.
- [30] Sahand Negahban and Martin J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Ann. Statist.*, 39(2):1069–1097, 04 2011.
- [31] Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *J. Mach. Learn. Res.*, 13:1665–1697, May 2012.
- [32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [33] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, August 2010.
- [34] R.D. Reiss. *Approximate distributions of order statistics: with applications to nonparametric statistics*. Springer series in statistics. Springer, 1989.

- [35] Alexander Ritchie, Robert A. Vandermeulen, and Clayton Scott. Consistent estimation of identifiable nonparametric mixture models from grouped observations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11676–11686. Curran Associates, Inc., 2020.
- [36] Aaron Schein, Mingyuan Zhou, David Blei, and Hanna Wallach. Bayesian poisson tucker decomposition for learning the structure of international relations. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2810–2819, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [37] Eugene F. Schuster. Incorporating support constraints into nonparametric estimators of densities. *Communications in Statistics - Theory and Methods*, 14(5):1123–1136, 1985.
- [38] Amnon Shashua and Tamir Hazan. Non-negative tensor factorization with applications to statistics and computer vision. In *Proceedings of the 22Nd International Conference on Machine Learning*, ICML ’05, pages 792–799, New York, NY, USA, 2005. ACM.
- [39] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London, 1986.
- [40] Le Song, Animashree Anandkumar, Bo Dai, and Bo Xie. Nonparametric estimation of multi-view latent variable models. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 640–648, Beijing, China, 22–24 Jun 2014. PMLR.
- [41] Le Song and Bo Dai. Robust low rank kernel embeddings of multivariate distributions. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [42] George R. Terrell and David W. Scott. Variable kernel density estimation. *Ann. Statist.*, 20(3):1236–1265, 09 1992.
- [43] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008.
- [44] Robert Vandermeulen and Clayton Scott. Consistency of robust kernel density estimators. In Shai Shalev-Shwartz and Ingo Steinwart, editors, *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 568–591, Princeton, NJ, USA, 12–14 Jun 2013. PMLR.
- [45] Robert A. Vandermeulen. Improving Nonparametric Density Estimation with Tensor Decompositions. *arXiv e-prints*, page arXiv:2010.02425, October 2020.
- [46] Robert A. Vandermeulen, René Saitenmacher, and Alexander Ritchie. A proposal for supervised density estimation. In *NeurIPS 2020 Workshop on Pre-registration in Machine Learning*, 2020.
- [47] Robert A. Vandermeulen and Clayton Scott. Robust kernel density estimation by scaling and projection in hilbert space. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems* 27, pages 433–441. Curran Associates, Inc., 2014.
- [48] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Number 47 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- [49] Yannis G. Yatracos. Rates of Convergence of Minimum Distance Estimators and Kolmogorov’s Entropy. *The Annals of Statistics*, 13(2):768 – 774, 1985.