

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/370370659>

Sample Complexity Bounds for Learning High-dimensional Simplices in Noisy Regimes

Conference Paper · April 2023

CITATIONS

0

READS

29

4 authors:



Seyed amir hossein Saberi

Sharif University of Technology

4 PUBLICATIONS 1 CITATION

[SEE PROFILE](#)



Amir Najafi

Institute for Research in Fundamental Sciences (IPM)

21 PUBLICATIONS 663 CITATIONS

[SEE PROFILE](#)



Abolfazl Motahari

Sharif University of Technology

56 PUBLICATIONS 442 CITATIONS

[SEE PROFILE](#)



Babak Khalaj

Sharif University of Technology

253 PUBLICATIONS 2,505 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



5G wireless networks [View project](#)



Cell Identity Codes: Understanding Cell Identity from Gene Expression Profiles using Deep Neural Networks [View project](#)

Sample Complexity Bounds for Learning High-dimensional Simplices in Noisy Regimes

Amir Hossein Saberi¹ Amir Najafi² Seyed Abolfazl Motahari³ Babak H. Khalaj¹

Abstract

In this paper, we find a sample complexity bound for learning a simplex from noisy samples. Assume a dataset of size n is given which includes i.i.d. samples drawn from a uniform distribution over an unknown simplex in \mathbb{R}^K , where samples are assumed to be corrupted by a multi-variate additive Gaussian noise of an arbitrary magnitude. We prove the existence of an algorithm that with high probability outputs a simplex having a ℓ_2 distance of at most ε from the true simplex (for any $\varepsilon > 0$). Also, we theoretically show that in order to achieve this bound, it is sufficient to have $n \geq (K^2/\varepsilon^2) e^{\Omega(K/\text{SNR}^2)}$ samples, where SNR stands for the signal-to-noise ratio. This result solves an important open problem and shows as long as $\text{SNR} \geq \Omega(K^{1/2})$, the sample complexity of the noisy regime has the same order to that of the noiseless case. Our proofs are a combination of the so-called sample compression technique in (Ashtiani et al., 2018), mathematical tools from high-dimensional geometry, and Fourier analysis. In particular, we have proposed a general Fourier-based technique for recovery of a more general class of distribution families from additive Gaussian noise, which can be further used in a variety of other related problems.

1. Introduction

Many practical problems in machine learning and data science can be naturally modeled by learning of high-

dimensional geometric shapes from a given set of unlabeled data points. In particular, learning of simplices from randomly scattered points arises in many fields ranging from bioinformatics to remote sensing (Chan et al., 2009; Schwartz & Shackney, 2010; Satas & Raphael, 2017). A simplex is defined as the set of all convex combinations of $K + 1$ points in \mathbb{R}^K , for $K \in \mathbb{N}$. Formally speaking, our problem can be stated as follows: assume n i.i.d. samples are drawn from a uniform measure over an unknown simplex in \mathbb{R}^K . Also, each sample is assumed to be corrupted by an additive multi-variate white Gaussian noise with a covariance of $\sigma^2 \mathbf{I}_{K \times K}$, for an unknown σ . In this regard, the main question that we try to address in this paper is: How large n needs to be in terms of parameters such as K , σ , SNR and etc., such that the true simplex can be consistently approximated with an arbitrarily high probability?

Learning of simplices is a well-studied task. Due to its practical importance, many heuristic algorithms are proposed to solve this problem in different scenarios such as noisy and sparse cases. In this paper, we focus on the theoretical aspects of this research area which still contains several open problems. There exists at least one efficient algorithm for learning of simplices that comes with a theoretical guarantee (see Section 1.1), but still suffers from a very large sample complexity, e.g., it requires $n \geq \Omega(K^{22})$. In a concurrent line of work, researchers work on deriving information-theoretic bounds on the sample complexity of this problem irrespective of their time and memory complexity. Recently, optimal information-theoretic sample complexity bounds are derived for the noiseless case where Maximum Likelihood Estimator (MLE) is used. The output of MLE for this problem is the minimum volume simplex which contains all the samples which runs in exponential time¹. Najafi et al. (2021) proved that MLE is a PAC-learning algorithm for finding a K -simplex with an asymptotically decreasing Total Variation (TV) distance from the true one. They showed that a maximum of $\tilde{\Omega}(K^2/\varepsilon)$ (Logarithmic terms are ignored) samples are sufficient to estimate the true simplex up to a TV distance of $\varepsilon > 0$. However, the noiseless case is unrealistic since we do not have access to clean data in real-

¹Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran ²School of Mathematics, Institute for Research in Fundamental Sciences (IPM), P.O. Box: 19395-5746, Tehran, Iran. ³Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. Correspondence to: Amir Hossein Saberi <saberi.sah@ee.sharif.edu>, Amir Najafi <najafi@ipm.ir>, Seyed Abolfazl Motahari <motahari@sharif.edu>, Babak H. Khalaj <khalaj@sharif.edu>.

¹Finding a polynomial algorithm with provable guarantees that comes with a reasonable dependency on K is still an open problem.

world situations. On the other hand, the minimum-volume inclusive simplex which is proposed by [Najafi et al. \(2021\)](#) is no longer a valid solution when samples are noisy, since the corrupted samples are not necessarily contained inside the true simplex.

In this work, we aim at finding the information-theoretic sample complexity of learning simplices in noisy regimes. Mathematically speaking, we assume samples are generated according to the following equation:

$$\mathbf{y}_i = \mathbf{V}\phi_i + \mathbf{z}_i, \quad i = 1, \dots, n, \quad (1)$$

where \mathbf{V} is a $K \times (K+1)$ matrix that includes the vertices of the true simplex \mathcal{S}_T as its columns. Each $\phi_i \in \mathbb{R}^K$ is drawn from a uniform Dirichlet distribution and $\mathbf{z}_i \in \mathbb{R}^K$ are sampled from a multivariate Gaussian distribution $\mathcal{N}(0, \sigma^2 \mathbf{I}_K)$. We tackle the task of estimating the true simplex via an exponential-time estimator $\hat{\mathcal{S}} = \hat{\mathcal{S}}(\mathbf{y}_{1:n})$. In particular, we prove that having at most $\tilde{\Omega}(K^2/\varepsilon^2) e^{\Omega(K/\text{SNR}^2)}$ noisy samples from a K -simplex is sufficient to estimate the simplex with a ℓ_2 -distance of at most ε , for any $\varepsilon > 0$. Here, SNR denotes the ratio of the signal, i.e., maximum standard deviation of the uniform measure over \mathcal{S}_T , to the standard deviation of the noise per component, i.e., σ . This result solves one of the main open problems in this area and could justify an interesting previously observed phenomenon in practice: Almost all heuristic algorithms perform close the noiseless case as long as SNR is larger than a presumably sub-linear function of dimension K . However, they experience a sudden decline in performance as SNR is decreased further from that point.

Our approach to prove the above-mentioned bound has four main steps. 1) We find a ball in \mathbb{R}^K which with high probability contains the true simplex. 2) This ball is then quantized through covering with enough isolated points. Each $K+1$ combination of such points forms a candidate simplex, and consequently a candidate density to approximate the true simplex. As long as the covering is accomplished with a sufficiently high precision, we prove the existence of at least one candidate that has a small TV-distance from the main simplex. 3) In the next step, we choose a noise-corrupted density, i.e., the convolution of a uniform measure over a candidate simplex with Gaussian distribution, from the candidate set such that, with high probability, it has a small TV distance from that of the true simplex. This can be done using some aspects of the so-called ‘‘sample compression’’ technique. 4) In the final step, we propose a novel Fourier-based technique to show that as long as the noise-smoothed versions of any two simplices are close to each other in TV-distance, their underlying simplices are also close to each other in the sense of ℓ_2 -distance.

1.1. Related Works

We categorize the existing works based on their focus, which could be either the efficiency of the proposed algorithms, or the fundamental sample complexity of the problem irrespective of its time complexity. In the former, the main purpose is mostly to provide a heuristic solution with promising results in practice. However, papers in the latter category aim to find fundamental and information-theoretic limitations for the problem. In the remainder of this part, we review a number of works from both categories.

[Anderson et al. \(2013\)](#) proved that by having $\Omega(K^{22})$ noiseless samples, one can estimate the true simplex via an polynomial-time algorithm. The core idea is to utilize the third moment and local search techniques from Independent Component Analysis (ICA) research. To the best of our knowledge, finding polynomial-time algorithms with a more relaxed dependence to K is still an open problem. Concurrently, [Najafi et al. \(2021\)](#) proved that sample complexity of the MLE in noiseless case is $\tilde{\Omega}(K^2/\varepsilon)$, where ε is the permissible TV-distance of the output of the algorithm from the true simplex. Since MLE in this case runs in exponential time, they also provide an alternative heuristic approach as a practical surrogate to MLE, which still does not come with a rigid theoretical guarantee.

Theoretical attempts to tackle this problem in the noisy setting is limited to the work of [Bhattacharyya & Kannan \(2020\)](#). Their work is also based on the sample compression technique, originally used in the seminal work of [Ashtiani et al. \(2018\)](#) in learning high-dimensional Gaussian mixture models. [Bhattacharyya & Kannan \(2020\)](#) prove that a noisy simplex can be learned using $\Omega(K^2)$ samples only if there exists at least one sample near each vertex of the simplex. In fact, this is a very strong assumption, and a back of the envelope calculation shows that one needs around $\tilde{\Omega}((1/\varepsilon)^K)$ samples to guarantee the occurrence of this event.

From a practical point of view, several heuristic methods have been introduced so far in order to deal with real-world problems in bioinformatics, hyper-spectral imaging and etc ([Piper et al., 2004](#); [Bioucas-Dias et al., 2012](#); [Lin et al., 2013](#)). In hyper-spectral imaging, one aims at finding the distribution of the constituent elements in an area by examining remote hyper-spectral images. Each pixel in such an image can be thought as a random convex combination of a finite set of fixed prototypes which correspond to pure elements that exist in that region. Therefore, the problem of finding constituent minerals would natural translate into estimating an unknown simplex from a set of presumably uniform samples ([Ambikapathi et al., 2011](#); [Agathos et al., 2014](#); [Zhang et al., 2017](#)). Learning of simplices in bioinformatics usually arises in the study of complex tissues. A complex tissue is composed of multiple cell-types—a group of cells with similar characteristics—such as blood, brain,

or even tumor cells (Tolliver et al., 2010; Zuckerman et al., 2013). Bulk data from complex tissues, such as gene expression level vectors, can again be modeled by convex combinations of its constituent cell-types. In this line of work, researchers aim at finding the structure of tissues through studying the state of cell-types, which again leads into learning a high-dimensional simplex (Shoval et al., 2012; Korem et al., 2015).

The rest of the paper is organized as follows: In Section 2, we formally define the problem and present our notation and definitions. Our main theoretical results as well as the algorithm that achieves our bounds are discussed in Section 3. Finally, Section 4 concludes the paper and presents some suggestions for future works.

2. Notation and Definitions

We use the same notations as (Najafi et al., 2021). A K -simplex is defined as the set of all convex combination of $K + 1$ affinely independent points in \mathbb{R}^K . Let $\mathbf{V} = [\mathbf{v}_0 | \mathbf{v}_1 | \dots | \mathbf{v}_K] \in \mathbb{R}^{K \times (K+1)}$ be a matrix whose columns represent vertices of the simplex, then K -simplex \mathcal{S} can be defined as

$$\mathcal{S} = \left\{ \mathbf{V}\phi \mid \phi \in \mathbb{R}^{K+1}, \phi \succ \mathbf{0}, \phi^T \mathbf{1} = 1 \right\}.$$

Also, let us denote the set of all possible K -simplices in \mathbb{R}^K by \mathbb{S}_K . We denote the uniform probability measure over a simplex \mathcal{S} by $\mathbb{P}_{\mathcal{S}}$, and its probability density function by $f_{\mathcal{S}}(x)$. Thus, $f_{\mathcal{S}}(x)$ can be written as

$$f_{\mathcal{S}}(x) = \frac{\mathbf{1}(x \in \mathcal{S})}{\text{Vol}(\mathcal{S})},$$

where $\text{Vol}(\mathcal{S})$ denotes the Lebesgue measure (or volume) of \mathcal{S} .

The noisy simplex family, i.e., the class of distributions formed by *convolving* uniform probability density function over simplices in \mathbb{S}_K by $G_{\sigma} \triangleq \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, is denoted by $\mathbb{G}_{K,\sigma}$. Mathematically speaking,

$$\mathbb{G}_{K,\sigma} \triangleq \{f_{\mathcal{S}} * G_{\sigma} \mid \mathcal{S} \in \mathbb{S}_K\}, \quad (2)$$

where $*$ denotes the convolution operator. Obviously, the distribution of input data points \mathbf{y}_i lie in $\mathbb{G}_{K,\sigma}$. With a little abuse of notation, we use the term “class of simplices” to both refer to \mathbb{S}_K and also $\{f_{\mathcal{S}} \mid \mathcal{S} \in \mathbb{S}_K\}$ whenever the difference is clear from the context. In a similar fashion, we refer to $\mathbb{G}_{K,\sigma}$ as the class of “noisy simplices”.

In order to measure the difference between two distributions, we use both ℓ_2 and total variation distance. Consider two probability measures \mathbb{P}_1 and \mathbb{P}_2 , with respective density functions f_1 and f_2 , which are defined over \mathbb{R}^K . Then, TV

distance between \mathbb{P}_1 and \mathbb{P}_2 can be defined as

$$\text{TV}(\mathbb{P}_1, \mathbb{P}_2) \triangleq \sup_{A \in \mathcal{B}} |\mathbb{P}_1(A) - \mathbb{P}_2(A)| = \frac{1}{2} \|\mathbf{f}_1 - \mathbf{f}_2\|_1,$$

where \mathcal{B} is the set of all Borel sets in \mathbb{R}^K .

Definition 2.1 (PAC-Learnability of a distribution class in realizable setting). We say a class of distributions \mathcal{F} is PAC learnable in realizable setting, if there exists a learning method which for any distribution $g \in \mathcal{F}$ and any $\epsilon, \delta > 0$, outputs an estimator \hat{g} using $n \geq \text{poly}(1/\epsilon, 1/\delta)$ i.i.d. samples from g , which with probability at least $1 - \delta$ satisfies

$$\|\hat{g} - g\|_{\text{TV}} \leq \epsilon. \quad (3)$$

We also need to define a series of geometric restrictions for the simplex in noisy cases. In fact, those simplices that are significantly stretched toward on particular direction can be shown to be more prone to noise than those with some minimum levels of geometric regularity. We discuss the necessity of such definitions in later stages. Similar to Najafi et al. (2021), let us denote the volume of the largest facet of a K -simplex (here, volume needs be calculated in \mathbb{R}^{K-1}) by $\mathcal{A}_{\max}(\mathcal{S})$, and the length of the largest line segment inside the simplex (diameter) by $\mathcal{L}_{\max}(\mathcal{S})$. In this regard, we define the isoperimetricity of a K -simplex as follow:

Definition 2.2 ($(\underline{\theta}, \bar{\theta})$ -isoperimetricity of simplices). A K -simplex $\mathcal{S} \in \mathbb{S}_K$ is defined to be $(\underline{\theta}, \bar{\theta})$ -isoperimetric if the following inequalities hold:

$$\begin{aligned} \mathcal{A}_{\max}(\mathcal{S}) &\leq \bar{\theta} \text{Vol}(\mathcal{S})^{\frac{K-1}{K}}, \\ \mathcal{L}_{\max}(\mathcal{S}) &\leq \underline{\theta} K \text{Vol}(\mathcal{S})^{\frac{1}{K}}. \end{aligned}$$

The overall concept of isoperimetricity for simplices reflects the fact that for a simplex to be (even partially) recoverable from noisy data, it should not be stretched too much in any direction or having highly acute angles. In other words, sample complexity in noisy regimes is also affected by the geometric shape of the underlying simplex.

Definition 2.3 (ϵ -representative set). For any $\epsilon > 0$, we say that a finite set of distributions \mathcal{G} is an ϵ -representative set for a distribution class \mathcal{F} , if for any distribution $f \in \mathcal{F}$, there exists at least one $g \in \mathcal{G}$ that satisfies

$$\|f - g\|_{\text{TV}} \leq \epsilon.$$

Throughout the paper, we use bold lowercase letters to show scalars, bold letters to show vectors and bold uppercase letters to show matrices. We also use light uppercase letters to show random variables and light lower case letters to show realization of a random variable. In this paper and for the sake of simplicity in notation, for any fixed $\bar{\theta}, \underline{\theta} > 0$, whenever we say the class of simplices we actually mean the class of $(\underline{\theta}, \bar{\theta})$ -isoperimetric simplices.

2.1. Problem Definition

In the first three parts of the paper, we aim at proving the PAC-learnability of noisy simplices family $\mathbb{G}_{K,\sigma}$. The final part is devoted to showing that if two members in $\mathbb{G}_{K,\sigma}$ are close to each other, e.g., the true noisy simplex and its consistent estimator, then the simplices are also consistently close.

Based on definition 2.1, and in order to show that the class of noisy K -simplices is PAC-learnable, one should find an algorithm that for all $f_S * G_\sigma \in \mathbb{G}_{K,\sigma}$, any positive $\varepsilon, \delta > 0$, and given a dataset $\mathbf{D} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ of i.i.d. samples drawn from $f_S * G_\sigma \in \mathbb{G}_{K,\sigma}$, as long as $n \geq \text{poly}(1/\varepsilon, 1/\delta)$, outputs a noisy simplex $f_{\hat{S}} * G_\sigma$ such that

$$\mathbb{P}[\|(f_{\hat{S}} - f_S) * G_\sigma\|_{TV} \geq \varepsilon] \leq \delta. \quad (4)$$

This is due to the fact that we assume samples to be drawn from \mathbb{P}_S and then are corrupted by an additive independent zero mean Gaussian noise with a covariance matrix of $\sigma^2 \mathbf{I}$. Our ultimate goal in the first three parts of the paper is the following: To derive explicit polynomial forms for the lower-bound $n \geq \text{poly}(1/\varepsilon, \log(1/\delta))$, which (as we show in Theorem 3.5) turns out to be

$$n \geq \tilde{\Omega}\left[\frac{K^2}{\varepsilon^2} \log \frac{1}{\delta}\right].$$

The final part is dedicated to show that as long as (with high probability) $(f_{\hat{S}} - f_S) * G_\sigma \leq \varepsilon$, then we also have

$$\|f_{\hat{S}} - f_S\|_2 \leq \varepsilon e^{\Omega(\frac{K}{\text{SNR}^2})},$$

where $\text{SNR} \triangleq (\mathcal{L}_{\max}(\mathcal{S})/K)/\sigma$ denotes the signal-to-noise ratio. It can be easily checked that based on the definition of a uniform Dirichlet distribution, \mathcal{L}_{\max}/K represents the maximum component-wise standard deviation associated to f_S . This will prove our claims in Section 1.

3. Statistical Learning of Noisy Simplices

Before going through the details, let us first present an sketch of proof for PAC-learnability of noisy simplices.

(*Bounding the candidate set*): In order to estimate the true simplex from noisy i.i.d. samples, we first split our data in half and use the first half to restrict the set of all K -simplices in \mathbb{S}_K to a bounded set $\mathbb{S}_K^{\mathcal{D}}$ which consists of all the simplices that happen to entirely fall within a finite ball. This way, we can eliminate very far candidates and thus focus on simplices that are placed near the data samples. In this regard, we construct a bounded version of \mathbb{S}_K , denoted by $\mathbb{S}_K^{\mathcal{D}}$, and prove that it includes (with high probability) the true underlying simplex.

(*Quantization*): We quantize this bounded set and create a finite ϵ -representative set of K -simplices denoted by

$\hat{\mathbb{S}}_K^{\mathcal{D}} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_M\}$ for $M \in \mathbb{N}$, such that for each simplex $\mathcal{S} \in \mathbb{S}_K^{\mathcal{D}}$, there exists some $i \in \{1, 2, \dots, M\}$ where $\|\mathbb{P}_{\mathcal{S}_i} - \mathbb{P}_{\mathcal{S}}\|_{TV} \leq \epsilon$.

(*Density selection*): In this part, we use the second half of data and try to choose the best simplex in $\hat{\mathbb{S}}_K^{\mathcal{D}}$. By the best simplex, we mean the one with the minimum TV-distance from the true simplex. We show that as long as the above-mentioned sample complexity bound is satisfied, the output of this selection procedure fall within a ε TV-distance of $f_{S_T} * G_\sigma$.

(*Denoising*): Finally, we show that estimating a noisy version of a simplex leads to consistent estimation for the simplex as well. This completes our proof.

3.1. Bounding The Candidate Set

In this part, we show how the first half of the dataset can be utilized in order to bound the set of all candidate simplices into a ball with a finite radius in \mathbb{R}^K . This procedure is crucial for later stages of the proof. In this regard, first let us review the generation process of noisy samples \mathbf{y}_i for $i \in [n]$:

$$\begin{aligned} \mathbf{y} &\sim f_{S_T} * G_\sigma \implies \mathbf{y} = \mathbf{x} + \mathbf{z}, \\ \mathbf{x} &= \mathbf{V}_S \phi, \quad \phi \sim \text{Dir}(1, 1, \dots, 1), \\ \mathbf{z} &\sim \mathcal{N}(\mathbf{0}, \sigma \mathbf{I}), \end{aligned} \quad (5)$$

where \mathbf{V}_S denotes the vertex matrix for \mathcal{S} .

The following lemma shows that having enough samples from a noisy simplex $f_S * G_\sigma$, one can find a hyper-sphere in \mathbb{R}^K which (with high probability) contains \mathcal{S} .

Lemma 3.1 (Creating $\mathbb{S}_K^{\mathcal{D}}$). *Suppose that we have a set of i.i.d. samples $\mathbf{D} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{2m}\}$ from $f_S * G_\sigma$ for $m \in \mathbb{N}$. If $m \geq 1000(K+1)(K+2) \log \frac{6}{\delta}$, then the true simplex \mathcal{S} with probability at least $1 - \delta$ is confined in a K -dimensional sphere with center point \mathbf{p} and radius R , where R and \mathbf{p} are defined as follows:*

$$R = 8\sqrt{(K+1)(K+2)D}, \quad \mathbf{p} = \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{y}_i, \quad (6)$$

and the variance of the noise can be upper-bounded as $\sigma^2 \leq \frac{D}{K-2} = R_n$, where D is defined as

$$D = \frac{1}{2m} \sum_{i=1}^m \|\mathbf{y}_{2i} - \mathbf{y}_{2i-1}\|_2^2.$$

Proof can be found in Section B of the Appendix. This way, the set $\mathbb{S}_K^{\mathcal{D}}$ can be fixed. Next, we discuss how to properly quantize this set in order to choose an appropriate final candidate for the true simplex, denoted by $\hat{\mathcal{S}}$.

3.2. Quantization

Let us call the hyper-sphere of Lemma 3.1 as $C^K(\mathbf{p}, R)$, which with high probability contains the true simplex. In this part, we choose a set of points $T_\epsilon^l(C^K(\mathbf{p}, R)) = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_l\}$ inside this sphere such that for any point $\mathbf{x} \in C^K(\mathbf{p}, R)$, there exists some $i \in \{1, 2, \dots, l\}$ such that $\|\mathbf{x} - \mathbf{p}_i\|_2 \leq \epsilon$. We call $T_\epsilon^l(C^K(\mathbf{p}, R))$ a covering set for $C^K(\mathbf{p}, R)$. One way to construct such covering set is to uniformly draw points from the sphere. We show that if the number of drawn random points exceeds $O\left(\left(1 + \frac{2R}{\epsilon}\right)^{2K}\right)$, then with high probability, they form an ϵ -covering set for $C^K(\mathbf{p}, R)$.

Suppose that we construct a random covering set as described above and denote it with $T_\epsilon(C^K(\mathbf{p}, R))$. Now, using each $K + 1$ distinct points in $T_\epsilon(C^K(\mathbf{p}, R))$ we can build a K -simplex. Assume we collect all such simplices in a set called $\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$, i.e.,

$$\widehat{\mathbb{S}}(C^K(\mathbf{p}, R)) = \left\{ \mathcal{S}(\mathbf{x}_1, \dots, \mathbf{x}_{K+1}) \mid \mathbf{x}_i \in T_\epsilon(C^K(\mathbf{p}, R)), i \in [K + 1] \right\}.$$

Obviously, there exist at most $\binom{|T_\epsilon(C^K(\mathbf{p}, R))|}{K+1}$ simplices in $\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$. The following lemma states that this set is a sufficiently good representative for all $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices inside $C^K(\mathbf{p}, R)$.

Lemma 3.2. *For any $\epsilon \in (0, 1)$, denote the set of all possible K -simplices with the vertices in $T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R))$ as $\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$. Then, the resulting set is an ϵ -representative set for all $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices in $C^K(\mathbf{p}, R)$, as long as we have:*

$$\alpha = \frac{\text{Vol}(\mathcal{S})^{1/K}}{5\bar{\theta}}.$$

The proof can be found in Section B of the Appendix. This way, the finite candidate set $\mathbb{S}_K^{\mathcal{D}}$ can be formed and we can jump to the next stage of our algorithm for finding a “good” candidate for the true simplex.

3.3. Density Selection

We take advantage of the fundamental result in Devroye & Lugosi (2012) which plays a central role in the remainder of our derivations in this subsection. At this stage, we have already created a finite set of representative simplices $\mathbb{S}_K^{\mathcal{D}}$ and set out to find the “best” simplex in this set. This can be done using the following theorem:

Theorem 3.3 (Theorem 6.3 of (Devroye & Lugosi, 2012)). *Let \mathcal{F} be a finite class of distributions consisting of M distinct members $\{f_1, f_2, \dots, f_M\}$. Also, suppose we have*

$n \geq \frac{\log(3M^2/\delta)}{2\epsilon^2}$ i.i.d. samples from an arbitrary distribution g for some $\epsilon, \delta > 0$. Then, there exist a deterministic algorithm \mathcal{A} , which outputs a number $j \in \{1, \dots, M\}$ satisfying the following inequality with probability at least $1 - \delta$:

$$\|f_j - g\|_{TV} \leq 3 \cdot \min_{i \in \{1, 2, \dots, M\}} \|f_i - g\|_{TV} + 4\epsilon. \quad (7)$$

Proof can be found inside the reference. Combining the results of Theorem 3.3 and Lemmas 3.1 and 3.2, we can present one of our main results as follows: The set of all shape-restricted simplices in \mathbb{R}^K which entirely fall inside a sphere with radius R is PAC-learnable:

Theorem 3.4 (PAC-Learnability of Simplices in $C^K(\mathbf{p}, R)$ Corrupted by Gaussain Noise with Bounded Variance). *Assume $\mathbf{p} \in \mathbb{R}^K$, and $R, R_n > 0$. Then, the class of $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices contained in the K -dimensional hyper-sphere $C^K(\mathbf{p}, R)$ and convolved with an isotropic Gaussian noise $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, with $\sigma \leq R_n$, is PAC-learnable. Specifically, for some $\epsilon, \delta > 0$, assume we have at least n i.i.d. samples from a distribution $f_S * G_\sigma$ with $\sigma \leq R_n$ and $S \in C^K(\mathbf{p}, R)$, where the the following bound is satisfied:*

$$n \geq 50 \frac{\log \frac{30R_n\sqrt{K}}{\delta\epsilon} + 2(K+1)^2 \log \left(1 + \frac{100R\bar{\theta}(K+1)}{\epsilon \text{Vol}(\mathcal{S})^{\frac{1}{K}}} \right)}{\epsilon^2}. \quad (8)$$

Then, there exists an algorithm \mathcal{A} whose outputs $\mathcal{S}_{\mathcal{A}}$ and $\sigma_{\mathcal{A}}$ satisfy the following inequality with probability at least $1 - \delta$:

$$\|f_S * G_\sigma - f_{\mathcal{S}_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}}\|_{TV} \leq \epsilon. \quad (9)$$

The proof can be found in Section A of the Appendix. Let $\mathbb{G}_{K,\sigma}(\mathbf{p}, R)$ represent the distribution set

$$\{f_S * G_\sigma \mid S \subseteq C^K(\mathbf{p}, R)\}.$$

In an agnostic setting, so far we have actually proved that for any $\epsilon, R, R_n > 0$, $\mathbf{p} \in \mathbb{R}^K$, and having $n \geq \tilde{\Omega}(K^2 \log R/\epsilon^2)$ samples from “any” distribution in

$$\bigcup_{\sigma \leq R_n} \mathbb{G}_{K,\sigma}(\mathbf{p}, R),$$

there exists an algorithm \mathcal{A} that outputs a simplex $\mathcal{S}_{\mathcal{A}}$ and noise standard deviation $\sigma_{\mathcal{A}}$ such that with a positive probability we have

$$\begin{aligned} & \|f_{\mathcal{S}_{\mathcal{A}}} * G_{K,\sigma_{\mathcal{A}}} - f_S * G_\sigma\|_{TV} \\ & \leq 4 \min_{\substack{S^* \subseteq C^K(\mathbf{p}, R) \\ \sigma^* \leq R_n}} \|f_{S^*} * G_{\sigma^*} - f_S * G_\sigma\|_{TV} + \epsilon, \end{aligned} \quad (10)$$

while the first term in the r.h.s. becomes zero in the realizable case. Also, note that $\bar{\theta}$ which represents the regularity level in the shape of simplices is also present inside the

bounds. In fact, trying to learn a highly stretched simplex from noisy data can become very tricky since even a small amount of noise can shoot almost all the samples outside of the simplex.

Another limitation of Theorem 9 is that it requires the simplices to be inside an sphere of radius R . Using Lemma 3.1, we show that this is not an actual necessity. The following theorem completes our first main result in this paper, which is the complete PAC-learnability of the class of noisy simplices:

Theorem 3.5 (PAC-Learnability of Noisy Simplices in General). *The class of $(\underline{\theta}, \theta)$ -isoperimetric K -simplices which are corrupted with Gaussian white noise, i.e.,*

$$\bigcup_{\sigma > 0} \mathbb{G}_{K, \sigma}$$

*is PAC-learnable. In other words, assume $n \geq \tilde{O}(K^2/\varepsilon^2)$ i.i.d. samples from a noisy simplex $f_S * G_\sigma$, for any $S \in \mathbb{S}_K$ and $\sigma > 0$, are given. Then, there exists an algorithm \mathcal{A} which outputs a noisy simplex $f_{S_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}}$ which with high probability satisfies*

$$\|f_{S_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}} - f_S * G_\sigma\|_{\text{TV}} \leq \varepsilon. \quad (11)$$

Proof of the above theorem can be found in Section A of the Appendix. Obviously, the $\log R$ dependence inside the sample complexity of Theorem 9 has been disappeared in that of Theorem 3.5.

3.4. Denoising

So far, we have shown that for a simplex $S \in \mathbb{S}_K$ with some levels of geometric regularity, one can learn $f_S * G_\sigma$ up to an arbitrarily small error as long as n satisfies the bound in Theorem 3.4. What remains to prove is that a consistent estimation of $f_S * G_\sigma$ leads to the same type of estimate for f_S as well. Our approach is based on showing the following three important properties:

- The difference function between any two distinct and geometrically regular members of $\{f_S | S \in \mathbb{S}_K\}$ corresponds to a *low-frequency* function, where “low-frequency” refers to the Fourier domain. In other words, it means that the majority of the energy of the difference function lies in those regions of the Fourier domain which are close to the origin.
- Being corrupted by an additive noise is equivalent to convolving each f_S with G_σ . Moreover, convolution transforms into point-wise multiplication in the Fourier domain.
- Finally, convolution with a Gaussian kernel G_σ preserves the low-frequency parts of the difference func-

tion. Therefore, the two simplices remain distinguishable even after corruption via additive Gaussian noise.

Mathematically speaking, assume $S_1, S_2 \in \mathbb{S}_K$ are distinct and have a minimum degree of geometric regularity, e.g., $\mathcal{L}_{\max}(S_i)$ is bounded for $i = 1, 2$. Our aim is to show that if $f_{S_1} * G_\sigma$ and $f_{S_2} * G_\sigma$ have a total variation distance of at least $\epsilon > 0$, then the ℓ_2 -distance between f_{S_1} and f_{S_2} is also bounded away from zero according to a function of ϵ, σ and the geometric regularity of simplices S_1 and S_2 . The theoretical core behind our method is stated in the following general theorem.

Theorem 3.6 (Recovery of Low-Frequency Objects from Additive Noise). *For $K \in \mathbb{N}$, consider a probability density function family \mathcal{F} which is supported over \mathbb{R}^K . Assume for sufficiently large $\alpha > 0$, the following bound holds for all $f, g \in \mathcal{F}$:*

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} |\mathcal{F}\{f - g\}(\omega)|^2 \leq \zeta(\alpha^{-1}) \int_{\mathbb{R}^K} |f - g|^2,$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier transform and ζ is an increasing function with $\zeta(0) = 0$ and continuity at 0. Also, assume the probability density function Q (also supported over \mathbb{R}^K) has the following property:

$$\inf_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}\{Q\}(\omega)| \geq \eta(\alpha),$$

where $\eta(\cdot)$ is a non-negative decreasing function. Then, there exists a non-negative constant C where for any $\sigma, \varepsilon > 0$ and $f, g \in \mathcal{F}$ with $\|f - g\|_2 \geq \varepsilon$, we have

$$\|(f - g) * Q\|_2 \geq \frac{\varepsilon}{(2\pi)^{K/2}} \left(\sup_{\alpha \geq 0} \eta(\alpha) \sqrt{1 - \zeta(\alpha^{-1})} \right),$$

Proof. For the sake of simplicity in notations, let $\mathcal{F}, \mathcal{G}, \mathcal{Q} : \mathbb{R}^K \rightarrow \mathbb{C}$ denote the Fourier transforms of f, g and Q , respectively. Due to Parseval’s theorem, we have

$$\|f - g\|_2^2 = \frac{1}{(2\pi)^K} \|\mathcal{F} - \mathcal{G}\|_2^2.$$

Also, due to the properties of the Fourier transform, which is the transformation of convolution into direct multiplication, one can write

$$\mathcal{F}\{(f - g) * Q\} = \mathcal{Q}(\mathcal{F} - \mathcal{G}).$$

Thus, for sufficiently large α we have:

$$\begin{aligned} (2\pi)^K \|(f - g) * Q\|_2^2 &= \int_{\mathbb{R}^K} |\mathcal{Q}(\omega) (\mathcal{F}(\omega) - \mathcal{G}(\omega))|^2 \\ &\geq \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{Q}(\omega)|^2 |\mathcal{F}(\omega) - \mathcal{G}(\omega)|^2 \\ &\geq \eta^2(\alpha) \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}(\omega) - \mathcal{G}(\omega)|^2 \\ &\geq \varepsilon^2 \eta^2(\alpha) [1 - \zeta(\alpha^{-1})]. \end{aligned} \quad (12)$$

The above chain of inequalities hold for all $\alpha > C$, therefore we have:

$$\|(f - g) * Q\|_2 \geq \frac{\varepsilon}{(2\pi)^{K/2}} \sup_{\alpha} \eta(\alpha) \sqrt{1 - \zeta(\alpha^{-1})},$$

which completes the proof. \square

Theorem 3.6 presents a general approach to prove the recoverability of latent functions (or objects, which are the main focus in this work) from a certain class of independent additive noise distributions. This approach works as long as the function class as well as the noise distribution are mostly comprised of low-frequency components in the Fourier domain. For example, the Gaussian noise hurts low-frequency parts of a geometric object far less than its high-frequency details. More specifically, we prove the following corollary for Theorem 3.6:

Corollary 3.7 (Recoverability from Additive Gaussian Noise $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$). *Consider the setting in Theorem 3.6, and assume the noise distribution follows $Q \triangleq \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ for $\sigma > 0$. Then, as long as for $f, g \in \mathcal{F}$ we have $\|f - g\|_2 \geq \varepsilon$ for some $\varepsilon \geq 0$, we also have*

$$\|(f - g) * Q\|_2 \geq \frac{\varepsilon}{(2\pi)^{K/2}} \left(\sup_{\alpha > C} \sqrt{1 - \zeta\left(\frac{1}{\alpha}\right)} e^{-K(\sigma\alpha)^2/2} \right) \quad (13)$$

Proof. The Fourier transform of $Q = \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ can be computed as follows:

$$\mathcal{F}\{Q\}(\omega) = \prod_{i=1}^K \mathcal{F}\{\mathcal{N}(0, \sigma^2)\}(\omega_i) = e^{-\sigma^2 \|\omega\|_2^2/2}.$$

Also, it can be easily checked that

$$\inf_{\|\omega\|_{\infty} \leq \alpha} e^{-\sigma^2 \|\omega\|_2^2/2} = e^{-\sigma^2/2(\alpha^2 + \dots + \alpha^2)} = e^{-K(\sigma\alpha)^2/2}.$$

By substitution into the end result of Theorem 3.6, the claimed bounds can be achieved and the proof is complete. \square

In this regard, our main explicit theoretical contribution in this section with respect to simplices has been stated in the following theorem:

Theorem 3.8 (Recoverability of Simplices from Additive Noise). *For any two $(\bar{\theta}, \underline{\theta})$ -isoperimetric simplices $\mathcal{S}_1, \mathcal{S}_2 \in \mathbb{S}_K$, given that*

$$\|f_{\mathcal{S}_1} * G_{\sigma} - f_{\mathcal{S}_2} * G_{\sigma}\|_{\text{TV}} \leq \varepsilon$$

for some $\varepsilon \geq 0$, we have

$$\|f_{\mathcal{S}_1} - f_{\mathcal{S}_2}\|_2 \leq \varepsilon e^{\Omega\left(\frac{K}{\text{SNR}^2}\right)},$$

where SNR denotes the effective signal-to-noise ratio which is defined earlier.

Proof is given in Section C of the supplementary document. The core idea is to first showing that simplices, in general, are low-frequency objects in the Fourier domain which holds due to their convexity and simple geometric shape. More specifically, Lemma C.3 shows that for a geometrically regular simplex $\mathcal{S} \in \mathbb{S}_K$, we have

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_{\infty} \geq \alpha} |\mathcal{F}\{f_{\mathcal{S}}\}(\omega)|^2 \leq \frac{1}{\text{Vol}(\mathcal{S})} \mathcal{O}\left(\frac{K}{\alpha}\right),$$

for a sufficiently large $\alpha > 0$, where constants in $\mathcal{O}(\cdot)$ only depend on regularity parameters of \mathcal{S} . This inequality is trivial when $K = 1$, since a one-dimensional simplex is a pulse function whose Fourier transform is a $\text{sinc}(x) = \frac{\sin ax}{ax}$ for some $a > 0$. However, showing this relation for larger values of K requires more mathematical effort which is already carried out in the proof of Theorem 3.8. Furthermore, we show a similar inequality holds for the difference of simplices as well. This would ultimately enable us to use the result of Corollary 3.7 to prove the main claim of the paper.

4. Conclusions

We present the first sample complexity bounds for consistent learning of high-dimensional simplices in noisy regimes. Formally speaking, we prove that given a sufficient amount of noisy data, one can estimate the true simplex up to an arbitrarily small ℓ_2 -error. Also, a presumably optimal polynomial dependence on a number of parameters, such as dimension has been achieved which matches those of the already-solved noiseless case. More interestingly, we have found theoretical justification for an already observed phenomenon in practice: The performance of most heuristic methods in learning of high-dimensional simplices undergo a rather sharp phase transition w.r.t. the noise power, where for $\text{SNR} \geq \Omega(K^{1/2})$ they behave almost identical to the noiseless case, but their performance severely degrades when SNR becomes smaller. Our proofs are based on a number of recent techniques which have been previously used on learning of Gaussian mixture models, a number of tools from high-dimensional geometry and a novel Fourier-based technique for denoising. The latter technique has applications a broader spectra of problems where general recovery of distribution families with certain low-frequency properties from additive Gaussian noise is under question. For future directions in this area, one can think of analyzing this problem under a broader noise and distortion model which might match the practice even more. Also, finding the first non-trivial and amenable polynomial-time algorithms for this problem is an interesting line of work which has remained open until this day.

5. Acknowledgements

This research was partially supported by a grant from IPM.

References

- Agathos, A., Li, J., Bioucas-Dias, J. M., and Plaza, A. Robust minimum volume simplex analysis for hyperspectral unmixing. In *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*, pp. 1582–1586. IEEE, 2014.
- Ambikapathi, A., Chan, T.-H., Ma, W.-K., and Chi, C.-Y. Chance-constrained robust minimum-volume enclosing simplex algorithm for hyperspectral unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 49(11): 4194–4209, 2011.
- Anderson, J., Goyal, N., and Rademacher, L. Efficient learning of simplices. In *Conference on Learning Theory*, pp. 1020–1045, 2013.
- Ashtiani, H., Ben-David, S., Harvey, N., Liaw, C., Mehrobian, A., and Plan, Y. Nearly tight sample complexity bounds for learning mixtures of gaussians via sample compression schemes. *Advances in Neural Information Processing Systems*, 31, 2018.
- Bhattacharyya, C. and Kannan, R. Near-optimal sample complexity bounds for learning latent k -polytopes and applications to ad-mixtures. In *International Conference on Machine Learning*, pp. 854–863. PMLR, 2020.
- Bioucas-Dias, J. M., Plaza, A., Dobigeon, N., Parente, M., Du, Q., Gader, P., and Chanussot, J. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE journal of selected topics in applied earth observations and remote sensing*, 5(2):354–379, 2012.
- Chan, T.-H., Chi, C.-Y., Huang, Y.-M., and Ma, W.-K. A convex analysis-based minimum-volume enclosing simplex algorithm for hyperspectral unmixing. *IEEE Transactions on Signal Processing*, 57(11):4418–4432, 2009.
- Devroye, L. and Lugosi, G. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.
- Devroye, L., Mehrabian, A., and Reddad, T. The total variation distance between high-dimensional gaussians. *arXiv preprint arXiv:1810.08693*, 2018.
- Korem, Y., Szekely, P., Hart, Y., Sheftel, H., Hausser, J., Mayo, A., Rothenberg, M. E., Kalisky, T., and Alon, U. Geometry of the gene expression space of individual cells. *PLoS computational biology*, 11(7):e1004224, 2015.
- Lin, C.-H., Ambikapathi, A., Li, W.-C., and Chi, C.-Y. On the endmember identifiability of craig’s criterion for hyperspectral unmixing: A statistical analysis for three-source case. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pp. 2139–2143. IEEE, 2013.
- Najafi, A., Ilchi, S., Saberi, A. H., Motahari, S. A., Khalaj, B. H., and Rabiee, H. R. On statistical learning of simplices: Unmixing problem revisited. *The Annals of Statistics*, 49(3):1626–1655, 2021.
- Piper, J., Pauca, V. P., Plemmons, R. J., and Giffin, M. Object characterization from spectral data using nonnegative factorization and information theory. In *Proceedings of AMOS technical conference*, 2004.
- Satas, G. and Raphael, B. J. Tumor phylogeny inference using tree-constrained importance sampling. *Bioinformatics*, 33(14):i152–i160, 2017.
- Schwartz, R. and Shackney, S. E. Applying unmixing to gene expression data for tumor phylogeny inference. *BMC bioinformatics*, 11(1):42, 2010.
- Shoval, O., Sheftel, H., Shinar, G., Hart, Y., Ramote, O., Mayo, A., Dekel, E., Kavanagh, K., and Alon, U. Evolutionary trade-offs, pareto optimality, and the geometry of phenotype space. *Science*, 336(6085):1157–1160, 2012.
- Tolliver, D., Tsourakakis, C., Subramanian, A., Shackney, S., and Schwartz, R. Robust unmixing of tumor states in array comparative genomic hybridization data. *Bioinformatics*, 26(12):i106–i114, 2010.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Zhang, S., Agathos, A., and Li, J. Robust minimum volume simplex analysis for hyperspectral unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 55(11): 6431–6439, 2017.
- Zuckerman, N. S., Noam, Y., Goldsmith, A. J., and Lee, P. P. A self-directed method for cell-type identification and separation of gene expression microarrays. *PLoS computational biology*, 9(8):e1003189, 2013.

A. Proof of Theorems

proof of Theorem 3.4. Suppose that \mathcal{S} is a $(\underline{\theta}, \bar{\theta})$ -isoperimetric simplex in \mathbb{R}^K which is confined in a K -dimensional hyper-sphere $C^K(\mathbf{p}, R)$. From Lemma 3.2, we know how to build an ϵ -representative set, for any $\epsilon > 0$, for all $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices in $C^K(\mathbf{p}, R)$. Let us denote this set with $\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$.

Now Consider the class of all isotropic Gaussian noise, $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, with $\sigma \leq R_n$. We denote this class of distributions with $\mathfrak{N}^K(R_n)$. Consider a set $C_{R_n}^{\epsilon/\sqrt{K}} = \left\{0, \frac{\epsilon}{\sqrt{K}}, \frac{2\epsilon}{\sqrt{K}}, \dots, R_n\right\}$, which $\frac{\epsilon}{\sqrt{K}}$ -covers the interval $[0, R_n]$. Now for all $\sigma_i \in C_{R_n}^{\epsilon/\sqrt{K}}$ we put $\mathcal{N}(\mathbf{0}, \sigma_i^2 \mathbf{I})$ in a set called $\widehat{\mathfrak{N}}^K(R_n)$. From Theorem 1.1 in (Devroye et al., 2018) it can be shown that $\widehat{\mathfrak{N}}^K(R_n)$ is an ϵ -representative set for $\mathfrak{N}^K(R_n)$. According to the way we build $\widehat{\mathfrak{N}}^K(R_n)$, we have $|\widehat{\mathfrak{N}}^K(R_n)| \leq \frac{R_n \sqrt{K}}{\epsilon}$. Now we build a set of noisy simplices as follows:

$$\widehat{\mathbb{G}}_{K,\sigma}(R, R_n) \triangleq \left\{ f_{\mathcal{S}}(\mathbf{x}) * G_{\sigma}(\mathbf{x}) \mid \mathcal{S} \in \widehat{\mathbb{S}}(C^K(\mathbf{p}, R)), G_{\sigma} \in \widehat{\mathfrak{N}}^K(R_n) \right\}. \quad (14)$$

Now for any density function $f_{\mathcal{S}}(\mathbf{x}) * G_{\sigma}(\mathbf{x})$, where $\mathcal{S} \in C^K(\mathbf{p}, R)$ and $\sigma \leq R_n$, we can find some $\mathcal{S}^* \in \widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$, and $G_{\sigma^*} \in \widehat{\mathfrak{N}}^K(R_n)$ such that

$$\begin{aligned} \|f_{\mathcal{S}} - f_{\mathcal{S}^*}\|_{\text{TV}} &\leq \epsilon, \\ \|G_{\sigma} - G_{\sigma^*}\|_{\text{TV}} &\leq \epsilon, \\ f_{\mathcal{S}^*} * G_{\sigma^*} &\in \widehat{\mathbb{G}}_{K,\sigma}(R, R_n). \end{aligned}$$

Now for the distance between $f_{\mathcal{S}^*}^{\sigma^*}$ and $f_{\mathcal{S}}^{\sigma}$ we have:

$$\begin{aligned} \|f_{\mathcal{S}^*} * G_{\sigma^*} - f_{\mathcal{S}} * G_{\sigma}\|_{\text{TV}} &\leq \|f_{\mathcal{S}^*} * G_{\sigma^*} - f_{\mathcal{S}^*} * G_{\sigma}\|_{\text{TV}} + \|f_{\mathcal{S}^*} * G_{\sigma} - f_{\mathcal{S}} * G_{\sigma}\|_{\text{TV}} \\ &\leq \|G_{\sigma^*} - G_{\sigma}\|_{\text{TV}} + \|f_{\mathcal{S}^*} - f_{\mathcal{S}}\|_{\text{TV}} \\ &\leq 2\epsilon. \end{aligned} \quad (15)$$

From the above inequalities it can be seen that for any density function $f_{\mathcal{S}} * G_{\sigma}$, where $\mathcal{S} \in C^K(\mathbf{p}, R)$ and $\sigma \leq R_n$, there exist some density function $f^* \in \widehat{\mathbb{S}}_n(C^K(\mathbf{p}, R), R_n)$ where $\|f^* - f_{\mathcal{S}} * G_{\sigma}\|_{\text{TV}} \leq 2\epsilon$. Therefore the set $\widehat{\mathbb{G}}_{K,\sigma}(R, R_n)$ is a 2ϵ -representative set for the class of K -simplices confined in a hyper-sphere with radius R which convolved with a Gaussian noise with variance $\sigma \leq R_n$. We show this class with

$$\mathbb{G}_{K,\sigma}(R, R_n) \triangleq \{f_{\mathcal{S}} * G_{\sigma} \mid \mathcal{S} \in \mathbb{S}_K, \mathcal{S} \in C^K(\mathbf{p}, R), \sigma \leq R_n\} \quad (16)$$

Assume that we have a set of i.i.d. samples from some distribution $f_{\mathcal{S}} * G_{\sigma} \in \mathbb{G}_{K,\sigma}(R, R_n)$. From Theorem 3.3, we know that there exists a deterministic algorithm \mathcal{A} such that given

$$n \geq \frac{\log(3|\widehat{\mathbb{G}}_{K,\sigma}(R, R_n)|^2/\delta)}{2\epsilon^2}$$

i.i.d. samples from $f_{\mathcal{S}} * G_{\sigma}$, the output of the algorithm denoted by $f_{\mathcal{S}_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}}$, with probability at least $1 - \delta$, satisfies:

$$\begin{aligned} \|f_{\mathcal{S}_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}} - f_{\mathcal{S}} * G_{\sigma}\|_{\text{TV}} &\leq 3 \min_{f \in \widehat{\mathbb{G}}_{K,\sigma}(R, R_n)} \|f - f_{\mathcal{S}} * G_{\sigma}\|_{\text{TV}} + 4\epsilon \\ &\leq 6\epsilon + 4\epsilon = 10\epsilon. \end{aligned} \quad (17)$$

For the cardinality of $\widehat{\mathbb{G}}_{K,\sigma}(R, R_n)$ we have:

$$\begin{aligned} |\widehat{\mathbb{G}}_{K,\sigma}(R, R_n)| &= |\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))| |\widehat{\mathfrak{N}}^K(R_n)| \\ &\leq |\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))| \frac{R_n \sqrt{K}}{\epsilon}. \end{aligned} \quad (18)$$

And, according to Lemma 3.2 the followings hold for the cardinality of $\widehat{\mathbb{S}}(C^K(\mathbf{p}, R))$:

$$\begin{aligned} \left| \widehat{\mathbb{S}}(C^K(\mathbf{p}, R)) \right| &= \binom{\left| T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R)) \right|}{K+1} \\ &\leq \left| T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R)) \right|^{K+1} \\ &\leq \left(\left(1 + \frac{2(K+1)R}{\alpha\epsilon} \right)^K \right)^{K+1} \\ &= \left(1 + \frac{2(K+1)R}{\alpha\epsilon} \right)^{K(K+1)}. \end{aligned} \quad (19)$$

Now from 19 and 18 we have:

$$|\widehat{\mathbb{G}}_{K,\sigma}(R, R_n)| \leq \frac{R_n \sqrt{K}}{\epsilon} \left(1 + \frac{2(K+1)R}{\alpha\epsilon} \right)^{K(K+1)}. \quad (20)$$

Then using 20 and 17, we can say that for any $\epsilon, \delta > 0$, there exists a PAC-learning algorithm \mathcal{A} for the class of noisy $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices in $\mathbb{G}_{K,\sigma}(R, R_n)$, whose sample complexity is bounded as follows:

$$\begin{aligned} n &\geq 50 \frac{\log \frac{10R_n \sqrt{K}}{\epsilon} + 2(K+1)^2 \log \left(1 + \frac{20(K+1)R}{\alpha\epsilon} \right) + \log \frac{3}{\delta}}{\epsilon^2} \\ &= 50 \frac{\log \frac{30R_n \sqrt{K}}{\delta\epsilon} + 2(K+1)^2 \log \left(1 + \frac{100R\bar{\theta}(K+1)}{\epsilon \text{Vol}(S)^{\frac{1}{K}}} \right)}{\epsilon^2}. \end{aligned} \quad (21)$$

In other words, given that the number of samples n satisfies the above lower-bound, then with probability at least $1 - \delta$ we have

$$\|f_{S_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}} - f_S * G_{\sigma}\|_{\text{TV}} \leq \epsilon.$$

This completes the proof. Using simple algebra, the bound can be further simplified into $n \geq O\left(\frac{K^2}{\epsilon^2} \log \frac{K}{\epsilon}\right)$ which completes the proof. \square

proof of Theorem 3.5. From Theorem 3.4, we know that the class of $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplices contained in the K -dimensional hyper-sphere $C^K(\mathbf{p}, R)$ and convolved with an isotropic Gaussian noise $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, with $\sigma \leq R_n$, is PAC-learnable, with sample complexity $O\left(\frac{K^2}{\epsilon^2} \log \frac{K}{\epsilon}\right)$, where ϵ is defined accordingly. And from Lemma 3.1, we know that if we have $O(K^2)$ samples from a noisy simplex $f_S * \mathbb{G}_{\sigma}$, we can find a K -dimensional sphere which with probability at least $1 - \delta/2$ contains the true simplex as long as the radius R of the sphere and the upper-bound of the noise variance R_n satisfy

$$R \leq 4\sqrt{K+1} \left(1 + \frac{K+2}{d_S/\sigma} \right) d_S \quad (22)$$

$$R_n \leq \frac{K+2}{K-3} \left(1 + \frac{d_S/\sigma}{K+2} \right) \sigma. \quad (23)$$

Therefore, it can be shown that there exists an algorithm \mathcal{A} such that given n i.i.d samples from \mathbb{G}_S with

$$\begin{aligned} n &\geq 50 \frac{\log \frac{30R_n \sqrt{K}}{\delta\epsilon} + 2(K+1)^2 \log \left(1 + \frac{100R\bar{\theta}(K+1)}{\epsilon \text{Vol}(S)^{\frac{1}{K}}} \right)}{\epsilon^2} + 2000(K+1)(K+2) \log \frac{6}{\delta} \\ &= 100 \frac{\log 6/\delta + (K+2)^2 \log \left(1 + \frac{100\bar{\theta}(K+1)^{3/2}}{\epsilon \text{Vol}(S)^{\frac{1}{K}}} \left(1 + \frac{K+2}{d_S/\sigma} \right) d_S \right)}{\epsilon^2} \\ &= O\left(\frac{K^2}{\epsilon^2} \log \frac{K}{\epsilon}\right), \end{aligned} \quad (24)$$

the output of the algorithm $f_{\mathcal{S}_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}}$ with probability at least $1 - \delta$ satisfies

$$\|f_{\mathcal{S}_{\mathcal{A}}} * G_{\sigma_{\mathcal{A}}} - f_{\mathcal{S}} * \mathbb{G}_{\sigma}\|_{\text{TV}} \leq \epsilon.$$

It should be noted that R and R_n in (24) is already replaced with the bound in the r.h.s of (22) and (23). This way, the proof is completed. \square

B. Proof of Lemmas

proof of Lemma 3.1. To find a K -dimensional sphere containing the true simplex, it suffices to find a point \mathbf{p} inside the simplex and an upper-bound R for its diameter, i.e., the maximum distance between two points in the simplex. Obviously, this ensures that the K -dimensional sphere with center point \mathbf{p} and radius R contains the true simplex. For the maximum distance between any two points inside $\mathcal{S} \in \mathbb{S}_K$, denoted by $d_{\mathcal{S}}$, we have

$$\begin{aligned} d_{\mathcal{S}} &= \max_{\mathbf{x}, \mathbf{y} \in \mathcal{S}} \|\mathbf{x} - \mathbf{y}\|_2 \\ &= \max_{\phi_{\mathbf{x}}, \phi_{\mathbf{y}} \in \mathcal{S}_S^K} \|\mathbf{V}_{\mathcal{S}} \phi_{\mathbf{x}} - \mathbf{V}_{\mathcal{S}} \phi_{\mathbf{y}}\|_2 \\ &= \max_{\phi_{\mathbf{x}}, \phi_{\mathbf{y}} \in \mathcal{S}_S^K} \|\mathbf{V}_{\mathcal{S}} (\phi_{\mathbf{x}} - \phi_{\mathbf{y}})\|_2, \end{aligned} \quad (25)$$

where \mathcal{S}_S^K represents the standard simplex in \mathbb{R}^K and $\mathbf{V}_{\mathcal{S}}$ denotes the vertex matrix of \mathcal{S} (Equation 4 in the main paper). Let $\sigma_{\max}(\mathcal{S})$ denote the maximum singular-value for the matrix $\mathbf{V}_{\mathcal{S}}$, and assume \mathbf{v}_{\max} be the unitary eigenvector that corresponds to $\sigma_{\max}(\mathcal{S})$. To be more precise, let $\mathbf{V}_{\mathcal{S}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ be the singular value decomposition of the vertex matrix. Then, $\sigma_{\max}(\mathcal{S})$ refers to the largest singular value on the main diagonal of $\mathbf{\Sigma}$, and \mathbf{v}_{\max} represents the corresponding column in \mathbf{V} . Then, for any point $\mathbf{x} \in \mathbb{R}^{K+1}$, one can write

$$\|\mathbf{V}_{\mathcal{S}} \mathbf{x}\|_2^2 = \sum_i \Sigma_{i,i}^2 \left\| \mathbf{U}_i \mathbf{V}_i^T \mathbf{x} \right\|_2^2 = \sum_i \Sigma_{i,i}^2 \left(\mathbf{V}_i^T \mathbf{x} \right)^2,$$

which results in the following useful inequality:

$$\|\mathbf{V}_{\mathcal{S}} \mathbf{x}\|_2 \geq \sigma_{\max}(\mathcal{S}) |\mathbf{v}_{\max}^T \mathbf{x}|.$$

For $\phi_{\mathbf{x}} \neq \phi_{\mathbf{y}}$, let $\mathbf{r} \triangleq (\phi_{\mathbf{x}} - \phi_{\mathbf{y}}) / \|\phi_{\mathbf{x}} - \phi_{\mathbf{y}}\|_2$, and $\ell \triangleq \|\phi_{\mathbf{x}} - \phi_{\mathbf{y}}\|_2$. In other words, let \mathbf{r} represent the unitary direction vector for the difference, and ℓ to denote the corresponding length. In this regard, we have

$$\max_{\phi_{\mathbf{x}}, \phi_{\mathbf{y}} \in \mathcal{S}_S^K} \|\mathbf{V}_{\mathcal{S}} (\phi_{\mathbf{x}} - \phi_{\mathbf{y}})\|_2 \geq \sigma_{\max}(\mathcal{S}) \max_{\mathbf{r}, \ell} \ell |\mathbf{v}_{\max}^T \mathbf{r}| \geq \sigma_{\max}(\mathcal{S}) \max_{\ell} \ell. \quad (26)$$

Now we should find a lower-bound for maximum possible ℓ that can be achieved for some fixed direction $\mathbf{r} = \mathbf{v}_{\max}$. To this aim, assume the difference vector $\phi_{\mathbf{x}} - \phi_{\mathbf{y}}$ starts from any arbitrary vertex, i.e., $\phi_{\mathbf{x}}$ is a one-hot vector with K components equal to zero and the remaining one equal to 1. In this regard, the minimum possible ℓ that can be achieved (minimum is taken w.r.t. direction of the difference vector) occurs when the difference vector becomes perpendicular to its front facet. Since all ϕ vectors belong to the standard simplex, such difference vectors can be easily shown to be of the form Δ with

$$\Delta_i = 1 \quad \text{and} \quad \Delta_{j|j \neq i} = -1/K,$$

for any $i \in 0, \dots, K$. The ℓ_2 -norm of all such vectors equals to $\sqrt{1 + 1/K}$. Thus, we have

$$d_{\mathcal{S}} \geq \sqrt{\frac{K+1}{K}} \sigma_{\max}(\mathcal{S}) \geq \sigma_{\max}(\mathcal{S}). \quad (27)$$

Assume we have $2m$ i.i.d. samples drawn from a noisy simplex $f_{\mathcal{S}} * \mathbb{G}_{\sigma}$, denoted by $\{\mathbf{y}\}_{i=1}^{2m}$. This way, we have $\mathbf{y}_i = \mathbf{V}_{\mathcal{S}} \phi_i + \mathbf{z}_i$ for $i = 1, \dots, 2m$, where ϕ_i represents the weight vector for the i th sample (drawn from a uniform Dirichlet distribution), while $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ is an independently drawn noise vector. Let us define \mathbf{D} as

$$\mathbf{D} = \frac{1}{2m} \sum_{i=1}^m \|\mathbf{y}_{2i} - \mathbf{y}_{2i-1}\|_2^2. \quad (28)$$

To find an upper-bound for d_S , it is enough to find a high probability lower-bound for D in terms of d_S . To this aim, we rewrite D as

$$\begin{aligned}
 D &= \frac{1}{2m} \sum_{i=1}^m \|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1} + \mathbf{z}_{2i} - \mathbf{z}_{2i-1}\|_2^2 \\
 &= \frac{1}{2m} \sum_{i=1}^m \|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2 + \frac{1}{2m} \sum_{i=1}^m \|\mathbf{z}_{2i} - \mathbf{z}_{2i-1}\|_2^2 \\
 &\quad + \frac{1}{m} \sum_{i=1}^m (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T (\mathbf{z}_{2i} - \mathbf{z}_{2i-1}) \\
 &= \frac{1}{2m} \sum_{i=1}^m \|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2 + \frac{2\sigma^2}{2m} \sum_{i=1}^m \|\tilde{\mathbf{z}}_{2i} - \tilde{\mathbf{z}}_{2i-1}\|_2^2 \\
 &\quad + \frac{\sigma\sqrt{2}}{m} \sum_{i=1}^m (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T (\tilde{\mathbf{z}}_{2i} - \tilde{\mathbf{z}}_{2i-1}). \tag{29}
 \end{aligned}$$

Let us denote the three terms in r.h.s. of (29) as I, II and III, respectively. Also, we have $\tilde{\mathbf{z}}_i = \mathbf{z}_i/\sigma$. To find a lower-bound for D , we should find respective lower-bounds for I, II and III. First, let us discuss about (I):

$$\text{I} = \frac{1}{2m} \sum_{i=1}^m \|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2 = f_{\mathbf{V}_S}(\phi_1, \phi_2, \dots, \phi_{2m}). \tag{30}$$

For the term inside of the summation, we have

$$\|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2 \stackrel{a.s.}{\leq} d_S^2, \tag{31}$$

$$\text{Var}(\|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2) \leq \frac{12\sigma_{\max}^4(\mathcal{S})}{K^3} \leq \frac{12d_S^4}{K^3} \tag{32}$$

From the above inequalities we can see that $\|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2$ satisfies the Bernstein's condition with $b = 2d_S^2$ (see eq. 2.15 in (Wainwright, 2019)). Then (from proposition 2.10 in (Wainwright, 2019)) the following inequality holds with probability at least $1 - \delta/6$, for any $\delta \in (0, 1)$:

$$f_{\mathbf{V}_S}(\phi_1, \dots, \phi_{2m}) \geq \mathbb{E}[f_{\mathbf{V}_S}(\phi_1, \dots, \phi_{2m})] - 2\sqrt{\max\left\{\frac{4\log\frac{6}{\delta}}{m}, \frac{3}{K^3}\right\} \cdot \frac{d_S^4}{m} \log\frac{6}{\delta}}, \tag{33}$$

where for the expected value of the function f we have:

$$\begin{aligned}
 \mathbb{E}[f_{\mathbf{V}_S}(\phi_1, \dots, \phi_{2m})] &= \frac{1}{2m} \sum_{i=1}^m \mathbb{E}[\|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2^2] \\
 &= \frac{1}{2} \mathbb{E}[\|\mathbf{V}_S \phi - \mathbf{V}_S \phi'\|_2^2] \\
 &= \frac{1}{2} \mathbb{E}[(\mathbf{V}_S \phi - \mathbf{V}_S \phi')^T (\mathbf{V}_S \phi - \mathbf{V}_S \phi')] \\
 &= \frac{1}{2} \mathbb{E}[(\phi - \phi')^T \mathbf{V}_S^T \mathbf{V}_S (\phi - \phi')] \\
 &= \frac{1}{2} \mathbb{E}[\text{Tr}(\mathbf{V}_S^T \mathbf{V}_S (\phi - \phi') (\phi - \phi')^T)] \\
 &= \frac{1}{2} \text{Tr}(\mathbf{V}_S^T \mathbf{V}_S \mathbb{E}[(\phi - \phi') (\phi - \phi')^T]) \\
 &= \frac{1}{2} \text{Tr} \left(\mathbf{V}_S^T \mathbf{V}_S \frac{1}{(K+1)^2(K+2)} \begin{bmatrix} K & -1 & \dots & -1 \\ -1 & K & \dots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \dots & K \end{bmatrix} \right) \\
 &= \frac{1}{2(K+1)^2(K+2)} \text{Tr}(\mathbf{V}_S^T \mathbf{V}_S ((K+1)\mathbf{I} - \mathbf{1}\mathbf{1}^T)) \\
 &= \frac{1}{2(K+1)^2(K+2)} \left((K+1) \text{Tr}(\mathbf{V}_S^T \mathbf{V}_S) - \text{Tr}(\mathbf{1}^T \mathbf{V}_S^T \mathbf{V}_S \mathbf{1}) \right) \\
 &= \frac{1}{2(K+2)} \left(\sum_{i=1}^K \frac{1}{K+1} \sum_{j=1}^{K+1} \mathbf{V}_S^2(i, j) - \sum_{i=1}^K \left(\frac{1}{K+1} \sum_{j=1}^{K+1} \mathbf{V}_S(i, j) \right)^2 \right) \\
 &= \frac{1}{2(K+2)} \sum_{i=1}^K \frac{1}{K+1} \sum_{j=1}^{K+1} \left(\mathbf{V}_S(i, j) - \frac{1}{K+1} \sum_{j=1}^{K+1} \mathbf{V}_S(i, j) \right)^2 \\
 &= \frac{1}{2(K+2)} \sum_{i=1}^K \frac{1}{K+1} \sum_{j=1}^{K+1} \left(\hat{\mathbf{V}}_S(i, j) - \frac{1}{K+1} \sum_{j=1}^{K+1} \hat{\mathbf{V}}_S(i, j) \right)^2 \\
 &= \frac{1}{2(K+2)} \sum_{i=1}^K \frac{1}{K+1} \sum_{j=1}^{K+1} \tilde{\mathbf{V}}_S(i, j)^2 \\
 &= \frac{1}{2(K+2)(K+1)} \sum_{j=1}^{K+1} \|\tilde{\mathbf{V}}_S^j\|_2^2
 \end{aligned} \tag{34}$$

where $\tilde{\mathbf{V}}_S$ and $\hat{\mathbf{V}}_S$ are defined as

$$\mathbf{V}_S = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{K+1}] \tag{35}$$

$$\begin{aligned}
 \hat{\mathbf{V}}_S &= [\mathbf{0}, \mathbf{v}_2 - \mathbf{v}_1, \mathbf{v}_3 - \mathbf{v}_1, \dots, \mathbf{v}_{K+1} - \mathbf{v}_1] \\
 &= [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_{K+1}]
 \end{aligned} \tag{36}$$

$$\begin{aligned}
 \tilde{\mathbf{V}}_S &= \left[\hat{\mathbf{v}}_1 - \frac{\sum_1^K \hat{\mathbf{v}}_i}{K+1}, \hat{\mathbf{v}}_2 - \frac{\sum_1^K \hat{\mathbf{v}}_i}{K+1}, \dots, \hat{\mathbf{v}}_K - \frac{\sum_1^K \hat{\mathbf{v}}_i}{K+1} \right] \\
 &= [\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2, \dots, \tilde{\mathbf{v}}_K, \tilde{\mathbf{v}}_{K+1}],
 \end{aligned} \tag{37}$$

And $\tilde{\mathbf{V}}_S^j$ is the j th column of $\tilde{\mathbf{V}}_S$. We know that the maximum distance between any two points inside a simplex is equal to the length of the biggest edge of the simplex. Now without loss of generality suppose that \mathbf{v}_2 and \mathbf{v}_3 are the vertices

related to this edge (if more than one edge have the maximum length, suppose that, v_2 and v_3 are belong to one of them). Therefore the length of the vector $v_3 - v_2$ associated with the maximum distance satisfies the following

$$\begin{aligned}
 d_S^2 &= \|v_3 - v_2\|_2^2 \\
 &= \|\tilde{v}_3 - \tilde{v}_2\|_2^2 \\
 &\leq \|\tilde{v}_3\|_2^2 + \|\tilde{v}_2\|_2^2 + 2\|\tilde{v}_3\|_2\|\tilde{v}_2\|_2 \\
 &\leq 2(\|\tilde{v}_3\|_2^2 + \|\tilde{v}_2\|_2^2).
 \end{aligned} \tag{38}$$

Now the following inequality can be deduced from 34 and 38

$$\mathbb{E}[f_{V_S}(\phi_1, \dots, \phi_{2m})] \geq \frac{1}{4(K+1)(K+2)} d_S^2. \tag{39}$$

From 33 and 39, it is easy to show that if we have $2m \geq 2000(K+1)(K+2) \log \frac{6}{\delta}$ i.i.d. samples from \mathbb{P}_S , then with probability at least $1 - \frac{\delta}{6}$ (for any $0 < \delta \leq 1$), we have

$$f(\phi_1, \phi_2, \dots, \phi_{2m}) \geq \frac{d_S^2}{8(K+1)(K+2)}. \tag{40}$$

Next, we should find a lower-bound for the second term in the r.h.s. of 29, i.e., II.

In 29, z_i s are drawn from a multivariate Gaussian distribution with mean vector $\mathbf{0}$ and covariance matrix $\sigma^2 \mathbf{I}$. Then, $z'_i = \tilde{z}_{2i} - \tilde{z}_{2i-1}$ is also a zero-mean Gaussian random vector with covariance matrix \mathbf{I} , and as a result $\|z'_i\|_2^2$ is a Chi-squared random variable with K degrees of freedom. We know that a Chi-squared random variable with K degrees of freedom is sub-exponential with parameters $(2\sqrt{K}, 4)$ and therefore $\frac{1}{2m} \sum_{i=1}^m \|z'_i\|_2^2$ is a sub-exponential random variable with parameters $(\sqrt{K/m}, 2/m)$. Then the following holds with probability at least $1 - \delta/6$:

$$\begin{aligned}
 \frac{1}{2m} \sum_{i=1}^m \|\tilde{z}_{2i} - \tilde{z}_{2i-1}\|_2^2 &\geq \frac{1}{2m} \sum_{i=1}^m \mathbb{E}[\|\tilde{z}_{2i} - \tilde{z}_{2i-1}\|_2^2] - \sqrt{\frac{2K}{m} \log \frac{6}{\delta}} \\
 &\geq \frac{K}{2} - \sqrt{\frac{2K}{m} \log \frac{6}{\delta}} \\
 &\geq \frac{K}{2} - \sqrt{\frac{K}{1000(K+1)(K+2)}} \\
 &\geq \frac{K}{2} - 1 = \frac{K-2}{2}.
 \end{aligned} \tag{41}$$

Based on the above inequalities, we can give a lower-bound for the second term in the r.h.s. of 29:

$$\text{II} = \frac{2\sigma^2}{2m} \sum_{i=1}^m \|\tilde{z}_{2i} - \tilde{z}_{2i-1}\|_2^2 \geq (K-2)\sigma^2. \tag{42}$$

Now to find a lower bound for D(S), we only need to give a lower bound for III, in the last inequality of 29. For this part we have:

$$\begin{aligned}
 \text{III} &= \frac{\sigma\sqrt{2}}{m} \sum_{i=1}^m (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T (\tilde{z}_{2i} - \tilde{z}_{2i-1}) \\
 &= \frac{\sigma\sqrt{2}}{m} \sum_{i=1}^m (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T z'_i \\
 &= \frac{\sigma\sqrt{2}}{m} \sum_{i=1}^m g_{\phi_{2i}, \phi_{2i-1}}(z_i'^1, z_i'^2, \dots, z_i'^m),
 \end{aligned} \tag{43}$$

where in the above equations we have:

$$\mathbb{E} (g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)) = \mathbb{E} (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T \mathbb{E} (z_i') = 0. \quad (44)$$

To find a lower bound for III we try to bound $g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)$ from below. To do so, we write the concentration inequality for this function as follow:

$$\begin{aligned} \mathbb{P} (g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m) \geq \epsilon) &\leq \frac{\mathbb{E}_{\phi_{2i}, \phi_{2i-1}, z_i'} [e^{\lambda g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)}]}{e^{\lambda \epsilon}} \\ &\leq \frac{\mathbb{E}_{\phi_{2i}, \phi_{2i-1}, z_i'} [e^{\lambda g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)}]}{e^{\lambda \epsilon}} \\ &= \frac{\mathbb{E}_{\phi_{2i}, \phi_{2i-1}} [\mathbb{E}_{z_i'} [e^{\lambda g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)}]]}{e^{\lambda \epsilon}}, \end{aligned} \quad (45)$$

where $g_{\phi_{2i}, \phi_{2i-1}} (z)$ is a lipschitz function with respect to z' :

$$\begin{aligned} |g_{\phi_{2i}, \phi_{2i-1}} (z_1) - g_{\phi_{2i}, \phi_{2i-1}} (z_2)| &\leq \|\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1}\|_2 \|z_1 - z_2\|_2 \\ &\leq 2d_S \|z_1 - z_2\|_2. \end{aligned} \quad (46)$$

From Lemma (2.27) in (Wainwright, 2019) we know that any L -lipschitz function of a Gaussian R.V. is sub-Gaussian. As a result, we have

$$\mathbb{E}_{z_i'} [e^{\lambda g_{\phi_{2i}, \phi_{2i-1}} (z_i'^1, z_i'^2, \dots, z_i'^m)}] \leq e^{\frac{4\pi^2 d_S^2 \lambda^2}{8}}. \quad (47)$$

From inequalities 47, 45 and equation 43 we have:

$$\mathbb{P} \left(\frac{\sigma\sqrt{2}}{2m} \sum_{i=1}^m g_{\phi_{2i}, \phi_{2i-1}} (z_i) \leq -\epsilon \right) \leq \min_{\lambda \geq 0} \frac{\mathbb{E}_{\phi_1, \dots, \phi_{2m}} [e^{\frac{\lambda^2 \pi^2 d_S^2 \sigma^2}{2m}}]}{e^{\lambda \epsilon}} \leq e^{-\frac{m\epsilon^2}{2\pi^2 d_S^2 \sigma^2}}. \quad (48)$$

Then if we have $2m \geq 2000(K+1)(K+2) \log \frac{6}{\delta}$, i.i.d samples, then with probability at least $1 - \frac{\delta}{6}$ we have:

$$\frac{1}{2m} \sum_{i=1}^m (\mathbf{V}_S \phi_{2i} - \mathbf{V}_S \phi_{2i-1})^T (z_{2i} - z_{2i-1}) \geq -\frac{\pi \sigma d_S}{20\sqrt{(K+2)(K+1)}}. \quad (49)$$

Now from inequalities 40, 42, 49 and equation 29 we have:

$$\begin{aligned} D &\geq \frac{d_S^2}{8(K+1)(K+2)} + (K-2)\sigma^2 - \frac{\pi \sigma d_S}{20\sqrt{(K+2)(K+1)}} \\ &= \frac{d_S^2}{8(K+1)(K+2)} + \left(\sigma\sqrt{K-2} - \frac{d_S \sqrt{\pi}}{40\sqrt{(K+1)(K+2)(K-2)}} \right)^2 - \\ &\quad \frac{d_S^2 \pi}{1600(K+1)(K+2)(K-2)} \end{aligned} \quad (50)$$

$$= (K-2)\sigma^2 + \left(\frac{\pi \sigma}{10\sqrt{2}} - \frac{d_S}{2\sqrt{2}\sqrt{(K+1)(K+2)}} \right)^2 - \frac{\pi^2 \sigma^2}{200}. \quad (51)$$

From the inequalities 51 and 50 it can be seen that if we have $2m \geq 2000(K+1)(K+2) \log \frac{6}{\delta}$ then the following inequalities with probability at least $1 - \delta/2$ give upper-bounds for the maximum distance between two points in the simplex \mathcal{S} , and the noise variance σ :

$$d_S \leq 4\sqrt{(K+1)(K+2)}D = R. \quad (52)$$

$$\sigma^2 \leq \frac{D}{K-3} = R_n \quad (53)$$

In the same way we could also give an upper bound for D . The following inequalities hold with probability more than $1 - \delta/2$:

$$\begin{aligned} D &\leq \frac{d_S^2}{K+2} + \sigma^2(K+2) + \frac{\pi\sigma d_S}{20\sqrt{(K+2)(K+1)}} \\ &\leq \left(\frac{d_S}{\sqrt{K+2}} + \sigma\sqrt{K+2} \right)^2 \\ &= \frac{d_S^2}{K+2} \left(1 + \frac{K+2}{d_S/\sigma} \right)^2 \end{aligned} \quad (54)$$

$$= \sigma^2 (K+2) \left(1 + \frac{d_S/\sigma}{K+2} \right)^2. \quad (55)$$

Based on equation 52 and inequations in 54 and 55, we can give an upper bound for the radius R in 52 and noise radius R_n . The following inequality holds with probability at least $1 - \delta$:

$$R \leq 4\sqrt{K+1} \left(1 + \frac{K+2}{d_S/\sigma} \right) d_S \quad (56)$$

$$R_n \leq \frac{K+2}{K-3} \left(1 + \frac{d_S/\sigma}{K+2} \right) \sigma. \quad (57)$$

Now the only thing we should do is to find a point inside of the simplex. To do this we define a new statistic \mathbf{p} as follows:

$$\mathbf{p} = \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{x}_i = \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{V}_S \phi_i + \mathbf{z}_i = \mathbf{V}_S \left(\frac{1}{2m} \sum_{i=1}^{2m} \phi_i \right) + \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{z}_i = \mathbf{I} + \mathbf{II} \quad (58)$$

It is clear that “ \mathbf{I} ” is placed inside of the true simplex. So the distance between \mathbf{p} and a point inside the main simplex can be calculated as follows:

$$\min_{\mathbf{x} \in S} \|\mathbf{p} - \mathbf{x}\|_2 \leq \left\| \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{z}_i \right\|_2. \quad (59)$$

In the above equation $\frac{1}{2m} \sum_{i=1}^{2m} \mathbf{z}_i$, is a zero mean Gaussian random vector with covariance matrix $\frac{\sigma^2}{2m} \mathbf{I}$, therefore $\left\| \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{z}_i \right\|_2^2$ is a chi-squared random variable with K degrees of freedom, and so we can write concentration inequality for it:

$$\mathbb{P} \left[\left| \left\| \frac{1}{2m} \sum_{i=1}^{2m} \mathbf{z}_i \right\|_2^2 - \frac{K\sigma^2}{2m} \right| \geq \epsilon \right] \leq 2e^{-\frac{m^2\epsilon^2}{2K\sigma^4}}. \quad (60)$$

Then with probability at least $1 - \delta$ we have the following inequality:

$$\begin{aligned} \min_{\mathbf{x} \in S} \|\mathbf{p} - \mathbf{x}\|_2^2 &\leq \frac{K\sigma^2}{2m} + \frac{\sigma^2}{m} \sqrt{2K \log \frac{2}{\delta}} \\ &\leq \frac{K\sigma^2}{m} \log \frac{2}{\delta} \end{aligned} \quad (61)$$

Now if we have $2m \geq 2000(K+1)(K+2) \log \frac{6}{\delta}$, then we can rewrite the above inequality as follows:

$$\min_{\mathbf{x} \in S} \|\mathbf{p} - \mathbf{x}\|_2 \leq \frac{\sigma}{10\sqrt{10}\sqrt{K+2}}. \quad (62)$$

As we mentioned earlier, the first part in the right hand side of the inequality 58, $\frac{1}{2m} \sum_{i=1}^{2m} \mathbf{V}_S \phi_i$, belongs to the interior of the simplex. Then a K -dimensional sphere with radius R and centered at this point, with probability at least $1 - \delta$ contains

the simplex \mathcal{S} . However, in the process of learning we do not have access to the noiseless data and thus cannot have such a point as the center of the sphere. Let us show the distance between \mathbf{p} and $\frac{1}{2m} \sum_{i=1}^{2m} \mathbf{V}_{\mathcal{S}}$ with d . Then, it is clear to see that the K -dimensional sphere with radius $R + d$ and center point \mathbf{p} , contains the sphere with center point at $\frac{1}{2m} \sum_{i=1}^{2m} \mathbf{V}_{\mathcal{S}} \phi_i$ and radius R . So any simplex which is placed in sphere $C^K(\frac{1}{2m} \sum_{i=1}^{2m} \mathbf{V}_{\mathcal{S}} \phi_i, R)$ is also placed in sphere $C^K(\mathbf{p}, R + d)$.

So we conclude that if we have $2m \geq 2000(K+1)(K+2) \log \frac{6}{\delta}$ i.i.d. samples from $\mathbb{G}_{\mathcal{S}}$ then the main simplex \mathcal{S} with probability more than or equal to $1 - \delta$ will be confined in a K -dimensional sphere with center point at \mathbf{p} and with radius R :

$$\begin{aligned} R &= 4\sqrt{(K+1)(K+2)D} + \frac{\sqrt{D}}{40\sqrt{K+2}} \\ &\leq 4\sqrt{(K+1)(K+2)D} \left(1 + \frac{1}{160(K+2)\sqrt{K+1}}\right) \\ &\leq 8\sqrt{(K+1)(K+2)D}. \end{aligned} \quad (63)$$

This completes the proof. \square

proof of Lemma 3.2. Consider a $(\underline{\theta}, \bar{\theta})$ -isoperimetric K -simplex \mathcal{S} , which is bounded in a sphere $C^K(\mathbf{p}, R)$. It is clear that all vertices of this simplex $\{\mathbf{v}_1, \dots, \mathbf{v}_{K+1}\}$ are placed in $C^K(\mathbf{p}, R)$. From the definition of $T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R))$, we know that for each vertex of the simplex \mathbf{v}_i , there exists some $\mathbf{v}'_i \in T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R))$ such that $\|\mathbf{v}_i - \mathbf{v}'_i\|_2 \leq \frac{\alpha\epsilon}{K+1}$. Assume for each vertex \mathbf{v}_i of \mathcal{S} , we denote its closest point in $T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R))$ as

$$\hat{\mathbf{v}}_i = \operatorname{argmin} \left\{ \|\mathbf{v}_i - \hat{\mathbf{v}}\|_2 \mid \hat{\mathbf{v}} \in T_{\frac{\alpha\epsilon}{K+1}}(C^K(\mathbf{p}, R)) \right\}.$$

Using these points we make a new simplex $\hat{\mathcal{S}}$. It is clear that $\hat{\mathcal{S}}$ belongs to $\hat{\mathbb{S}}(C^K(\mathbf{p}, R))$. Assume $f_{\mathcal{S}}$ and $f_{\hat{\mathcal{S}}}$ denote the probability density functions that correspond to \mathcal{S} and $\hat{\mathcal{S}}$, respectively. Then, the TV-distance between $f_{\mathcal{S}}$ and $f_{\hat{\mathcal{S}}}$ can be written and bounded as

$$\begin{aligned} \text{TV}(\mathbb{P}_{\mathcal{S}}, \mathbb{P}_{\hat{\mathcal{S}}}) &= \sup_A \mathbb{P}_{\mathcal{S}}(A) - \mathbb{P}_{\hat{\mathcal{S}}}(A) \\ &= \int_{\mathbf{x} \in \{\mathbf{x}' : f_{\mathcal{S}}(\mathbf{x}') \geq f_{\hat{\mathcal{S}}}(\mathbf{x}')\}} f_{\mathcal{S}}(\mathbf{x}) - f_{\hat{\mathcal{S}}}(\mathbf{x}) \\ &= \int_{\mathbf{x} \in \mathcal{S} - \hat{\mathcal{S}}} f_{\mathcal{S}}(\mathbf{x}) - f_{\hat{\mathcal{S}}}(\mathbf{x}) + \left(\int_{\mathbf{x} \in \mathcal{S} \cap \hat{\mathcal{S}}} f_{\mathcal{S}}(\mathbf{x}) - f_{\hat{\mathcal{S}}}(\mathbf{x}) \right) \mathbf{1}(\text{Vol}(\mathcal{S}) \leq \text{Vol}(\hat{\mathcal{S}})) \\ &\leq \int_{\mathbf{x} \in \mathcal{S} - \hat{\mathcal{S}}} f_{\mathcal{S}}(\mathbf{x}) - f_{\hat{\mathcal{S}}}(\mathbf{x}) + \int_{\mathbf{x} \in \mathcal{S} \cap \hat{\mathcal{S}}} |f_{\mathcal{S}}(\mathbf{x}) - f_{\hat{\mathcal{S}}}(\mathbf{x})| \\ &= \frac{\text{Vol}(\mathcal{S} - \hat{\mathcal{S}})}{\text{Vol}(\mathcal{S})} + \text{Vol}(\mathcal{S} \cap \hat{\mathcal{S}}) \left| \frac{1}{\text{Vol}(\mathcal{S})} - \frac{1}{\text{Vol}(\hat{\mathcal{S}})} \right|, \end{aligned} \quad (64)$$

where $\mathcal{S} - \hat{\mathcal{S}}$ shows the set difference between \mathcal{S} and $\hat{\mathcal{S}}$. Let us denote the two terms in the r.h.s. of 64 as I and II, respectively. To find an upper-bound for the TV-distance in 64 we find respective upper-bounds for I and II. In order to do so, first let us discuss about I:

$$\text{I} = \frac{\text{Vol}(\mathcal{S} - \hat{\mathcal{S}})}{\text{Vol}(\mathcal{S})}. \quad (65)$$

The distance between vertices of \mathcal{S} and $\hat{\mathcal{S}}$ are less than $\frac{\alpha\epsilon}{K+1}$, then the maximal difference set (in terms of volume) between \mathcal{S} and $\hat{\mathcal{S}}$ can be bounded as follows: simplex has $K+1$ facets, and the difference set that can occur from altering each of

them is upper-bounded as $\leq \mathcal{A}_i(\mathcal{S}) \times \alpha\epsilon/(K+1)$, where $i = 1, \dots, K+1$ denotes the facet index. This way, for I we have:

$$\begin{aligned} \text{I} &\leq \sum_{i=1}^{K+1} \frac{\alpha\epsilon}{K+1} \frac{\mathcal{A}_i(\mathcal{S})}{\text{Vol}(\mathcal{S})} \\ &\leq \alpha\epsilon \mathcal{A}_{\max}(\mathcal{S}) \\ &\leq \alpha\epsilon \bar{\theta} \text{Vol}^{\frac{1}{K}}, \end{aligned} \quad (66)$$

where $\mathcal{A}_{\max}(\mathcal{S})$ is the volume of the largest facet of \mathcal{S} . In the final inequality of 66, we take advantage of the isoperimetricity property of \mathcal{S} . Next, we should find an upper-bound for the second term in the r.h.s. of 64, i.e., II:

$$\text{II} = \text{Vol}(\mathcal{S} \cap \hat{\mathcal{S}}) \left| \frac{1}{\text{Vol}(\mathcal{S})} - \frac{1}{\text{Vol}(\hat{\mathcal{S}})} \right|. \quad (67)$$

To find an upper-bound for II we should find an upper-bound for $\text{Vol}(\hat{\mathcal{S}})$. To do so, we create a “rounded” simplex \mathcal{S}^r which forms by adding a K -dimensional sphere with radius $r = \frac{\alpha\epsilon}{K+1}$ to the original simplex \mathcal{S} . It can be shown that $\hat{\mathcal{S}}$ is definitely placed inside \mathcal{S}^r and therefore $\text{Vol}(\hat{\mathcal{S}}) \leq \text{Vol}(\mathcal{S}^r)$. In this regard, we have

$$\begin{aligned} \text{II} &= \text{Vol}(\mathcal{S} \cap \hat{\mathcal{S}}) \left| \frac{1}{\text{Vol}(\mathcal{S})} - \frac{1}{\text{Vol}(\hat{\mathcal{S}})} \right| \\ &\leq \left| \frac{\text{Vol}(\hat{\mathcal{S}}) - \text{Vol}(\mathcal{S})}{\text{Vol}(\mathcal{S})} \right| \\ &\leq \left| \frac{\text{Vol}(\mathcal{S}^r) - \text{Vol}(\mathcal{S})}{\text{Vol}(\mathcal{S})} \right| \\ &\leq \left(\sum_{i=1}^{K+1} \frac{\alpha\epsilon}{K+1} \mathcal{A}_i(\mathcal{S}) + (K+1) \mathcal{C}_\alpha \left(\frac{\alpha\epsilon}{K+1} \right)^K \right) \frac{1}{\text{Vol}(\mathcal{S})} \\ &\leq \left((K+1) \frac{\alpha\epsilon}{K+1} \mathcal{A}_{\max}(\mathcal{S}) + (K+1) \mathcal{C}_\alpha \left(\frac{\alpha\epsilon}{K+1} \right)^K \right) \frac{1}{\text{Vol}(\mathcal{S})} \\ &\leq 2\alpha\epsilon \bar{\theta} \text{Vol}^{\frac{1}{K}}, \end{aligned} \quad (68)$$

where $\mathcal{C}_\alpha \left(\frac{\alpha\epsilon}{K+1} \right)^K$ is the volume of a K -dimensional sphere with radius $r = \frac{\alpha\epsilon}{K+1}$. Now, using 66 and 68 we have

$$\text{TV}(\mathbb{P}_{\mathcal{S}}, \mathbb{P}_{\hat{\mathcal{S}}}) \leq 3\alpha\epsilon \bar{\theta} \text{Vol}^{\frac{1}{K}} \leq \frac{3}{5}\epsilon \leq \epsilon. \quad (69)$$

From 69, it can be seen that for any $(\underline{\theta}, \bar{\theta})$ -isoperimetric simplex \mathcal{S} which is bounded in a K -dimensional sphere $C^K(\mathbf{p}, R)$, there exists at least one simplex in $\hat{\mathbb{S}}(C^K(\mathbf{p}, R))$ within a TV-distance of ϵ from \mathcal{S} . Thus the proof is complete.

In the end, it is worth mentioning one possible method to build an ϵ -covering set for a K -dimensional sphere $C^K(\mathbf{p}, R)$. Although there exists deterministic ways to build this set, an alternative approach is by uniformly sampling points from the sphere. Consider an $\epsilon/2$ -packing set with size L for the sphere. Based on the result of the “coupon collector problem” we know that having $L \log(L)$ uniform samples from the sphere guarantees with high probability that there exists at least one sample within a distance of $\epsilon/2$ from each element of the packing set. Therefore, the resulting samples are an ϵ -covering set for the sphere. Hence, the cardinality of an ϵ -covering set built in this way is at most $(1 + \frac{4R}{\epsilon})^{2K}$. \square

C. Consistency Analysis

Assume $\mathcal{S}_1, \mathcal{S}_2 \in \mathbb{S}_K$ represent two arbitrary simplices in \mathbb{R}^K . In this regard, let $\mathbb{P}_{\mathcal{S}_1}$ and $\mathbb{P}_{\mathcal{S}_2}$ denote the probability measures, and $f_{\mathcal{S}_1}$ and $f_{\mathcal{S}_2}$ represent the probability density functions associated to \mathcal{S}_1 and \mathcal{S}_2 , respectively. Let us assume that $\mathcal{S}_1, \mathcal{S}_2$ have a minimum degree of geometric regularity in the following sense.

Definition C.1. For a simplex $\mathcal{S} \in \mathbb{S}_K$ with vertices $\theta_0, \dots, \theta_K \in \mathbb{R}^K$, we say \mathcal{S} is $(\bar{\lambda}, \underline{\lambda})$ -regular if

$$\underline{\lambda} \leq \lambda_{\min}(\Theta) \leq \lambda_{\max}(\Theta) \leq \bar{\lambda},$$

where $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ denotes the largest and smallest eigenvalues of a matrix, respectively. Here, Θ represents the zero-centered vertex matrix of \mathcal{S} , i.e.,

$$\Theta \triangleq [\theta_1 - \theta_0 | \dots | \theta_K - \theta_0].$$

This definition has tight connections to the previously used notion of $(\bar{\theta}, \underline{\theta})$ -isoperimetricity which is already used in the main body of the manuscript. In fact, it can be shown that $\mathcal{L}_{\max} = \mathcal{O}(\bar{\lambda})$ and vice versa.

Our aim is to show that if the noisy versions of \mathcal{S}_1 and \mathcal{S}_2 , which we denote by $f_{\mathcal{S}_1} * G_\sigma$ and $f_{\mathcal{S}_2} * G_\sigma$, respectively, have a maximum total variation distance of at least $\epsilon > 0$, then the TV distance between $\mathbb{P}_{\mathcal{S}_1}$ and $\mathbb{P}_{\mathcal{S}_2}$ is also bounded away from zero according to a function of ϵ, σ and the geometric regularity of simplices \mathcal{S}_1 and \mathcal{S}_2 . The theoretical core behind our method is stated in the following general theorem.

Theorem C.2 (Recovery of Low-Frequency Objects from Additive Noise). *For $K \in \mathbb{N}$, consider a probability density function family $\mathcal{F} \subseteq \mathcal{M}(\mathbb{R}^K)$, i.e., a subset of distributions supported over \mathbb{R}^K . Assume for sufficiently large $\alpha > 0$, the following bound holds for all $f, g \in \mathcal{F}$:*

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} |\mathcal{F}\{f\}(\omega) - \mathcal{F}\{g\}(\omega)|^2 \leq \zeta(\alpha^{-1}) \int_{\mathbb{R}^K} |f - g|^2,$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier transform, and ζ is an increasing function with $\zeta(0) = 0$ and continuity at 0. Also, assume the probability density function $Q \in \mathcal{M}(\mathbb{R}^K)$ again for a sufficiently large $\alpha > 0$ has the following property:

$$\inf_{\|\omega\|_\infty \leq \alpha} \|\mathcal{F}\{Q\}(\omega)\| \geq \eta(\alpha),$$

where $\eta(\cdot)$ is a non-negative decreasing function. Then, there exists a non-negative constant C where for any $\sigma, \epsilon > 0$ and $f, g \in \mathcal{F}$ with $\|f - g\|_2 \geq \epsilon$, we have

$$\|(f - g) * Q\|_2 \geq \frac{\epsilon}{(2\pi)^K} \left(\sup_{\alpha > C} \eta(\alpha) \sqrt{1 - \zeta(\alpha^{-1})} \right),$$

with $*$ denoting the multi-dimensional convolution operator.

Proof. For the sake of simplicity in notations, let $\mathcal{F}, \mathcal{G}, \mathcal{Q} : \mathbb{R}^K \rightarrow \mathbb{C}$ denote the Fourier transforms of f, g and Q , respectively. Due to Parseval's theorem, we have

$$\|f - g\|_2^2 = \frac{1}{(2\pi)^K} \|\mathcal{F} - \mathcal{G}\|_2^2.$$

Also, due to the properties of the Fourier transform, which is the transformation of convolution into direct multiplication, one can write

$$\mathcal{F}\{(f - g) * Q\} = \mathcal{Q}(\mathcal{F} - \mathcal{G}).$$

Thus, there exists universal constant $C > 0$ such that for any $\alpha > C$:

$$\begin{aligned} (2\pi)^K \|(f - g) * Q\|_2^2 &= \int_{\mathbb{R}^K} |\mathcal{Q}(\omega) (\mathcal{F}(\omega) - \mathcal{G}(\omega))|^2 \\ &\geq \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{Q}(\omega)|^2 |\mathcal{F}(\omega) - \mathcal{G}(\omega)|^2 \\ &\geq \eta^2(\alpha) \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}(\omega) - \mathcal{G}(\omega)|^2 \\ &\geq \epsilon^2 \eta^2(\alpha) [1 - \zeta(\alpha^{-1})]. \end{aligned} \tag{70}$$

The above chain of inequalities hold for all $\alpha > C$, therefore we have:

$$\|(f - g) * Q\|_2^2 \geq \frac{\varepsilon}{(2\pi)^K} \sup_{\alpha > C} \eta(\alpha) \sqrt{1 - \zeta(\alpha^{-1})}, \quad (71)$$

which completes the proof. \square

Theorem C.2 presents a general approach to prove the recoverability of latent functions (or objects, which are the main focus in this work) from a certain class of independent additive noise. This approach works as long as the function class as well as the noise distribution are mostly comprised of low-frequency components in the Forier domain. For example, the Gaussian noise hurts low-frequency parts of a geometric object far less than its high-frequency details. More specifically, we prove the following corollary for Theorem C.2:

Corollary C.3 (Recoverability from Additive Gaussian Noise $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$). *Consider the setting in Theorem C.2, and assume the noise distribution follows $Q \triangleq \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ for $\sigma > 0$. Then, as long as for $f, g \in \mathcal{F}$ we have $\|f - g\|_2 \geq \varepsilon$ for some $\varepsilon \geq 0$, we also have*

$$\|(f - g) * Q\|_2 \geq \frac{\varepsilon}{(2\pi)^K} \left(\sup_{\alpha > C} \sqrt{1 - \zeta\left(\frac{1}{\alpha}\right)} e^{-K(\sigma\alpha)^2/2} \right)$$

Proof. The Fourier transform of $Q = \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ can be computed as follows:

$$\mathcal{F}\{Q\}(\omega) = \prod_{i=1}^K \mathcal{F}\{\mathcal{N}(0, \sigma^2)\}(\omega_i) = e^{-\sigma^2 \|\omega\|_2^2/2}. \quad (72)$$

Also, it can be easily checked that

$$\inf_{\|\omega\|_\infty \leq \alpha} e^{-\sigma^2 \|\omega\|_2^2/2} = e^{-\sigma^2/2(\alpha^2 + \dots + \alpha^2)} = e^{-K(\sigma\alpha)^2/2}.$$

By substitution into the end result of Theorem C.2, the claimed bounds can be achieved and the proof is complete. \square

In this regard, our main explicit theoretical contribution in this section with respect to simplices has been stated in the following theorem:

Theorem C.4 (Recoverability of Simplices from Additive Noise). *For any two $(\bar{\lambda}, \underline{\lambda})$ -regular simplices $\mathcal{S}_1, \mathcal{S}_2 \in \mathbb{S}_K$ with $\bar{\lambda}, \underline{\lambda} > 0$, given that*

$$\mathcal{D}_{\text{TV}}(\mathbb{P}_{\mathcal{S}_1 * G_\sigma}, \mathbb{P}_{\mathcal{S}_2 * G_\sigma}) = \frac{1}{2} \int |(f_{\mathcal{S}_1} - f_{\mathcal{S}_2}) * G_\sigma| \leq \varepsilon$$

for some $\varepsilon \geq 0$, where G_σ (for $\sigma \geq 0$) represents the density function associated to a Gaussian measure with zero mean and covariance matrix of $\sigma^2 \mathbf{I}_{K \times K}$ in \mathbb{R}^K , i.e., $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Then, we have

$$\mathcal{D}_{\text{TV}}(\mathbb{P}_{\mathcal{S}_1}, \mathbb{P}_{\mathcal{S}_2}) \leq \varepsilon e^{\Omega\left(\frac{K}{\text{SNR}^2}\right)},$$

where $\text{SNR} \triangleq \frac{\bar{\lambda}}{K\sigma}$ denotes the effective signal-to-noise ratio, which is the ratio of the standard deviation of a scaled uniform Dirichlet distribution to that of the noise, per dimension.

Proof. Proof is based on properties of the Fourier transforms of simplices \mathcal{S}_1 and \mathcal{S}_2 . For $K \in \mathbb{N}$, and any integrable function $f : \mathbb{R}^K \rightarrow \mathbb{R}$, the Fourier transform of f , denoted by $\mathcal{F}\{f\}(\omega) : \mathbb{R}^K \rightarrow \mathbb{C}$ is defined as follows:

$$\mathcal{F}\{f\}(\omega) \triangleq \int_{\mathbf{x} \in \mathbb{R}^K} f(\mathbf{x}) e^{-i\omega^T \mathbf{x}} d\mathbf{x}. \quad (73)$$

Throughout this proof, for $\mathcal{S} \in \mathbb{S}_K$, let us denote by $\mathcal{F}_{\mathcal{S}}$ the Fourier transform of $f_{\mathcal{S}}$, i.e., the uniform proability density function over \mathcal{S} . Also, the inverse Fourier transform which recovers $f_{\mathcal{S}}$ from $\mathcal{F}_{\mathcal{S}}$ can be written as

$$f_{\mathcal{S}}(\mathbf{x}) = \frac{1}{2\pi} \int_{\omega \in \mathbb{R}^K} \mathcal{F}_{\mathcal{S}}(\omega) e^{i\omega^T \mathbf{x}} d\omega.$$

In this regard, let Δ_K represent the standard simplex in \mathbb{R}^K which means a K -simplex with $\theta_0 = \mathbf{0}$ and $\Theta = \mathbf{I}$. We begin by deriving a number of useful properties for Δ_K through the following lemmas.

Lemma C.5. *Let $\mathcal{F}_{\Delta_k}(\omega_1, \dots, \omega_k) : \mathbb{R}^k \rightarrow \mathbb{C}$ for $k \in [K]$ represent the Fourier transform of f_{Δ_k} in \mathbb{R}^k . Also, we have $\omega_{1:k} \triangleq (\omega_1, \dots, \omega_k)$. Then, the following recursive relation holds for $k > 1$:*

$$\mathcal{F}_{\Delta_k}(\omega_{1:k}) = \frac{k}{i\omega_k} \left[\mathcal{F}_{\Delta_{k-1}}(\omega_{1:k-1}) - e^{-i\omega_k} \mathcal{F}_{\Delta_{k-1}}(\omega_1 - \omega_k, \dots, \omega_{k-1} - \omega_k) \right].$$

Proof. For any $(x_1, \dots, x_k) \in \Delta_k$, Let $Z_i \triangleq x_1 + \dots + x_i$ for $i \in [k]$. Then

$$\begin{aligned} \mathcal{F}_{\Delta_k}(\omega_{1:k}) &= \int_0^1 \int_0^{1-Z_1} \dots \int_0^{1-Z_{k-1}} k! e^{-i(\omega_1 x_1 + \dots + \omega_k x_k)} dx_1 \dots dx_k \\ &= \int_0^1 \dots \int_0^{1-Z_{k-2}} k! e^{-i(\omega_1 x_1 + \dots + \omega_{k-1} x_{k-1})} \left(\int_0^{1-Z_{k-1}} e^{-i\omega_k x_k} dx_k \right) dx_1 \dots dx_{k-1}, \end{aligned} \quad (74)$$

where a uniform probability density function over Δ_k has been assumed to be $f_{\Delta_k}(\mathbf{x}) = (k!) \mathbf{1}(\mathbf{x} \in \Delta_k)$ for all $\mathbf{x} \in \mathbb{R}^k$, due to the fact the Lebesgue measure (or volume) of the standard simplex is $1/k!$. Since we have

$$\int_0^{1-Z_{k-1}} e^{-i\omega_k x_k} dx_k = \frac{1}{i\omega_k} \left(1 - e^{-i\omega_k(1-Z_{k-1})} \right),$$

the Fourier transform can be rewritten as follows:

$$\begin{aligned} \mathcal{F}_{\Delta_k}(\omega_{1:k}) &= \frac{k}{i\omega_k} \int_0^1 \dots \int_0^{1-Z_{k-2}} (k-1)! e^{-i(\omega_1 x_1 + \dots + \omega_{k-1} x_{k-1})} dx_1 \dots dx_{k-1} \\ &\quad - \frac{k e^{-i\omega_k}}{i\omega_k} \int_0^1 \dots \int_0^{1-Z_{k-2}} (k-1)! \left(\prod_{i=1}^{k-1} e^{-i(\omega_i - \omega_k) x_i} \right) dx_1 \dots dx_{k-1} \\ &= \frac{k}{i\omega_k} \left[\mathcal{F}_{\Delta_{k-1}}(\omega_{1:k-1}) - e^{-i\omega_k} \mathcal{F}_{\Delta_{k-1}}(\omega_1 - \omega_k, \dots, \omega_{k-1} - \omega_k) \right], \end{aligned} \quad (75)$$

which completes the proof. \square

Lemma C.5 gives us a recursive procedure to produce the Fourier transform, or derive usefull properties for K -dimensional standard simplex through tools such as induction. The following lemma establishes a relation between the Fourier transform of the standard simplex and that of an arbitrary simplex in \mathbb{S}_K with a non-zero Lebesgue measure.

Lemma C.6. *For a simplex $S \in \mathbb{S}_K$ with vertices $\theta_0, \dots, \theta_K$ and a reversible zero-translated vertex matrix Θ , we have*

$$\mathcal{F}_S(\omega_{1:k}) = e^{-i\omega_{1:k}^T \theta_0} \mathcal{F}_{\Delta_K}(\Theta^T \omega_{1:k}).$$

Proof. Let $f_S : \mathbb{R}^K \rightarrow \mathbb{R}_{\geq 0}$ denote the probability density function associated to S . Then, it can be seen that

$$f_S(\mathbf{x}) = \frac{1}{\det(\Theta)} f_{\Delta_K}(\Theta^{-1}(\mathbf{x} - \theta_0)), \quad \forall \mathbf{x} \in \mathbb{R}^K.$$

In this regard, one just needs to write down the definition of Fourier transform for simplex S and utilizes the change of variables technique as follows:

$$\mathcal{F}_S(\omega_{1:k}) = \int_{\mathbb{R}^K} f_S(\mathbf{x}_{1:k}) e^{-i\omega_{1:k}^T \mathbf{x}_{1:k}} dx_1 \dots dx_k \quad (76)$$

$$\begin{aligned} &= \frac{1}{\det(\Theta)} \int_{\mathbb{R}^K} f_{\Delta_K}(\Theta^{-1}(\mathbf{x} - \theta_0)) e^{-i\omega_{1:k}^T \mathbf{x}_{1:k}} dx_1 \dots dx_k \\ &= e^{-i\omega_{1:k}^T \theta_0} \int_{\mathbb{R}^K} f_{\Delta_K}(\mathbf{u}_{1:k}) e^{-i\omega_{1:k}^T \Theta \mathbf{u}_{1:k}} du_1 \dots du_k \\ &= e^{-i\omega_{1:k}^T \theta_0} \mathcal{F}_{\Delta_K}(\Theta^T \omega_{1:k}). \end{aligned} \quad (77)$$

Therefore, the proof is complete. \square

In the following lemma, we show that the uniform measure over the standard simplex Δ_K corresponds to a low-frequency probability density function. This would be the first step toward using Corollary C.3, in order to prove Theorem C.4.

Lemma C.7 (Low-Pass Property of Δ_K). *There exists a universal constant $C > 0$, such that for $\alpha > C$ and $K \in \mathbb{N}$ the uniform probability density function over Δ_K , i.e., $f_{\Delta_K} : \mathbb{R}^K \rightarrow \mathbb{R}_{\geq 0}$, is a low-frequency function in the following sense:*

$$\frac{\text{Vol}(\Delta_K)}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} |\mathcal{F}_{\Delta_K}(\omega)|^2 \leq \mathcal{O}\left(\frac{K}{\alpha}\right), \quad (78)$$

where by $\|\omega\|_\infty$ we simply mean $\max_{i \in [K]} |\omega_i|$.

Proof. Proof is based on the direct analysis of the Fourier transform of f_{Δ_K} , which we denoted as $\mathcal{F}_{\Delta_K}(\omega)$. We take advantage of the fact that the K -dimensional unit hypercube $[0, 1]^K$ can be thought as the union of $K!$ properly rotated and translated versions of Δ_K .

Mathematically speaking, assume the ordered tuple of unit axis-aligned vectors $E = (\mathbf{1}_1, \dots, \mathbf{1}_K)$, where $\mathbf{1}_i$ for $i \in [K]$ denotes the one-hot vector over the i th component. Then, let $\mathbf{V}_1, \dots, \mathbf{V}_{K!} \in \mathbb{R}^{K \times K}$ be the set of orthonormal matrices, where each matrix transforms E into one of its $K!$ possible permutations, i.e.,

$$(\mathbf{V}_i \mathbf{1}_1, \dots, \mathbf{V}_i \mathbf{1}_K) = (\mathbf{1}_{p_{i,1}}, \dots, \mathbf{1}_{p_{i,K}}), \quad \mathbf{p}_i \in \text{Perm}([K]).$$

In this regard, we already know that there exist $K!$ corresponding vectors $\mathbf{b}_1, \dots, \mathbf{b}_{K!} \in \mathbb{R}^K$, such that the following combined probability density function

$$\frac{1}{K!} \sum_{i=1}^{K!} f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]} = f_{\square_K}, \quad (79)$$

where f_{\square_K} denotes the pdf over the K -dimensional unit hypercube. Also, it should be noted that any two distinct summands in (79) have an empty overlap. In fact, each $\mathbf{V}_i[\Delta_K + \mathbf{b}_i]$ represents a translated and rotated version of the standard simplex Δ_K .

Next, one can see that

$$\begin{aligned} \mathcal{F}\left\{\frac{1}{K!} \sum_{i=1}^{K!} f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]}\right\} &= \frac{1}{K!} \sum_{i=1}^{K!} \mathcal{F}\{f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]}\} \\ &= \frac{1}{K!} \sum_{i=1}^{K!} e^{-i\omega^T \mathbf{b}_i} \mathcal{F}_{\Delta_K}(\mathbf{V}_i^{-1} \omega), \end{aligned} \quad (80)$$

where we have used the result of Lemma C.6. Also, note that $\mathbf{V}_i^{-1} = \mathbf{V}_i$ for all i , and $\mathbf{V}_i \omega$ represents a permutation of the components of ω . As a result and due to the symmetry of the standard simplex Δ_K and also the ℓ_∞ -norm w.r.t. the ordering of the edges, we have

$$\int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}\{f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]}\}(\omega)|^2 = \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}_{\Delta_K}(\mathbf{V}_i \omega)|^2 = \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}_{\Delta_K}(\omega)|^2, \quad \forall i \in [K]. \quad (81)$$

Now, using (79) we have:

$$\begin{aligned} \int_{\|\omega\|_\infty \leq \alpha} \left| \frac{1}{K!} \sum_{i=1}^{K!} \mathcal{F}\{f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]}\} \right|^2 &= \frac{1}{(K!)^2} \sum_{i,j} \int_{\|\omega\|_\infty \leq \alpha} \left| \mathcal{F}_{\mathbf{V}_i^{-1} \Delta_K} \mathcal{F}_{\mathbf{V}_j^{-1} \Delta_K} \right| \left| e^{-i\omega^T (\mathbf{b}_i - \mathbf{b}_j)} \right| \\ &= \frac{1}{K!} \int_{\|\omega\|_\infty \leq \alpha} |\mathcal{F}_{\Delta_K}(\omega)|^2 + \left(1 - \frac{1}{K!}\right) \mathcal{O}\left(\frac{1}{\alpha}\right). \end{aligned} \quad (82)$$

The latter term in the r.h.s. of (82) corresponds to the $(K!)^2 - K!$ summands for which we have $i \neq j$. In fact, if $i \neq j$ and based on the fact that Fourier transform preserves inner product², we have

$$\int_{\mathbb{R}^K} \left| \mathcal{F}_{\mathbf{V}_i \Delta_K} \mathcal{F}_{\mathbf{V}_j \Delta_K} e^{-i\omega^T (\mathbf{b}_i - \mathbf{b}_j)} \right| = \int_{\mathbb{R}^K} f_{\mathbf{V}_i[\Delta_K + \mathbf{b}_i]} f_{\mathbf{V}_j[\Delta_K + \mathbf{b}_j]} = 0,$$

²This is a direct result of the fact that Fourier transform is an ‘‘orthonormal’’ transformation.

which holds since $\Delta_i \triangleq \mathbf{V}_i [\Delta_K + \mathbf{b}_i]$ and $\Delta_j \triangleq \mathbf{V}_j [\Delta_K + \mathbf{b}_j]$ do not overlap with each other. However, when we add the constraint $\|\boldsymbol{\omega}\|_\infty \leq \alpha$, we effectively take the integral over the hypercube $[-\alpha, \alpha]^K$ instead of the whole \mathbb{R}^K . This procedure is equivalent to the inner product of the two “smoothed” versions of f_{Δ_i} and f_{Δ_j} in the spatial domain. Here, by “smoothed” we simply mean being convolved with a K -dimensional sinc function with parameter α , i.e.,

$$\text{sinc}_{\alpha, K}(\mathbf{x}) \triangleq \prod_{i=1}^K \frac{\sin(\alpha x_i)}{\alpha x_i}.$$

Hence, we have

$$\int_{\|\boldsymbol{\omega}\|_\infty \leq \alpha} |\mathcal{F}_{\Delta_i} \mathcal{F}_{\Delta_j}| = \int_{\mathbb{R}^K} [f_{\Delta_i} * \text{sinc}_{\alpha, K}] [f_{\Delta_j} * \text{sinc}_{\alpha, K}], \quad (83)$$

which is at most $\mathcal{O}(1/\alpha)$, since only the leakages that are due to $\text{sinc}_{\alpha, K}$ overlap with each other.

On the other hand, the l.h.s. of (82) represents the integration of the Fourier transform of \square_K within $[-\alpha, \alpha]^K$. Therefore, we have

$$\begin{aligned} \frac{1}{(2\pi)^K K!} \int_{\|\boldsymbol{\omega}\|_\infty \leq \alpha} |\mathcal{F}_{\Delta_K}(\boldsymbol{\omega})|^2 &\geq \frac{1}{(2\pi)^K} \int_{\|\boldsymbol{\omega}\|_\infty \leq \alpha} |\mathcal{F}_{\square_K}(\boldsymbol{\omega})|^2 - \mathcal{O}\left(\frac{1}{\alpha}\right) \\ &= \prod_{i=1}^K \frac{1}{2\pi} \left(\int_{-\alpha}^{\alpha} \left| \int_0^1 e^{-i\omega_i x_i} \right|^2 \right) - \mathcal{O}\left(\frac{1}{\alpha}\right) \\ &= \left(1 - \frac{2}{\pi\alpha} + o(\alpha^{-1}) \right)^K - \mathcal{O}\left(\frac{1}{\alpha}\right) \\ &= 1 - \mathcal{O}\left(\frac{K}{\alpha}\right), \end{aligned} \quad (84)$$

where we have used Laurent series expansion for the integral of sinc function to derive the bound. Finally, noting the fact that we have $\text{Vol}(\Delta_K) = 1/K!$, and also the Parseval’s theorem:

$$\frac{1}{(2\pi)^K} \int_{\mathbb{R}^K} |\mathcal{F}_{\Delta_K}(\boldsymbol{\omega})|^2 = \int_{\mathbb{R}^K} f_{\Delta_K}^2(\mathbf{x}) = \frac{\text{Vol}(\Delta_K)}{\text{Vol}^2(\Delta_K)} = \text{Vol}^{-1}(\Delta_K)$$

completes the proof. \square

So far, we have managed to show that the standard simplex Δ_K is associated to a low-frequency PDF and thus would preserve a minimum level of information even after getting corrupted by additive Gaussian noise. Using Lemma C.6, one can simply extend this notion to any $(\bar{\lambda}, \underline{\lambda})$ -regular simplex. Before that, let us also extend this notion of “low-frequency” property to the difference function

$$f_{\Delta_K} - f_{\mathbf{A}_\varepsilon(\Delta_K + \mathbf{b}_\varepsilon)}, \quad (85)$$

where the linear transformation matrix \mathbf{A}_ε and translation vector \mathbf{b}_ε (for $\varepsilon > 0$) are ε -controlled perturbations that alter the standard simplex with the following magnitude

$$\mathcal{D}_{\text{TV}}(\mathbb{P}_{\Delta_K}, \mathbb{P}_{\mathbf{A}_\varepsilon(\Delta_K + \mathbf{b}_\varepsilon)}) = \varepsilon.$$

This task is more challenging since the difference PDF function in (85) can be both positive and negative, and in fact, integrates to zero over \mathbb{R}^K . However, the following lemma states that such functions for any \mathbf{A}_ε and \mathbf{b}_ε are “band-pass” in nature: even though in the Fourier domain they become exactly zero at the origin, they also become infinitesimally small in terms of magnitude as one asymptotically increases $\|\boldsymbol{\omega}\|_\infty$.

Lemma C.8 (Low-Frequency Property for Difference of Simplices). *For $\varepsilon \geq 0$, Assume the linear transformation matrix $\mathbf{A}_\varepsilon \in \mathbb{R}^{K \times K}$ where $A_{i,j} = \delta_{ij} + \mathcal{O}(\varepsilon)$, with δ_{ij} denoting the kroecker’s delta function. Also, assume the translation vector $\mathbf{b}_\varepsilon \in \mathbb{R}^K$ where $b_i = \mathcal{O}(\varepsilon)$. Let $\hat{\Delta}_{K,\varepsilon}$ denote the simplex $\mathbf{A}_\varepsilon(\Delta_K + \mathbf{b}_K)$, and assume the total variation distance between \mathbb{P}_{Δ_K} and $\mathbb{P}_{\hat{\Delta}_{K,\varepsilon}}$ is ε . Then*

$$\frac{1}{(2\pi)^K} \int_{\|\boldsymbol{\omega}\|_\infty \geq \alpha} \left\| \mathcal{F}_{\Delta_K} - \mathcal{F}_{\hat{\Delta}_{K,\varepsilon}} \right\|^2 \leq \mathcal{O}\left(\frac{K\varepsilon}{\alpha}\right),$$

for sufficiently large $\alpha > 0$.

Proof. Based on Lemmas C.7 and C.6 and for sufficiently small ε , we already know that both \mathcal{F}_{Δ_K} and $\mathcal{F}_{\hat{\Delta}_{K,\varepsilon}}$ are low-frequency functions in the following sense: their normalized ℓ_2 energy outside of the hypercube $[-\alpha, \alpha]^K$ is at most $\mathcal{O}(K/\alpha)$. We also know that

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} \left\| \mathcal{F}_{\Delta_K} - \mathcal{F}_{\hat{\Delta}_{K,\varepsilon}} \right\|^2 = \mathcal{O}(\varepsilon),$$

due to the assumption of the lemma. However, in order to prove the $\leq \mathcal{O}(K\varepsilon/\alpha)$ bound, we need to prove another important property: that the energy distribution of the difference function is not significantly controlled by ε , e.g., it does not concentrate on high-frequency regions with $\|\omega\|_\infty \geq \mathcal{O}(1/\varepsilon)$.

Before going further into the mathematical details of the above statement, let us present a better view on the Fourier transform of a simplex. We start this procedure with computing \mathcal{F}_{Δ_K} . In this regard, we show the following relation holds for $k \in \mathbb{N}$:

$$\mathcal{F}_{\Delta_k} = \frac{k!}{i^k} \sum_{\ell=1}^k (1 - e^{-i\omega_\ell}) \left[\omega_\ell \prod_{j \neq \ell} (\omega_j - \omega_\ell) \right]^{-1}. \quad (86)$$

Proof of (86) is by induction. **Base:** First, we show the equality holds for the simple case of $k = 1$. **Step:** Then, we prove that for $k \geq 2$, if the relation holds for $k - 1$, then it also holds for k .

Base For $k = 1$, it can be readily seen that a uniform measure over Δ_1 is a *unit pulse function* over the interval $[0, 1]$. Therefore, the Fourier transform of f_{Δ_1} is already known to be a sinc function, i.e.,

$$\mathcal{F}_{\Delta_1}(\omega) = \int_0^1 e^{-i\omega x} dx = \frac{1 - e^{-i\omega}}{i\omega}, \quad (87)$$

which matches the formulation of (86) if one sets $k = 1$.

Induction step Assume (86) holds for $k - 1$. By taking advantage of Lemma C.5, we have:

$$\begin{aligned} \mathcal{F}_{\Delta_k} &= \frac{k}{i\omega_k} \left(\mathcal{F}_{\Delta_{k-1}}(\omega_1, \dots, \omega_{k-1}) - e^{-i\omega_k} \mathcal{F}_{\Delta_{k-1}}(\omega_1 - \omega_k, \dots, \omega_{k-1} - \omega_k) \right) \\ &= \frac{(k-1)!}{i^{k-1}} \frac{k}{i\omega_k} \left(\sum_{\ell=1}^{k-1} (1 - e^{-i\omega_\ell}) \left[\omega_\ell \prod_{j \neq \ell}^{k-1} (\omega_j - \omega_\ell) \right]^{-1} \right. \\ &\quad \left. - e^{-i\omega_k} \sum_{\ell=1}^{k-1} (1 - e^{-i(\omega_\ell - \omega_k)}) \left[(\omega_\ell - \omega_k) \prod_{j \neq \ell}^{k-1} (\omega_j - \omega_\ell) \right]^{-1} \right) \\ &= \frac{k!}{i^k} \sum_{\ell=1}^{k-1} ((\omega_k - \omega_\ell) (1 - e^{-i\omega_\ell}) - \omega_\ell (e^{-i\omega_\ell} - e^{-i\omega_k})) \left[\omega_\ell \omega_k \prod_{j \neq \ell}^{k-1} (\omega_j - \omega_\ell) (\omega_k - \omega_\ell) \right]^{-1} \\ &= \frac{k!}{i^k} \sum_{\ell=1}^{k-1} (\omega_k (1 - e^{-i\omega_\ell}) - \omega_\ell (1 - e^{-i\omega_k})) \left[\omega_\ell \omega_k \prod_{j \neq \ell}^{k-1} (\omega_j - \omega_\ell) (\omega_k - \omega_\ell) \right]^{-1} \\ &= \frac{k!}{i^k} \sum_{\ell=1}^{k-1} (1 - e^{-i\omega_\ell}) \left[\omega_\ell \prod_{j \neq \ell}^k (\omega_j - \omega_\ell) \right]^{-1} - \frac{k!}{i^k} \left(\frac{1 - e^{-i\omega_k}}{\omega_k} \right) \sum_{\ell=1}^{k-1} \prod_{j \neq \ell}^k \frac{1}{\omega_j - \omega_\ell} \\ &= \frac{k!}{i^k} \sum_{\ell=1}^k \left(\frac{1 - e^{-i\omega_\ell}}{\omega_\ell} \right) \prod_{j \neq \ell}^k \frac{1}{\omega_j - \omega_\ell}, \end{aligned} \quad (88)$$

which again matches with the formulation of (86) and therefore completes the induction step. Here, for the last equality we have used the following mathematical identity:

$$\sum_{\ell=1}^k \prod_{j \neq \ell}^k \frac{1}{\omega_j - \omega_\ell} = 0.$$

End of induction

Based on this explicit formula, the difference function $\mathcal{F}_{\Delta_K} - \mathcal{F}_{\hat{\Delta}_{K,\varepsilon}}$ can be written as follows:

$$\mathcal{F}_{\Delta_K} - \mathcal{F}_{\hat{\Delta}_{K,\varepsilon}} = \frac{K!}{i^K} \sum_{\ell=1}^K (1 - e^{-i\omega_\ell}) \left[\omega_\ell \prod_{j \neq \ell}^K (\omega_j - \omega_\ell) \right]^{-1} (1 - r_\ell(\omega)), \quad (89)$$

where

$$\begin{aligned} r_\ell(\omega) &\triangleq e^{-i\omega^T \mathbf{b}_\varepsilon} \left(\frac{1 - e^{-i\omega^T \mathbf{a}_\ell}}{1 - e^{-i\omega_\ell}} \right) \frac{\omega_\ell}{\omega^T \mathbf{a}_\ell} \prod_{j \neq \ell} \frac{\omega_j - \omega_\ell}{\omega^T (\mathbf{a}_j - \mathbf{a}_\ell)} \\ &= e^{-i\omega^T (\mathbf{b}_\varepsilon - \delta_\ell)} \frac{\sin(\omega_\ell/2 + \delta_\ell^T \omega/2)}{\sin(\omega_\ell/2)} \frac{1}{1 + (\omega/\omega_\ell)^T \delta_\ell} \prod_{j \neq \ell} \frac{1}{1 + \frac{(\delta_j - \delta_\ell)^T \omega}{(\omega_j - \omega_\ell)}}, \end{aligned} \quad (90)$$

where \mathbf{a}_i denotes the i th row of \mathbf{A}_ε . Also, we have assumed $\mathbf{a}_i = \mathbf{1}_i + \delta_i$ with $\|\delta_i\|_\infty$, $i \in [K]$ being upper bounded by $\mathcal{O}(\varepsilon)$ according to the definition of \mathbf{A}_ε . It can be readily checked that $r_\ell(\omega) = 1 + \mathcal{O}(\varepsilon)$. Also, the multiplicative factors that form $r_\ell(\omega)$ are either not dependent on $\|\omega\|_2$ (dependence is only on the direction, and not the magnitude) or they oscillate as a function of magnitude. This implies that $r_\ell(\omega)$ which appears solely due to the difference between the two simplices does not behave differently in regions which are far from the origin or the areas in its vicinity.

Due to the above derivations, the difference function has a similar behaviour to that of $\mathcal{F}_{\Delta_K}(\omega)$ and thus we have

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} \left\| \mathcal{F}_{\Delta_K} - \mathcal{F}_{\hat{\Delta}_{K,\varepsilon}} \right\|^2 \leq \mathcal{O}\left(\frac{K\varepsilon}{\alpha}\right),$$

which completes the proof. \square

With the help of Corollary C.3 and Lemmas C.6, C.7 and C.8, we can finally prove the claimed bound. Let simplex \mathcal{S}_1 and \mathcal{S}_2 to have the vertex matrices Θ_1 and Θ_2 , respectively. Also, we have already assumed

$$\underline{\lambda} \leq \lambda_{\min}(\Theta_i) \quad , \quad \lambda_{\max}(\Theta_i) \leq \bar{\lambda} \quad \text{for } i = 1, 2.$$

Then, based on Lemma C.6 and for a sufficiently large $\alpha > 0$, we have

$$\begin{aligned} \int_{\|\omega\|_\infty \geq \alpha} |\mathcal{F}_{\mathcal{S}_i}(\omega)|^2 &= \int_{\|\omega\|_\infty \geq \alpha} \left| \mathcal{F}_{\Delta_K}(\Theta_i^T \omega) \right|^2 \\ &\leq \frac{1}{\det(\Theta)} \int_{\|\omega\|_\infty \geq \lambda\alpha} |\mathcal{F}_{\Delta_K}(\omega)|^2 \\ &\leq \frac{(2\pi)^K}{\text{Vol}(\mathcal{S}_i)} \mathcal{O}\left(\frac{K}{\lambda\alpha}\right), \end{aligned} \quad (91)$$

for $i = 1, 2$. Therefore, Lemma C.8 would consequently imply the following relation for the difference function of the two simplices:

$$\frac{1}{(2\pi)^K} \int_{\|\omega\|_\infty \geq \alpha} |\mathcal{F}_{\mathcal{S}_1} - \mathcal{F}_{\mathcal{S}_2}|^2 \leq \mathcal{O}\left(\frac{K}{\lambda\alpha}\right) \int_{\mathbb{R}^K} (f_{\mathcal{S}_1} - f_{\mathcal{S}_2})^2, \quad (92)$$

which facilitates the usage of our main general result in Theorem C.2, and more specifically Corollary C.3 which leads to the following relation:

$$\begin{aligned} \|(f_{S_1} - f_{S_2}) * G_\sigma\|_2 &\geq \|f_{S_1} - f_{S_2}\|_2 \sup_\alpha \left(\sqrt{1 - \mathcal{O}\left(\frac{K}{\lambda\alpha}\right)} e^{-K(\sigma\alpha)^2} \right) \\ &\geq \|f_{S_1} - f_{S_2}\|_2 e^{-\mathcal{O}\left(\frac{K}{\text{SNR}^2}\right)}, \end{aligned} \quad (93)$$

since $\alpha \geq \Omega(K/\bar{\lambda})$ guarantees that the l.h.s. of the bound remains positive, while gives the minimum possible exponent (at least order-wise) to the exponential term. Combining the fact ℓ_2 -norm $\leq \ell_1$ -norm together with theorem's assumptions, we have

$$\|(f_{S_1} - f_{S_2}) * G_\sigma\|_2 \leq \|(f_{S_1} - f_{S_2}) * G_\sigma\|_1 \leq 2\varepsilon.$$

Consequently, (93) implies the following upper-bound on the ℓ_2 -norm of the difference PDF function $f_{S_1} - f_{S_2}$ as follows:

$$\|f_{S_1} - f_{S_2}\|_2 \leq \varepsilon e^{\mathcal{O}\left(\frac{K}{\text{SNR}^2}\right)},$$

and completes the proof. □